


## Copyright Declaration

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.,

Fisher, G. J. (2023), 'On the Emergence of Organic Economic Institutions and the Impact of Legal Rules', University of Southampton, Faculty of Engineering and Physical Sciences, Electronics and Computer Science, Institute for Complex Systems Simulation, PhD Thesis, pp. 1-469.



University of Southampton

Faculty of Engineering and Physical Sciences  
Electronics and Computer Science  
Institute for Complex Systems Simulation

# On the Emergence of Organic Economic Institutions and the Impact of Legal Rules

by

Gregory James Fisher

(Orchid: 0009-0009-0522-5587)

A thesis submitted for the degree of Doctor of Philosophy

16 May 2023

University of Southampton

**Abstract**

Faculty of Engineering and Physical Sciences  
Electronics and Computer Science  
Institute for Complex Systems Simulation

Thesis for the degree of Doctor of Philosophy

On the Emergence of Organic Economic Institutions and the Impact of Legal Rules

Gregory James Fisher

Unplanned institutions appear to play an important role in economic activity but we lack a good understanding of how they emerge. Also, can legal rules enable the endogenous emergence of these types of institution, i.e., can they be planned? This seems paradoxical but we find that it is possible.

These issues are important because collective behaviour, of which institutions are one type, can be viewed as the other side of the coin to competitive behaviour in economies; however, far fewer research resources have been devoted to the former than the latter over the past two centuries.

The approach taken in this thesis is to apply the recent advances in Complexity Economics and Agent-Based Computational Economics to the questions of unplanned institutional emergence and the potential role of legal rules.

More specifically, agent-based models are developed that approximate the ontology of complex economic systems (the resulting gap between theory and models is much narrower than equivalent work in game theory). These models are then used to explore unplanned institutional emergence and the impact of legal rules.

The resulting simulations show that unplanned institutions can emerge from co-adaptive reinforcement learning and that habits subsequently form. However, the environment has to be sufficiently enabling of such emergence. Legal rules can be used to foster an enabling environment when institutions do not emerge endogenously but any corruption must first be overcome.

The ultimate aim of this thesis is to improve our understanding of unplanned institutional emergence (and the role of legal rules) in order to help foster solutions to real-world institutional problems.

## Declaration of Authorship

I, Gregory James Fisher, declare that this thesis entitled *On the Emergence of Organic Economic Institutions and the Impact of Legal Rules* and the work presented in it are my own and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. None of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. None of the thesis is based on work done jointly by myself with others; and
7. Parts of this work have been published as:
  - a conference paper entitled “[A Computational Model Demonstrating Market Emergence and the Division of Labour](#)” by myself and Seth Bullock, presented at the [Social Simulation Conference](#) in Rome, September 2016.
  - a conference paper entitled “[Institutional Emergence: Lessons From Two Computational Models](#)” by myself and Seth Bullock, at the [Annual Conference of the World Interdisciplinary Network of Institutional Research \(WINIR\)](#) in Hong Kong, September 2018.

Gregory James Fisher

16 May 2023

## Acknowledgements

Without doubt, most of my thanks go to my long-suffering partner, Pooneh Ghoddoosi, who has supported me in many different ways throughout the challenging processes of research and thesis writing.

I am also indebted to my (only) supervisor, Seth Bullock, especially for keeping me as a student when he left Southampton for Bristol. And my thanks also go to Enrico Gerding as my ‘Admin Supervisor’ at Southampton: he has put up with more than we expected!

My gratitude also goes to the directors of funding at the EPSRC who chose to invest in a PhD programme in complex systems simulation, which enabled this thesis. More specifically, the work presented here was supported by the EPSRC Doctoral Training Centre grant EP/G03690X/1. My thanks also go to Seth (again) and Jonathan Essex, both of whom created and ran the programme.

Special mention should also go to my friend, Rhett Gayle, who spent a considerable amount of time reading through earlier drafts of work and providing constructive advice. A sarcastic, pedantic philosopher makes for an ideal feedback mechanism.

I would also like to thank the examiners of this thesis, Professor Geoffrey Hodgson and Professor Richard Watson, for their patience and expertise in evaluating the work.

Finally, the orientation of this thesis is consistent with the ‘social construction of self’ so I would be remiss if I did not thank the (other) very many people I have learnt from over the years, and with whom I have discussed various aspects of this thesis. They are too numerous to mention here but they know who they are.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Motivation	2
1.1.1	The Importance of Economic Institutions	2
1.1.2	Challenges in the Literature	4
1.1.3	Social Simulation Literature: Emergence and Immergence	9
1.2	Research Questions	11
1.2.1	Complexity and Computational Modelling	13
1.3	Executive Summary	15
1.4	Assumptions and Definitions	20
1.4.1	Institutions	21
1.4.2	Organic Institutions	22
1.4.3	Order	23
1.4.4	Legal Rules	24
1.4.5	Habits	25
1.4.6	Reinforcement Learning	25
1.4.7	Property Rights	26
1.4.8	Downward and Upward Causation / Effects	27
1.4.9	Types of Interaction Challenges	27
1.5	Added Value	28
<b>2</b>	<b>Complexity Science and Complexity Economics</b>	<b>31</b>
2.1	The Complexity Sciences: Defining Features	32
2.1.1	Multiple Heterogeneous Agents	33
2.1.1.1	Nested Complex Systems	33
2.1.1.2	Semi-Permeable Boundaries	34
2.1.1.3	Internal Models	34
2.1.1.4	Decentralised Control / Bottom-Up Focus	35
2.1.2	Interdependence	35
2.1.3	Time	36
2.2	The Complexity Sciences: Conceptual Features	37
2.2.1	Feedback	37
2.2.2	Non-Linearity	37
2.2.3	Space of Possibilities	38
2.2.4	Far-From-Equilibrium	38
2.2.5	Emergence	39
2.2.6	Self-Organisation	40
2.2.7	Lock-In	41
2.2.8	State of Paradox	43
2.2.9	Sensitivity to Initial Conditions / Predictability	44

2.2.10	Symmetry Breaking	44
2.3	Complexity Economics	45
2.3.1	Should Economies be Viewed as Complex Systems?	45
2.3.2	Dealing with Mapping and Other Challenges	46
2.3.3	The Eleven Principles of Complexity Economics	46
2.3.3.1	Economies are computational in nature	47
2.3.3.2	Economies exhibit non-equilibrium and dis-equilibrium features	49
2.3.3.3	Economies include the processes of <i>formation</i> and <i>allocation</i>	50
2.3.3.4	Economies have a stratified ontology	51
2.3.3.5	Complexity Economics is open to ‘generalized Darwinism’	54
2.3.3.6	Economies are non-ergodic	56
2.3.3.7	Uncertainty is an important feature in economic behaviour	57
2.3.3.8	Agents use mental models to make decisions	58
2.3.3.9	Complexity Economics is open to inter-disciplinarity	60
2.3.3.10	Complexity Economics values realism and is sympathetic to instrumentalism	61
2.3.3.11	Pragmatism	67
2.3.4	Neoclassical Economics as Restricted Complexity Economics	68
2.3.5	Choice of Framework in Economics	69
2.3.5.1	Rationale for Using Complexity Economics to Study Organic Institutions	71
2.3.6	Internal Consistencies: Complexity Economics and Neoclassical Theory	71
2.4	Criticisms of The Complexity Sciences & Complexity Economics	73
2.5	Chapter Summary	75
<b>3</b>	<b>Spontaneous Order - Rationale and History</b>	<b>77</b>
3.1	Introduction and Rationale	77
3.2	A Short History of Spontaneous Order	80
3.2.1	Pre-Scottish Enlightenment	82
3.2.2	The Scottish Enlightenment	84
3.2.3	Post-Scottish Enlightenment	85
3.3	Historical Influences on Hayek	89
<b>4</b>	<b>Hayek’s Spontaneous Order</b>	<b>93</b>
4.1	Interpreting Hayek	94
4.2	Hayek’s Ontology	96
4.2.1	Hayek’s Theory of Mind	96
4.2.2	Hayek’s Knowledge Problem	97
4.2.3	Instincts, Habits and Rules	98
4.2.4	Hayek on Complexity	99
4.3	Invisible Hand Explanations	105
4.3.1	Eligibility as an Invisible Hand Explanation	106
4.3.2	Characteristics of Invisible Hand Explanations	107
4.3.3	The Aggregate and Functional-Evolutionary Moulds	109
4.3.4	Links to Computational Models	110
4.4	Catallaxy and the ‘Market Order’	111
4.4.1	Criticisms of Hayek’s ‘Market Order’	113



4.5	Hayek's Theory of Cultural Evolution . . . . .	115
4.5.1	Variation . . . . .	116
4.5.2	Selection . . . . .	117
4.5.3	Durability . . . . .	119
4.5.4	Hayek's Explanation of Cultural Evolution . . . . .	120
4.5.5	Order and Efficiency . . . . .	121
4.5.6	Mapping from Individuals to Groups . . . . .	122
4.6	Hayek on Common Law and Legislation . . . . .	123
4.6.1	Hayek's Ahistoric View of Common Law . . . . .	124
4.6.2	Inconsistency in Hayek's Writings . . . . .	125
4.7	Vanberg's Reconciliation of Hayek Regarding Legislation . . . . .	126
4.7.1	Links to Computational Models . . . . .	130
4.8	Summary . . . . .	130
<b>5</b>	<b>Models of Organic Institutional Emergence</b>	<b>133</b>
5.1	Hodgson and Knudsen's Traffic Convention Model . . . . .	134
5.1.1	The Traffic Conventions Model . . . . .	135
5.1.2	Simulation Results . . . . .	138
5.1.3	Discussion and Critique . . . . .	140
5.1.4	Commonality of Rules within the Group . . . . .	143
5.2	Social Norms Within the Social Simulation Literature . . . . .	143
5.2.1	Norm Diffusion . . . . .	145
5.2.2	Multi-Agent Systems & Artificial Intelligence Models . . . . .	146
5.2.3	The Emergence of Social Norms . . . . .	146
5.2.3.1	Two Simple Coordination Models . . . . .	147
5.2.3.2	EMergence In the Loop (EMIL) . . . . .	148
5.3	Game Theoretic Approaches . . . . .	154
5.3.1	Institutions via Games . . . . .	155
5.3.2	The Rules-as-Equilibria Approach . . . . .	158
5.3.3	Field's Criticisms . . . . .	162
5.3.4	The Value of Game Theory Models of Institutions . . . . .	164
5.3.5	Implications for Modelling Organic Institutional Emergence . . . . .	167
5.4	Models of Monetary Emergence . . . . .	168
5.4.1	Kiyotaki & Wright's Model of Monetary Emergence . . . . .	169
5.4.2	Marimon et al's Computational Model . . . . .	173
5.4.3	Empirical Studies of Monetary Emergence . . . . .	177
5.4.4	Comparing the Literature . . . . .	179
5.5	Concluding Comments . . . . .	181
5.5.1	Reasoning, Learning and Habits . . . . .	181
5.5.2	Beyond First-Order Approximations . . . . .	182
5.5.3	Thoroughness of Analysis . . . . .	182
<b>6</b>	<b>Models: Rationale, Design and Results</b>	<b>185</b>
6.1	Research Questions . . . . .	186
6.2	Model Design . . . . .	190
6.2.1	Risks with ABMs . . . . .	193
6.3	Model Description . . . . .	195
6.4	Summary of Results . . . . .	198
6.4.1	Main Results . . . . .	198
6.4.2	Parameter Space Exploration and Enabling Environments . . . . .	201

6.4.3	Habituation . . . . .	202
6.4.4	The Division of Labour and Knowledge . . . . .	206
6.4.5	Liberal Legislation Experiments . . . . .	208
6.4.6	Invisible Hand Explanations: Characteristics . . . . .	209
6.4.7	Emergence and Immersion . . . . .	211
<b>7</b>	<b>Market Emergence Model</b>	<b>213</b>
7.1	Overview of the Model . . . . .	214
7.1.1	The Environment and Time . . . . .	214
7.1.2	The Agents . . . . .	215
7.1.3	Foraging . . . . .	218
7.1.4	Trading . . . . .	219
7.1.5	End of Round Adjustments and Communication . . . . .	222
7.1.5.1	Consumption & Metabolism . . . . .	222
7.1.5.2	Updating Foraging Skills . . . . .	222
7.1.5.3	Updating Foraging Strategies . . . . .	223
7.1.5.4	Communication . . . . .	225
7.2	Results: Null Scenario - No Trading . . . . .	225
7.3	Results: Null Scenario - No Memories with Specialisation . . . . .	226
7.4	Results: Null Scenario - Memories without Specialisation . . . . .	228
7.5	Results: Default Scenario - Memories with Specialisation . . . . .	230
7.5.1	The Emergence of the Market . . . . .	232
7.5.2	The Emergence of Specialization . . . . .	235
7.5.3	Proximity to the General Equilibrium Outcome . . . . .	239
7.6	Habituation . . . . .	240
7.6.1	Reinforcement Learning and Habituation . . . . .	241
7.6.1.1	Efficiency and Equality . . . . .	243
7.6.1.2	Agent Level Data: Reinforcement Learning and Habituation . . . . .	244
7.6.2	Reinforcement Learning Turned Off . . . . .	245
7.7	Discussion of Results . . . . .	246
7.7.1	Organic (Unplanned) Institutions . . . . .	246
7.7.2	Organic Institutions with Habituation? . . . . .	248
<b>8</b>	<b>Market Emergence Model: Exploring the Parameter Space</b>	<b>251</b>
8.1	Summary of the Results . . . . .	251
<b>9</b>	<b>Property Rights Model: Introduction</b>	<b>255</b>
9.1	Propensities to Steal and Defend . . . . .	256
9.2	Adjustments to Agent Interaction . . . . .	258
9.3	Similarity with Conventional Games . . . . .	262
9.4	The Evolution of the Agents' Propensities to Steal and Defend . . . . .	263
9.4.1	The Reduced Gain / Loss From Interactions and Adjustments . . . . .	263
9.4.2	Mapping $v_i$ and $v_j$ on to Propensities . . . . .	264
9.4.2.1	Expected Versus Absolute Pay-offs . . . . .	265
9.4.2.2	Cognitive Coarseness . . . . .	266
9.4.2.3	Rate of Change Coefficient . . . . .	267
9.4.2.4	Example . . . . .	267
9.4.3	Propensities, Learning and Probability Limits . . . . .	269
9.5	Moving Around the Grid: Interaction and Avoidance . . . . .	269
9.6	End of Round Communications & Grid Target Selection . . . . .	272

9.7	Changes to the Original Parameters	273
<b>10</b>	<b>Property Rights Model: Simulations and Discussion</b>	<b>275</b>
10.1	Property Rights Model: Null Simulations	275
10.1.1	Null Scenario: Replicating the Original Model	276
10.1.2	Null Scenario: All Agents Steal	277
10.1.3	Null Scenario: Substantive Rationality Approach	279
10.1.3.1	Resource Concentration Effect	281
10.1.4	Null Scenario: Fixing Agents' Propensities at 0.5	283
10.2	Property Rights Model: Default Parameter Set	285
10.2.1	Propensities to Defend	285
10.2.2	Propensities to Steal	287
10.2.3	Fights and Transactions on the Grid	291
10.2.4	Specialisation & Child Birth	293
10.3	Habituation Experiments	294
10.3.1	Reinforcement Learning 'Switched Off'	296
10.3.2	Reinforcement Learning and Habituation	301
10.4	Discussion of Results	306
10.4.1	Organic (Unplanned) Institutions	307
10.4.2	Organic Institutions with Habituation?	309
10.4.3	Generalized Darwinism?	310
<b>11</b>	<b>Property Rights Model: An Exploration of the Parameter Space</b>	<b>313</b>
11.1	Summary of Results	314
<b>12</b>	<b>Property Rights Model: Liberal Legislation</b>	<b>317</b>
12.1	Yellow Agents	320
12.1.1	Legal Rules without Corruption (Yellow Agents)	320
12.1.2	Legal Rules with Corruption (Yellow Agents)	322
12.1.2.1	The Emergence of Successful and Unsuccessful Economies	325
12.2	Low Cost of Fighting	327
12.2.1	Legal Rules without Corruption (Low Cost of Fighting)	328
12.2.2	Legal Rules with Corruption (Low Cost of Fighting)	331
12.3	Power from Fighting Skill	332
12.3.1	Legal Rules without Corruption (Power from Fighting Skill)	333
12.3.2	Legal Rules with Corruption (Power from Fighting Skill)	337
12.4	Power from Accumulated Resources: A Wealthy Agent	337
12.4.1	Legal Rules without Corruption (Power from Resources)	337
12.4.2	Legal Rules with Corruption (Power from Resources)	339
12.5	Conclusion	340
<b>13</b>	<b>Conclusion</b>	<b>343</b>
13.1	First Research Question: Symmetry Breaking via Immergence and Emergence	344
13.1.1	A Generalised Framework of Organic Institutional Emergence	344
13.1.1.1	An Emerging Property	347
13.1.1.2	Downward Effects	347
13.1.1.3	Upward Effects	348
13.1.1.4	Positive Feedback ('Outer Emergence')	348
13.1.1.5	Reconstitution of Mental Models	349
13.1.1.6	Positive Feedback (Immergence)	349

13.1.1.7	Enabling Environment . . . . .	350
13.1.1.8	Final Results . . . . .	350
13.1.1.9	Mutually Supporting Emergence and Immergence . . . . .	350
13.1.2	Types of Downward Effects . . . . .	351
13.1.3	First Model . . . . .	352
13.1.4	Second Model . . . . .	353
13.1.5	Commonality . . . . .	355
13.2	Second Research Question: Liberal Legislation . . . . .	356
13.3	Pre-Existing Institutions & Stratified Ontologies . . . . .	357
13.4	Research Involving Agent-Based Models . . . . .	358
13.5	Future Research . . . . .	359
13.5.1	Simplification . . . . .	359
13.5.2	Bottom-Up Development . . . . .	360
13.5.3	Development of the Existing Model . . . . .	360
13.5.4	Empirical Evidence . . . . .	360
13.5.5	Expansion . . . . .	361
13.5.6	Guiding Principles & Other Domains . . . . .	361
<b>A</b>	<b>Interaction Typologies</b> . . . . .	<b>395</b>
A.1	Schotter's Typology . . . . .	395
A.2	Schultz's Typology . . . . .	397
<b>B</b>	<b>Exploration of the Parameter Space: First Model</b> . . . . .	<b>401</b>
B.1	Memory Decay . . . . .	401
B.1.1	Low Values of $m_{dec}$ . . . . .	401
B.1.2	Intermediate Values of $m_{dec}$ . . . . .	402
B.1.3	High Values of $m_{dec}$ . . . . .	402
B.2	Population Density . . . . .	406
B.3	Communication . . . . .	406
B.4	Trading on the Way to the Target Location . . . . .	408
B.5	Speed of Foraging Skill Change . . . . .	408
B.6	Randomizing Agent Home Locations . . . . .	410
B.7	Expectations Accuracy When Adjusting Foraging Strategies . . . . .	411
B.8	The Weight of Other Agents' Transactions . . . . .	413
B.9	Travel Distance . . . . .	413
B.10	Selection of Target Location: Winner Takes All . . . . .	415
B.11	The Agents' Initial Resource Endowment . . . . .	417
B.12	The Fountains' Initial Endowments . . . . .	419
B.13	Geographically Locating Resource Fountains . . . . .	420
B.14	Agents with Limited Life Spans . . . . .	422
B.15	Skill Acquisition: A Linear Approach . . . . .	425
<b>C</b>	<b>Exploration of the Parameter Space: Second Model</b> . . . . .	<b>427</b>
C.1	Adjusting Agents' Starting Propensities . . . . .	427
C.2	Adjusting the Cost of Fighting . . . . .	428
C.2.1	Fight Cost of Zero . . . . .	429
C.2.2	Fight Cost of 2.0 . . . . .	431
C.2.3	Intermediate Fight Costs . . . . .	431
C.3	Adjusting the Influence of Other Agents' Experience . . . . .	432
C.4	Adjusting the Errors Made in Expected Pay-Offs . . . . .	433
C.4.1	Standard Deviation of Zero . . . . .	433

C.4.2	High Standard Deviation	434
C.4.3	Intermediate Sizes of the Standard Deviation	436
C.5	Adjusting the Standard Deviation of Children's Starting Propensities	436
C.6	Adjusting the Agents' Initial Resource Endowments	438
C.6.1	The Two Types of Simulation	439
C.6.2	Starting Resources of Below and Equal to 200 Units	439
C.6.3	Higher Values of Starting Resources	440
C.7	Changing the Nature of Agent Learning	442
C.7.1	Rate of Change Coefficient ( $r$ )	443
C.7.1.1	Low Values of $r$ ( $\leq 0.004$ )	443
C.7.1.2	Higher Values of $r$ ( $> 0.004$ )	444
C.7.2	Cognitive Coarseness ( $\delta$ )	444
C.7.2.1	Coarse-Grained Cognition ( $\delta = 0$ )	444
C.7.2.2	Fine-Grained Cognition ( $\delta = 1$ )	445
C.7.3	Relative and Absolute Approaches	446
C.7.3.1	Using Absolute Reduced Values When Adjusting the Propensities to Defend	447
C.7.3.2	Using Relative Reduced Values When Adjusting the Propensities to Steal	449
C.8	Adjusting the Reputations Architecture	450
C.8.1	Strangers: Use of Substantive Rationality	450
C.8.2	Strangers: Assuming the Worst	451
C.8.3	Adjusting the Memory Length of the Agents' Reputations Arrays	451
C.8.3.1	No Memories	452
C.8.3.2	Short Memory Lengths (1 - 4 Rounds)	452
C.9	Changing the Agents' Target Location Weights	452
C.10	Reducing Agent Clustering & Resource Concentration	453
C.10.1	Enhanced Avoidance of Previous Fight Locations	454
C.10.2	Allowing Agents to Go Home	455
C.10.3	Enhanced Avoidance of Previous Fight Locations and Allowing Agents to Go Home	455
<b>D</b>	<b>Property Rights Model: Experimentation</b>	<b>459</b>
D.1	Yellow Agents	460
D.2	Power from Fighting Skill	461
D.2.1	Simulations without Black Sheep	462
D.2.2	Simulations with Black Sheep	463
D.3	Power from Personal Resources	465
D.3.1	Equal Starting Resources	466
D.3.2	One Wealthy Agent	467



# List of Figures

5.1	Taken from Hodgson and Knudsen (2004), p. 29. The caption in the original paper reads “Fig. 2. Degrees of convergence with 200 runs for each level of habit and error $w_X = \{1.4, 0.9, 0.7, w_{Habit}\}$ .” The vertical axis shows Convergence (0.5 means the cars drive equally on the left and right side and 1 means all the cars drive on the same side); the y-axis (error) corresponds to $\epsilon$ above; and the x-axis corresponds to Habituation $_{n,t}$ above. Reproduced with permission from Elsevier B. V. . . . .	138
5.2	Taken from Hodgson and Knudsen (2004), p. 31. The caption in the original paper reads “Fig. 3. The habit effect in parameter space.” $C_0$ is convergence when $w_{Habit} = 0$ and $C_1$ is convergence when $w_{Habit} = 1$ . The difference between these two ( $C_1 - C_0$ ) is termed the <i>habit effect</i> . Reproduced with permission from Elsevier B. V. . . . .	139
5.3	A summary of Kiyotaki and Wright’s (1989) production and consumption environment (Model A). A group of agents is divided into 3 types. Each arrow relates to a specific resource: the squares at the base of each arrow correspond to production and the pointy ends correspond to its consumption by each agent type, i.e., Type 1 agents consume Resource 1 and produce Resource 2; Type 2 agents consume Resource 2 and produce Resource 3; and Type 3 agents consume Resource 3 and produce Resource 1. The design of the problem space ensures there is no double coincidence of wants vis-à-vis any two agents’ production and consumption resources.	170
5.4	Kiyotaki and Wright’s (1989) Fundamental Equilibrium. Types 1 and 3 only transact with Type 2 who essentially intermediates between them. Resource 1 is viewed as a form of money in this pattern. . . . .	171
5.5	Kiyotaki and Wright’s (1989) Speculative Equilibrium. This is identical to the Fundamental Equilibrium outcome except this time Type 1 agents will accept Resource 3 from either Type 2 or 3 agents. Both resource 1 and 3 are viewed as money-like: Resource 1 is used as ‘money’ by Type 2 agents and Resource 3 is used by Type 1 agents. . . . .	172
6.1	A breakdown of each round in the first model. In the first stage of each round, agents forage from resource fountains. In the second stage they try to find each other within a geographic region, in order to trade any resources they might have. At the end of the round, various adjustments are made and some of the agents communicate. . . . .	196

- 6.2 A time series of the number of agents targetting different locations on the grid during the first 50 rounds of a typical simulation that uses the default parameter set. This chart demonstrates a scenario when two markets dominate for a period but symmetry breaking means that, eventually, one dominates. . . . . 199
- 6.3 The ‘cloud’ of agents’ propensities to steal over 200 rounds in a typical simulation when the default parameter set was used. The chart shows the propensities to steal of the agents who survived until the end of the simulation as blue ‘dots’. The blue line is the mean of these values. The red dots represent the propensities to steal of the agents who died before the end of the simulation, and the red line is the mean of these values (it ends in Round 102 when the last of these agents died). The chart shows how, on the whole, the propensities to steal of the agents who died was sustained above those who survived. . . . . 200
- 6.4 A time series of the contributions to a target location’s weight in memory over the first 200 rounds of a typical simulation (for one agent). The two ‘areas’ denote contributions to the total weight over time: the blue area shows the contribution of reinforcement learning and the red area shows that of habituation. This chart shows how the agent ‘learned’ to visit the location depicted, i.e., reinforcement learning dominated during the emergence of the market. . . . . 204
- 6.5 The same time series as shown in Fig. 6.4 above but over 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the weight of this location over time. . . . 205
- 6.6 A time series of the agents’ propensities to steal over the first 400 rounds of a typical simulation when the agents’ propensities change as a result of reinforcement learning and habituation (here,  $ha_2 = 0.01$ ). The red lines depict ‘Al Capone’ agents, whose propensities increase above 1, and who all died. The blue lines depict passive-aggressive agents (who all survived). 206
- 6.7 A time series of a passive-aggressive agent’s propensity to steal (the black line) over the first 100 rounds of a typical simulation when propensities change as a result of reinforcement learning and habituation (here  $ha_2 = 0.005$ ). The y-axis measures the level of the agent’s propensity to steal and the contributions to this propensity. The four ‘areas’ denote contributions to the agent’s propensity to steal over time: the upper (dark red) area shows positive contributions by habituation; the second (dark blue) area shows positive contributions by reinforcement learning; the third (light blue) represents negative contributions by reinforcement learning; and the lower (light red) area represents the negative contribution by habituation. This chart shows how the agent’s propensity to steal declines below zero (this occurred in Round 26) mainly as a result of reinforcement learning. . 207
- 6.8 The same time series as shown in Fig. 6.7 above but for 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the agent’s propensity to steal over time. . 207
- 7.1 A breakdown of each round in the first model. In the first stage of each round, agents forage from resource fountains. In the second stage they try to find each other within a geographic region, in order to trade any resources they might have. At the end of the round, various adjustments are made and some of the agents communicate. . . . . 215



7.2	The location of agents' homes on the grid. In the default simulations the grid is a torus with dimensions of $50 \times 50$ . Each of the agents starts the trading phase from its home location. These 'homes' are shown above as black squares. If a new agent is born, its home location on the grid is chosen in such a way as to maximize the sparsity of the population (this is in order to ensure consistency between simulations). . . . .	218
7.3	Total agent population over 2,000 rounds when the agents are prevented from trading. Over 20 simulations the total population declines to approximately 15-16 agents, which can be viewed as the carrying capacity of the environment when the agents neither trade nor specialise. . . . .	226
7.4	Heatmap of transactions in each grid square (last 100 rounds of a typical simulation) when the agents hold no previous transactions in memory. The agents move around the grid in a random walk and interact only when they bump in to each other. This results in a 'noisy' dispersion of transactions. . . . .	227
7.5	Supply and demand curves and total transactions in one round of a typical simulation when the agents have no memories of previous transactions. The supply and demand curves are calculated after foraging and before trading, and are generated by discovering what each agent would have supplied or demanded for each resource over a range of prices; and then aggregating these over all agents. The red star indicates the total volume (measured on the horizontal axis) and the mean price (measured on the vertical axis) of the subsequent transactions in the same round. The mean ratio of actual transactions to the market clearing volume was 18% over 20 simulations. . . . .	227
7.6	Heatmap of transactions in each grid square (last 100 rounds only) when the agents use their memories of previous transactions to find other agents. Over each simulation they converge on a single grid location in order to trade (a 'market square'). Note, however that this location differed between simulations: it was an emergent property of each simulation. . . . .	229
7.7	Supply and demand curves and total transactions in one round of a typical simulation when the agents use their memories of previous transactions to find other agents. See Fig. 7.5 above for a description of the data. On average the ratio of actual transactions to the market clearing volume was 100.0% over 20 simulations. . . . .	229
7.8	A time series of the turnover ratio over 500 rounds when agents use their memories of previous transactions to find other agents. The turnover ratio is defined as actual transactions divided by the market clearing volume. It was measured and recorded in each round. In this chart we can see how the turnover ratio increased toward 1 as agents find each other and trade. The ratio is maintained at approximately 1 after a single market emerges. . . . .	230

- 7.9 Time series of the mean specialisation value in a typical simulation using the default parameter set. For each agent, the ‘specialisation value’ is the number of times an agent forages from the fountain it visits most often (given 2 fountains and 5 foraging opportunities, this value varies between 3 and 5). For the agents as a whole, an average value of 3 would mean all the agents were generalists and a value of 5 would mean they were all specialists (foraging from the same resource fountain in all five time slots during the foraging phase). The data here shows how the agents became more specialised over the first 200 rounds of the simulation, which correlates with the turnover ratio. . . . . 231
- 7.10 A time series of the number of agents targetting six locations on the grid during the first 50 rounds of a typical simulation that used the default parameter set. The chart shows how multiple markets emerged during the first 30 rounds but one of these ([12, 8]) was visited by more agents. Symmetry breaking meant all the agents eventually converged on this location. . . . . 232
- 7.11 A time series of the number of agents targetting different locations on the grid during the first 50 rounds of a typical simulation that uses the default parameter set. This chart demonstrates a scenario when two markets dominate for a period but symmetry breaking meant that, eventually, one dominates. Reproduced from Fig. 6.2 in Chapter 6 for convenience. . . . . 234
- 7.12 The specialisation threshold for an agent over 200 rounds of a typical simulation. The model keeps track of this metric (the red line) for each agent over each round: if the turnover ratio experienced by the agent and its neighbours (the black line) exceeds this value then agents tend to specialise, and vice versa. The blue line shows the foraging skill (as a probability of detection) of the resource in which the agent became specialised, and the green line shows the non-specialised resource. . . . . 237
- 7.13 Supply and demand curves and total transactions in Round 999 of a typical simulation when  $ha = \infty$ . See Fig. 7.5 above for a description of the data. On average the ratio of actual transactions to the market clearing volume during rounds 900 - 999 was 0.98 over 20 simulations. . . . . 242
- 7.14 A time series of the contributions to a target location’s weight in memory over the first 200 rounds of a typical simulation (for one agent). The two ‘areas’ denote contributions to the total weight over time: the dark blue area shows the contribution of reinforcement learning; and the dark red area shows that of habituation. This chart shows how the agent learned to visit the location depicted as a result of reinforcement learning. Reproduced from Chapter 6 for convenience. . . . . 245
- 7.15 The same time series as shown in Fig. 7.14 above but over 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the weight of this location over time. Reproduced from Chapter 6 for convenience. . . . . 246
- 9.1 Overview of agent interaction in the second model. Agents initially decide whether to steal or trade. If one of them wishes to steal and the other trade then this second agent has to decide whether to defend its resource or not. As a result there are six potential scenarios for each interaction. . . . . 259

9.2	Probabilities of each of the six potential scenarios in Fig. 9.1. These are joint probabilities derived from both of the agents' probabilities of stealing and defending. . . . .	262
9.3	A diagram depicting how an agent evaluates interacting or moving on the grid in the second model. The $3 \times 3$ grid shown here is a subset of the $50 \times 50$ torus on which agents could move during the interaction phase. The agent (in red) initially evaluates the expected gains and losses from potential interactions on its current square (in blue). If there are no agents on that square or none of the interactions are deemed beneficial, the agent evaluates the gains and losses on the other squares within sight (the grey squares). . . . .	270
10.1	Heatmap of transactions in the last 100 rounds of a typical simulation when the agents' propensities to steal are fixed at 0. The results replicate those of the default simulations of the first model except, here, markets are made up of small areas of adjacent squares. . . . .	277
10.2	Total agent population over 5,000 of an extended (and typical) simulation when the agents' propensities to steal are fixed at 0. The population plateaus at approximately 43 agents, as in the original default simulations.	277
10.3	Total agent population over 2,000 rounds of a typical simulation when the agents' propensities to steal are fixed at 1, i.e., the agents only ever steal. The population collapses in all 20 simulations owing to debilitating fight costs. . . . .	278
10.4	Agent interactions by scenario (5-round moving averages) when decisions are made via substantive rationality. Note that '1' refers to transactions; '2F' refers to interactions when the instigating agent attempts to trade but defends its resources when the counterpart tries to steal; '2A' is the same as 2F but the instigating agent acquiesces; '3F' is when the instigating agent attempts to steal and its counterpart fights back; '3A' is the same as 3F but the counterpart acquiesces; and '4' is when the agents both attempt to steal from each other. Substantive rationality leads agents to attempt theft from each other most of the time. . . . .	280
10.5	Total agent population when decisions were made via substantive rationality. The data are taken from the same simulation depicted in Fig. 10.4. The fact that most agents stole meant the population collapsed. . . . .	281
10.6	Agents' resource holdings at the start, middle, and end of the interaction phase of Round 30 of a typical simulation when agents made decisions via substantive rationality. In these charts, the agents' resource holdings were put in ranked order at the beginning of the interaction phase (the first time period), in the middle (the 25th time period) and at the end (the 50th time period). The charts show how the agents' resource holdings became more concentrated over this phase of the round. . . . .	282
10.7	Time series of the gini coefficient of agents' resources holdings over the first 200 rounds of a typical simulation when all the agents' propensities were fixed at 0.5. In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. The chart shows that resources became more concentrated in approximately the first 100 rounds (when the dark blue line was higher than the light blue line) but not thereafter (when most agents had died and those remaining generally avoided each other). . . . .	284

- 10.8 Fan chart of the agents' propensities to defend over the first 250 rounds of a typical simulation when the default parameters were used. The black line represents the mean and each red band represents one standard deviation (skew-adjusted) away from the mean. Note the data shows how the agents' propensities were allowed to increase above 1 and to decline below 0 (the probability an agent defended its resources was restricted to between 0 and 1, which are limits shown by two green dashed lines). In this simulation the mean propensity to defend increased above 1 in Round 28. . . . . 286
- 10.9 Time series of cumulative contributions to the agents' propensities to defend by scenario over 250 rounds in the simulation depicted in Fig. 10.8. The model records changes in the agents' propensities to defend by interaction (scenarios 2F, 2A, 3F, and 3A). These were aggregated in each round by scenario and the cumulative time series for one of these simulations is shown here. The chart shows that agents mostly learned to defend their resources in scenarios 2A and 3A. . . . . 286
- 10.10 Fan chart of the agents' propensities to steal over the same simulation as that depicted in figures 10.8 and 10.9. The black line represents the mean and each blue band represents one standard deviation (skew-adjusted) away from the mean. Note the data uses the agents' propensities, which were allowed to increase above 1 and to decline below 0 (the probability an agent attempted to steal was restricted to between 0 and 1, which are shown by two green dashed lines). The chart shows how the mean propensity to steal initially increased but then declined to below 0. . . . 287
- 10.11 Time series of contributions to the agents' propensities to steal by scenario over 250 rounds in the simulation depicted in figures 10.8 - 10.10. The model records changes in the agents' propensities to steal by interaction. These were aggregated in each round by scenario and the time series for one of these simulations is shown here. The chart shows that agents mostly learned to steal in scenarios 3A and 3F but they learned to trade in scenarios 2F and 4. Ultimately, the agents learned to respect each other's property. . . . . 288
- 10.12 The 'cloud' of agents' propensities to steal over 200 rounds in a typical simulation when the default parameter set was used (copied from Chapter 6 for convenience). The chart shows the propensities to steal of the agents who survived until the end of the simulation as blue 'dots'. The blue line is the mean of these values. The red dots represent the propensities to steal of the agents who died before the end of the simulation, and the red line is the mean of these values (it ends in Round 102 when the last of these agents died). The chart shows how, on the whole, the propensities to steal of the agents who died was sustained, initially, above those who survived. This figure is replicated from Chapter 6 for convenience. . . . 290

10.13	Time series of the gini coefficient of agents' resources holdings (means over 20 simulations). In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. Mean values were also recorded over all 20 simulations in a set, which are shown here. The chart shows that resources became more concentrated in approximately the first 200 rounds (when the dark blue line was higher than the light blue line) but the gap between the two declined until approximately Round 250. The decline in resource concentration this demonstrated was a result of the agents' propensities to steal declining. . . . .	291
10.14	Heatmap of fight locations in the first 100 rounds of a typical simulation when the default parameter set is used. The chart shows how agents congregated on the same area and fought. . . . .	292
10.15	Heatmap of transaction locations in the first 100 rounds of a typical simulation when the default parameter set is used. The chart shows how agents congregated in the same location shown in Fig. 10.14 above, i.e., where fights took place. . . . .	292
10.16	Heatmap of transaction locations in the last 100 rounds of the same simulation depicted in figures 10.14 and 10.15. The chart shows how agents congregated in the same location shown in those charts. However, agents only ever traded at this stage of the simulation. . . . .	293
10.17	Time series of the turnover ratio in the first 250 rounds of the same simulation depicted in figures 10.14 - 10.16. The ratio increased progressively toward 1 as the Al Capone agents died off. . . . .	294
10.18	Time series of the total agent population from the same simulation depicted in figures 10.14 - 10.17 above. The population initially declined to 15 agents but property rights emerged among these agents who met at the same market, specialised and then bore children. . . . .	294
10.19	A time series of the agents' propensities to steal over the first 400 rounds of a typical simulation when $ha_2 = 0.0001$ and when there was no reinforcement learning. Here, the agents' propensities changed too slowly for property rights to emerge. . . . .	297
10.20	A time series of the agents' propensities to steal over the first 400 rounds of a typical simulation when the agents' propensities changed as a result of reinforcement learning and habituation (here, $ha_2 = 0.005$ ). The red lines depict 'hawks', whose propensities increased above 1, and who all died. The blue lines depict 'doves' and passive-aggressive agents (who all survived). . . . .	298
10.21	A time series of the agents' total resources over 1,000 rounds (taken from the same simulation as that depicted in Fig. 10.20 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows that all the 'hawks' died; all the passive-aggressive agents (the lower group of lines above zero) survived; and the 'doves' (the higher group of lines above zero) also survived. . . . .	299
10.22	A time series of the agents' propensities to steal of an atypical simulation when the agents' propensities changed as a result of reinforcement learning and habituation (here, $ha_2 = 0.005$ ). The red lines depict 'hawks', whose propensities increased above 1, all of whom died with one exception. The blue lines depict doves passive-aggressive agents. This pattern is observed when a relatively large number of 'doves' are seen relative to passive-aggressive agents. The surviving hawk thrives in this environment. . . . .	299

10.23	A time series of the agents' total resources over 1,000 rounds (taken from the same simulation as that depicted in Fig. 10.22 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows that all the 'hawks' but one died - this surviving hawk was able to bully the doves (after all the other hawks died) allowing its resource stock to increase significantly. . . . .	300
10.24	The 'cloud' of the agents' propensities to steal over 1,000 rounds of a simulation when the agents' propensities change as a result of reinforcement learning but not habituation, i.e., $ha_2 = 0$ . Each blue 'dot' represents the propensity to steal of a living agent in each round. In this simulation, the agents' propensities fell below zero after Round 550 (approximately): property rights emerged. . . . .	301
10.25	The 'cloud' of the agents' propensities to steal over 1,000 rounds under the same conditions as those in Fig. 10.24 above. In this simulation, the agents' propensities did not fall below zero before the end of the simulation: property rights did not emerge. . . . .	302
10.26	A time series of the agents' propensities to steal over the first 400 rounds of a typical simulation when the agents' propensities change as a result of reinforcement learning and habituation (here, $ha_2 = 0.01$ ). The red lines depict 'hawks', whose propensities increase above 1, and who all died. The blue lines depict passive-aggressive agents (who all survived). In these simulations there were no doves. . . . .	304
10.27	A time series of the agents' total resources over 1,000 rounds (taken from the same simulation as that depicted in Fig. 10.26 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows that all the 'hawks' died and all the passive-aggressive agents survived. . . . .	304
10.28	A time series of a passive-aggressive agent's propensity to steal (the black line) over the first 100 rounds of a typical simulation when propensities change as a result of reinforcement learning and habituation (here $ha_2 = 0.005$ ). The four 'areas' denote contributions to the agent's propensity to steal over time: the upper (dark red) area shows positive contributions by habituation; the second (dark blue) area shows positive contributions by reinforcement learning; the third (light blue) represents negative contributions by reinforcement learning; and the lower (light red) area represents the negative contribution by habituation. This chart shows how the agent's propensity to steal declines below zero (this occurred in Round 26) mainly as a result of reinforcement learning. . . . .	306
10.29	The same time series as shown in Fig. 10.28 above but for 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the agent's propensity to steal over time. . . . .	307
10.30	A time series of a hawk's propensity to steal (the black line) over the 235 rounds it was alive. The data were taken from a typical simulation when propensities changed as a result of reinforcement learning and habituation (here $ha_2 = 0.005$ ). See Fig. 10.28 above for a description of the data. The chart shows that the agent's propensity to steal increased above 1 because the net impact of habituation exceeded that of reinforcement learning. . . . .	307

12.1	A diagram summarising which agents incur fight costs and fines when the legal rules do not include compensation payments to victims. . . . .	318
12.2	A diagram summarising which agents incur fight costs and fines, and which receive compensation, when the legal rules include compensation payments to victims. . . . .	319
12.3	Fan chart of living agents' propensities to steal when a fine of 0.5 was applied to stealing agents in a 'Yellow Agents' experiment. The black line represents the mean and each band shows one standard deviation (skew-adjusted) away from the mean. This chart shows the fine did not work in that agents did not come to respect each other's property. . . . .	321
12.4	The cloud of agents' propensities to steal over the first 500 rounds of a typical simulation in a 'Yellow Agents' experiment when 0.4 resource units was paid as a fine and compensation. Each dot represents a living agent's propensity to steal in each round. This chart shows that property rights emerged among the surviving agents. . . . .	321
12.5	A 3-dimensional surface showing the proportion of 'successful' simulations (the vertical axis) in 'Yellow Agents' experiments when the legal rule included fines and compensation of 0.6 resource units. The left-hand horizontal axis, labelled 'Lambda' ( $\lambda$ ), was the bribe stated as a proportion of the fine; and the right-hand horizontal axis, labelled 'Prop. Steal' ( $\bar{p}^s$ ), was the mean starting propensity to steal of the initial agents. Here, success meant the agents population did not collapse and property rights emerged in the surviving population. We can see that success increased monotonically in both $\lambda$ and $\bar{p}^s$ when $\lambda < 0.75$ and $\bar{p}^s > 0.25$ . . . . .	324
12.6	The cloud of agents' propensities to steal over the first 500 rounds of an 'unsuccessful' simulation in a 'Yellow Agents' experiment when fines and compensation of 0.6 resource units were applied and there was corruption ( $\lambda = 0.25$ ). Each dot represents a living agent's propensity to steal in each round and the black line is the median of these propensities. Property rights did not emerge in this simulation. . . . .	326
12.7	The cloud of agents' propensities to steal over the first 500 rounds of a 'successful' simulation in a 'Yellow Agents' experiment using the same parameters as Fig. 12.6. Each dot represents a living agent's propensity to steal in each round and the black line is the median of these propensities. Despite the same parameters as Fig. 12.6 being used, here property rights did emerge among the surviving agents. . . . .	326
12.8	Fan chart of living agents' propensities to steal over a typical simulation when the cost of fighting was 0.05 resource units. The blue line represents the mean and each colour band represents one standard deviation (skew-adjusted) away from this mean. Here, the cost of fighting was known to be below the threshold required for property rights to emerge, hence they did not. . . . .	328
12.9	The cloud of agents' propensities to steal over the first 500 rounds of a typical simulation when the cost of fighting was low and a fine of 0.1 resource units was applied without compensation to victims. There was no corruption. Each dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. The chart shows the fine was sufficient to encourage property rights to emerge. . . . .	329

- 12.10 The cloud of agents' propensities to steal over the first 500 rounds of a typical simulation when the cost of fighting was low and a fine of 3 resource units was applied along with compensation to victims. There was no corruption. Each dot represents a living agent's propensity to steal in each round. The chart shows property rights emerged swiftly when this legal rule was applied. . . . . 330
- 12.11 The cloud of agents' propensities to steal over the first 500 rounds of a typical simulation when the cost of fighting ( $c$ ) was 0.05 and a fine of 0.1 resource units was applied along with compensation to victims. Here, the harshest corruption conditions possible were applied at the beginning of the simulation. Each dot represents a living agent's propensity to steal in each round and the black line is the median of these propensities. The chart shows property rights emerged despite extreme corruption early in the simulation. This corruption was eliminated after the median propensity to steal declined to below zero in Round 113. . . . . 332
- 12.12 The cloud of agents' propensities to steal over a typical simulation when the outcome of fights was determined by relative fighting skills, black sheep were allowed, and a fine of 1 resource unit was applied. Each dot represents a living agent's propensity to steal in each round. The chart shows five black sheep were born but they all died. In addition, the agent population was resilient to these black sheep: their propensities to steal remained negative and they survived. . . . . 333
- 12.13 Time series of the agents' fighting skills in the simulation depicted in Fig. 12.12 above. The chart shows how the five black sheep became skilled at fighting; however, the combination of fines and fight costs meant they all died. . . . . 334
- 12.14 The cloud of agents' propensities to steal over a typical simulation when the outcome of fights was determined by relative fighting skills, black sheep were allowed, and a legal rule applied fines and compensation of 0.5 resource units. Each dot represents a living agent's propensity to steal in each round. In this simulation, seven black sheep were born: four of them were 'reformed' by the legal rule (their propensities to steal declined to below zero) but three of them died (those born with propensities above 2.3). . . . . 335
- 12.15 Time series of the agents' fighting skills in the simulation depicted in Fig. 12.14 above. The chart shows how the black sheep again became skilled at fighting. However, the four 'reformed' sheep saw their skills decline when their propensities to steal fell below zero. . . . . 336
- 12.16 The cloud of the agents' propensities to steal over the first 500 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 8,000 units of each resources) was included in the population. A fine of 0.4 resource units was applied to transgressors alongside compensation to victims. Each blue dot represents a (non-wealthy) living agent's propensity to steal in each round and the red line represents the wealthy agent. In this simulation the legal rule encouraged the wealthy agent to adopt a Passive Aggressive strategy, which meant its propensity to steal became negative quite quickly. Ultimately, property rights emerged across the whole population. 338



12.17	The cloud of the agents' propensities to steal over the first 500 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 8,000 units of each resources) was included in the population. A fine of 0.5 resource units was applied alongside compensation to victims and the simulation was started with the harshest corruption conditions possible. Each blue dot represents a (non-wealthy) living agent's propensity to steal in each round and the red line represents the wealthy agent. The black line is the agents' median propensity to steal. In this simulation the propensities to steal of the non-wealthy agents were driven lower by fight costs but, initially, the wealthy agent became an Al Capone agent. However, as the efficacy of the legal rule improved, this agent paid more fines and fewer bribes, which ultimately led it to respect property rights. . . . .	340
13.1	An illustration of the generalised framework of organic institutional emergence. This framework contains two 'learning loops': the first corresponds to the emergence of a property in the environment; and the second corresponds to the reconstitution of the agents' mental models. See the main text for a detailed description of the seven features this framework contains.	346
B.1	A time series of the total number of agents visiting specific target locations when the memory decay rate was zero and two markets of equal size emerged. These markets were maintained because symmetry breaking was weakened by the fact that all transactions were sustained in the agents' memories. . . . .	402
B.2	A time series of the total number of agents visiting specific target locations when the memory decay rate was 0.9501. No market was sustained until Round 326 when 3 agents met at location [25, 8] and transacted (10 times). This market was sustained after it was reported to several agents who met and transacted, which ensured the market was maintained despite the agents' poor memories. . . . .	404
B.3	A time series of the total number of agents visiting specific target locations when the memory decay rate was 0.9751. No markets were sustained until Round 929 when 3 agents met at location [32, 32] and transacted. This location was not, however, sustained because the agents had such poor memories that they eventually forgot about it. . . . .	405
B.4	Time series of the turnover ratio over 500 rounds when the agents did not communicate. The turnover ratio increased towards 1 more slowly than in Fig. 7.8: agents' convergence on markets was slowed down when they were prevented from communicating. . . . .	407
B.5	A time series of agent specialisation. This chart depicts an agent's skilled (blue line) and non-skilled (green line) resource detection probabilities, the probability threshold above which the agent chose to specialise (the red line), and the expected probability of transacting (the black line). The data is shown over the first 500 rounds of a typical simulation that used the default parameter set. The increase in the blue line to 1 depicts specialisation. . . . .	409
B.6	A time series of agent specialisation when the speed of adjustment of detection probabilities was set very low, at 0.0025. See Fig. B.5 above for a description of the data, which here is shown over 1,000 rounds. The chart shows how specialisation was slowed down in these simulations. . . .	410

B.7	Agent home locations after 10,000 rounds when the agents did not store transaction information and when their homes were allocated randomly. Agents walked around the grid randomly and tended to survive if their home locations were in close proximity to others. This is a type of agglomeration effect. . . . .	411
B.8	A time series of agent specialisation when the agents made inaccurate estimations of future foraging yields (here, $IQ_i = 0.4$ ). See Fig. B.5 above for a description of the data, which is shown in this figure over 1,000 rounds. The chart shows how specialisation was slowed down relative to the default simulations but the agent nonetheless specialised. . . . .	412
B.9	Heatmap of transactions in the last 100 rounds of a typical simulation when the agents' travel distance was limited to 10 grid squares. The chart shows that several 'local markets' emerged when the agents could not traverse the whole grid. . . . .	416
B.10	Market catchment areas in the last 100 rounds of the same simulation depicted in Fig. B.9. Each green square shows locations within 10 grid squares of the market; and the 'diamonds' represent agent homes. We can see that all but two agents lived within at least one catchment area. . . . .	416
B.11	The total agent population, births and deaths when agents were born with 5 units of each resource over 1,000 rounds of a typical simulation. The total population initially declined to below the carrying capacity, which enabled 11 agents to survive. Agents eventually specialised and the increased productivity meant they bore children. . . . .	418
B.12	Agents' marginal rates of substitution during Round 20 of a typical simulation when resource fountains were located on the grid and agents started the trading phase from the last fountain they visited. The chart shows time series for each agent during the trading phase: their marginal rates of substitution converged on two different values, corresponding to transactions at each fountain location (agents remained at these locations so there was no arbitrage). . . . .	421
B.13	Agents' marginal rates of substitution during Round 139 of the same simulation depicted in Fig. B.12. Resource fountains were located on the grid and agents started the trading phase from the last fountain they visited. The chart shows time series for each agent during the trading phase: here, the agents were not fully specialised but their marginal rates of substitution converged on a single value because all agents targetted the same location on the grid. This market was located at one of the two fountain sites. . . . .	422
B.14	Agents' marginal rates of substitution during Round 999 of the same simulation depicted in figures B.12 and B.13. Resource fountains were located on the grid and agents started the trading phase from the last fountain they visited. The chart shows time series for each agent during the trading phase: the agents were now fully specialised and their marginal rates of substitution converged because the single market location enabled arbitrage. Specialisation meant none of the agents could trade until they all met at the market square. . . . .	423

- B.15 Total agent population over 1,000 rounds in a typical simulation when the agents' lifespan was limited to 300 rounds. In this simulation the total population fell to just 3 agents but these lived long enough to specialise and bear children. Their children also specialised and bore children, etc., which allowed the total population to recover. . . . . 425
- C.1 The cloud of agents' propensities to steal over the first 150 rounds of a typical simulation when the agents' propensities were determined by a uniform distribution at the start. Each dot represents a living agent's propensity in each round and the blue line is the mean of these propensities. The propensities evolved as they did in the default simulations, approximately speaking. . . . . 428
- C.2 The cloud of agents' propensities to steal over the first 250 rounds of a typical simulation when the agents' propensities to steal and defend started relatively low (0.1 on average). Each dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. In this chart, the propensities to steal increased aggressively to approximately 0.5 - 1.0 before declining to below zero. . . . . 429
- C.3 The cloud of agents' propensities to steal over the first 1,000 rounds of a typical simulation when the cost of fighting was 0. Each dot represents a living agent's propensity to steal in each round. The chart shows how, in general, the agents' propensities to steal increased to and above 1: with no cost associated with fighting, agents did not learn to respect others' property. . . . . 430
- C.4 Agents' resource holdings at the start, middle, and end of the interaction phase of Round 800 of the same simulation depicted in Fig. C.3 above. In these charts, the agents' resource holdings were put in ranked order at the beginning of the trading phase (the first time period), in the middle (the 25th time period) and at the end (the 50th time period). The charts show how the agents' resource holdings became more concentrated over the round. . . . . 431
- C.5 Time series of the total agent population in a typical (extended) simulation when the cost of fighting was high. The agents ignored each other: the lack of transactions and specialisation meant the total population declined to 15 agents, which was the carrying capacity of the environment (in the absence of specialisation). . . . . 432
- C.6 The cloud of agents' propensities to steal over the first 175 rounds of a typical simulation when the agents accounted for the experiences of other agents as much as their own when learning ( $\beta = 1$ ). Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. The chart shows how more Al Capone agents emerged in these simulations but they eventually died: these are the four blue dotted lines that can be seen drifting higher in the chart. The surviving agents' propensities to steal declined to below zero as in the default simulations. . . . . 433

- C.7 The cloud of the agents' propensities to defend over the first 300 rounds in a typical simulation when the agents made very large errors in estimating their pay-offs in interactions. These propensities remained at approximately 0.5 on average over the whole simulation: in effect the agents interacted randomly. These simulations showed that for property rights to emerge it was necessary for agents to discern between other agents. . . . 434
- C.8 The cloud of the agents' propensities to steal over the first 1,000 rounds of a typical simulation when 'black sheep' children could be born (these were new agents instantiated with positive propensities to steal). Each blue dot represents a living agent's propensity to steal in each round. In this simulation, six of such children were born but they all died: the population proved resilient. . . . . 437
- C.9 The cloud of the agents' propensities to steal over the first 1,400 rounds of a typical simulation when agents were instantiated with a mean of 1,000 units of each resource. Each blue dot represents a living agent's propensity to steal in each round. It took approximately 1,200 rounds for all the agents' propensities to decline to below zero. . . . . 441
- C.10 A time series of the agents' total resources over 2,000 rounds (taken from the same simulation as that depicted in Fig. C.9 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows how all but 5 of the agents died but when property rights emerged, these surviving agents thrived. . . . . 443
- C.11 The cloud of the agents' propensities to steal over the first 80 rounds of a typical simulation when agents' cognition was 'coarse-grained' ( $\delta = 0$ ). This meant agents attached +/-1 values to beneficial and detrimental interactions, respectively. Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. All of the agents survived in this simulation and they came to respect property rights much more quickly than in the default simulations. 445
- C.12 The cloud of the agents' propensities to steal over the first 160 rounds of a typical simulation when agents' cognition was 'fine-grained' ( $\delta = 1$ ). Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. In this chart the agents' propensities declined more reluctantly than when their cognition was 'coarse-grained'. Nonetheless, property rights eventually emerged in this simulation. . . . . 446
- C.13 The cloud of the agents' propensities to defend over the first 100 rounds in a typical simulation when the agents used absolute gains / losses to adjust these propensities. Each red dot represents a living agent's propensity to defend in each round and the red line is the mean of these propensities. The agents' propensities increased as in the default simulations but they tended to remain below 1. . . . . 447
- C.14 The cloud of the agents' propensities to steal over the first 100 rounds of the same simulation depicted in Fig. C.13. Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. Here, because defence of property did not emerge, the agents' propensities to steal did not decline as they did in the default simulations. . . . . 448

- C.15 The cloud of the agents' propensities to steal over the first 100 rounds of a typical situation when these propensities were adjusted by outcomes relative to expectations. Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. Property rights did not emerge in this simulation. In fact, agents learned it was generally preferable to steal than to trade. . . . . 450
- C.16 A heatmap of fight locations in first 100 rounds of a typical simulation when agents used an enhanced target selection technique. This allowed them to better avoid locations they associated with their resources being stolen. Fights were much more dispersed in this simulation than in the same period in the default simulations (see Fig. 10.14 as a contrast). . . . 454
- C.17 Time series of the gini coefficient of agents' resources holdings (mean of 20 simulations) over the first 600 rounds when the agents removed themselves from the grid if they had gained more than 2 resource units during the interaction phase. In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. This chart shows that resources became more concentrated in approximately the first 250 rounds (when the dark blue line was higher than the light blue line) but less so than in the default simulations. . . . . 456
- C.18 Time series of the gini coefficient of agents' resources holdings (mean of 20 simulations) over the first 600 rounds when the agents removed themselves from the grid if they had gained more than 2 resource units during the interaction phase and when they used an enhanced target selection technique. In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. This chart shows that resources became more concentrated in approximately the first 400 rounds (when the dark blue line was higher than the light blue line) but less so than in the default simulations and in Fig. C.17 above. . . . . 457
- D.1 Time series of living agents' propensities to steal (mean of 20 simulations) over the first 200 rounds when the agents never defended their resources. The chart shows that agents learned it was generally preferable to steal rather than to trade. . . . . 461
- D.2 A time series of contributions to the agents' propensities to steal by scenario over the first 200 rounds of a typical simulation within the 'Yellow Agents' experiment. The chart shows agents learned it preferable to steal when the instigating agent attempted to steal and the counterpart acquiesced (Scenario 3A); but they learned it was preferable to trade when both agents tried to steal from each other (Scenario 4). The net effect was positive so the agents' propensities to steal increased. . . . . 461
- D.3 Time series of living agents' propensities to steal over the first 500 rounds of a typical simulation when the agents had fighting skills, the change coefficient used ( $r^f$ ) was 0.01, and the agent population did not collapse. Each blue dot represents a living agent's propensity to steal in each round. In this particular simulation the surviving agents came to respect each other's property. . . . . 463

- D.4 A time series of the agents' fighting skills over the first 500 rounds of the simulation depicted in Fig. D.3 above. Each line depicts the fighting skill of a single agent. The chart shows how all the agents' fighting skills increased and that the Al Capone agents with higher fighting skills (and propensities to steal) all died. The surviving agents' fighting skills declined toward zero as their propensities to steal also fell. . . . . 464
- D.5 Time series of living agents' propensities to steal over a typical simulation when the agents had fighting skills and 'black sheep' were born. Each blue dot represents a living agent's propensity to steal in each round. Here, a group of the initial agents came to respect each other's property; however, when the first 'black sheep' was born, it bullied the other agents until they all died. . . . . 465
- D.6 A time series of the agents' fighting skills over the simulation depicted in Fig. D.5 above. The fighting skills of the initial agents declined toward zero after property rights emerged. When the first 'black sheep' was born, its fighting skill increased more rapidly than the other agents because it bullied all of them, successfully, to the point of population collapse. . . . 465
- D.7 The cloud of the agents' propensities to steal over the first 250 rounds of a typical simulation when fight outcomes were determined by relative resource holdings. Each blue dot represents a living agent's propensity to steal in each round. Approximately speaking, these simulations replicated the results of the default simulations; however, the passive aggressive strategy was more successful in these simulations and, as a result, more of the initial 25 agents survived and specialised. . . . . 467
- D.8 The cloud of the agents' propensities to steal over the first 100 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 8,000 units of each resource) was included in the population. Each blue dot represents a (non-wealthy) living agent's propensity to steal in each round. The red line represents the wealthy agent, which adopted an Al Capone strategy in this simulation. This meant it bullied the other agents and the population collapsed. . . . . 468
- D.9 The cloud of the agents' propensities to steal over the first 1,000 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 4,000 units of each resources) was included in the population. Each blue dot represents a (non-wealthy) living agent's propensity to steal in each round. The red line represents the wealthy agent. Here, this wealthy agent adopted a passive aggressive strategy, which meant its propensity to steal became negative quite quickly. Ultimately, property rights emerged across the whole population. . . . . 469

# List of Tables

5.1	A Summary of Duffy & Ochs' Experiment Results (Model A). Two sets of results are presented: one where $u_i = 20$ (Table 4, p. 859) and another where $u_i = 100$ (Table 7, p. 865). Other parameters are identical for both: $\beta = 0.9$ , $c_1 = 1$ , $c_2 = 4$ , and $c_3 = 9$ ). The first experiment corresponds to Kiyotaki and Wright's (1989) Fundamental Equilibrium and the second to their Speculative Equilibrium. Values of 1 in this table mean that Type $i$ agents would always exchange Resource $(i + 1) \bmod 3$ for $(i + 2) \bmod 3$ (0 means they would not). The results referred to as "KW Exp" are the values expected in Kiyotaki and Wright's (1989) Nash equilibria for each agent type. Those referred to as "D&O" refer to the results observed in Duffy and Ochs's (1999) experiments (second half of results only): "D&O(p)" refers to these results when the subjects start each game holding their production goods, and "D&O(e)" when they start with the good expected in Kiyotaki and Wright's (1989) Nash equilibria. There are two main differences: (i) Type 1 agents do not behave as expected in Kiyotaki and Wright's (1989) Nash equilibria (in both experiments) - notably, their behaviour does not appear to change between the two experiments; and (ii) Type 3 agents do not behave as expected when $u_i = 100$ .	178
5.2	A Summary of Duffy's (2001) Experiment Results. The parameters are identical to Duffy and Ochs's (1999) Speculative Equilibrium experiments (RHS of Table 5.1 above). The results referred to as "KW Exp" are the values expected in Kiyotaki and Wright's (1989) Nash equilibria for each agent type. Those referred to as "D&O" refer to the results observed in Duffy and Ochs's (1999) experiments (stated in Table 5.1 above). "D Sims" refers to Duffy's (2001) main simulation results.	179
7.1	Proximity to general equilibrium over the last 100 rounds of 20 default simulations. If the default simulations replicated the general equilibrium model then the turnover ratio would be 1, the mean price difference (mean actual prices divided by the market clearing price) would be zero, and the mean daily standard deviation of price changes would also be zero. These data indicate that the results of the default simulations did indeed approximate the results of the General Equilibrium model.	239

7.2	A summary of results for ‘habituation’ experiments based on the first model. Data are mean values for the last 100 rounds of 20 simulations, for each value of $ha$ (standard deviations in parentheses). $ha$ refers to the habituation parameter added to the weight of an agent’s target location (if there is one), which does not decay over time (note in the default simulations, $ha = 0$ , which is represented in the top row). ‘No. markets’ refers to the mean number of market locations. ‘% on one square’ refers to the percentage of transactions on the square with the most transactions. The turnover ratio is the volume of actual transactions divided by the would-be volume at the market clearing price. The final row (‘No RL’) refers to a set of experiments when reinforcement learning is ‘switched off’: here, agents only ever visit the location where they first transacted during a random walk. . . . .	241
9.1	Adjustments to agents’ propensities to steal and defend following an interaction. . . . .	264
9.2	An example of changes in an instigating agent’s propensities to steal and defend when both agents hold 1 unit of each resource in their baskets. The data is shown for every one of the six scenarios. . . . .	268
10.1	Main results for simulations when there was no reinforcement learning, for selected values of $ha_2$ . The middle column shows the total population of agents at the end of 1,000 rounds, and the right hand column shows the percentage of simulations in which property rights emerged (defined as when the propensities to steal of all the surviving agents at the end of the simulation were negative). . . . .	297
10.2	Main results of simulations in which the agents’ propensities changed via reinforcement learning and habituation, for selected values of $ha_2$ (shown in the left hand column). Data correspond to 20 simulations for each value of $ha_2$ (standard deviations in parentheses). ‘Pop’n’ refers to the mean total population of agents at the end of 1,000 rounds; ‘% Prop Rights’ indicates the percentage of 20 simulations in which property rights emerged (this is when all the agents alive at the end of Round 1,000 had negative propensities to steal); ‘Max Prop Steal’ is the mean propensity to steal at which the agents’ propensities peaked in the simulations; ‘No. Doves’ is the mean number of doves alive at the end of 1,000 rounds; ‘No. PA’ refers to the mean number of passive-aggressive agents alive at the end of 1,000 rounds; and ‘No. Hawks’ shows the same data for the number of hawks alive. . . . .	303
A.1	A Pure Coordination Game: “Matching Pennies” . . . . .	396
A.2	The Prisoners’ Dilemma . . . . .	396
A.3	A Non-Pure Coordination Game: “Battle of the Sexes” . . . . .	396
A.4	A Game with No Nash Equilibrium But Coinciding Preferences: “Around the Houses” . . . . .	397



B.1	Results of simulations when the agents' travelling distance on the grid was adjusted (last 100 rounds only). The data show that as the distance was increased, agents congregated at fewer markets, the turnover ratio increased, more agents were within the catchment area of at least one market, and the agents became more specialised. Notes: 1. Mean max spec value is a scale from 3 to 5 where 3 means the agents were generalists and 5 means they were specialists. 2. Mean max detection probabilities measures the highest of each agent's two detection probabilities and then takes the mean of these for all agents. 3. Squares with transactions measures the total number of squares on the grid in which agents traded. 4. The turnover ratio is the actual net transactions divided by the market clearing volume of trade. 5. Serviced locations refers to the average percentage of agent homes locations that were within at least one market's catchment area, i.e., when the market was within the travel distance from an agent's home. . . . .	415
B.2	Adjusting life spans. The table shows the mean population in the last 100 of 1,000 rounds and the percentage of simulations in which the agent population did not collapse. Unsurprisingly, as the agent's life span was increased, the mean population increased and the population collapsed fewer times. . . . .	424
C.1	Adjusting the agents' starting resources and the resulting simulation types. The left hand column corresponds to the mean starting resources for each agent upon instantiation. The number of type 1 simulations (when the agent population collapsed) is shown in the middle column, and the number of type 2 simulations is shown in the right hand column (when the population survived). The data show that as the mean starting resources was increased, the agent population collapsed fewer time. . . . .	440
C.2	An example of changes to a counterpart agent's propensity to defend using (i) relative outcomes ( $v_j - E_j^{Q3}(v_j)$ ); and (ii) absolute outcomes ( $v_j$ ), when the agent defended its resources (scenarios 3F, win or lose), and when it acquiesced (scenario 3A). The initiating agent ( $i$ ) had no resources and the counterpart ( $j$ ) had 10 units in total, and we assumed $P_i^{FB} = 0.5$ . The most significant difference between the two was that $j$ 's propensity to defend <i>increased</i> when it won the fight and the relative outcome was employed but it <i>decreased</i> when the absolute outcome was used. . . . .	448



# List of Abbreviations

ABM	Agent-Based Model
IE	Institutional Economics
MRS	Marginal Rate of Substitution (the ratio of an agent's holding of one resource divided by its holding of the other)
NCE	Neoclassical Economics
OED	Oxford English Dictionary



# Chapter 1

## Introduction

Ford! There's an infinite number of monkeys outside who want to talk to us about this script for Hamlet they've worked out.

– *The Hitchhiker's Guide to the Galaxy* by Douglas Adams

This thesis is focused on identifying the mechanisms by which ‘organic’ (or ‘unplanned’) economic institutions might emerge, and the conditions under which this emergence might occur.

The definition of ‘institutions’ adopted in this thesis is that of “durable systems of established and embedded social rules that structure social interactions.” (Hodgson, 2006a, p. 13). This definition is unpacked further below. Moreover, the category of organic institutions, following Menger [1890] (1981), refers to institutions that emerge without any intentional design<sup>1</sup>.

The thesis also considers the role of legal rules (defined in Section 1.4 below) as potential enablers of institutional emergence.

These matters are considered in the context of the literature focused on ‘spontaneous order’. The meaning of this term that is adopted by the thesis follows Adam Ferguson who refers to order that is “the result of human action, but not the execution of any human design” (Ferguson, 1767, p. 205)<sup>2</sup>.

---

<sup>1</sup>The types of organic institutions most commonly referred to in the literature are conventions and social norms.

<sup>2</sup>The term ‘spontaneous order’ is attributed to Michael Polanyi who writes that when “order is achieved among human beings by allowing them to interact with each other on their own initiative ... we have a system of spontaneous order in society.” (Polanyi, 1951, p. 159).

The rest of this chapter is divided into five sections:

1. Motivation for the research;
2. Two research questions;
3. An executive summary of the thesis;
4. Some definitions and assumptions; and
5. A summary of the added value of this thesis to the literature.

## 1.1 Motivation

The main arguments of this section are that: (i) organic institutions have been and continue to be important for economic activity; and (ii) our understanding of the process of organic institutional emergence is under-researched. These two combine to provide powerful motivations for this thesis, which focuses on the mechanisms of institutional emergence and the conditions under which this occurs.

In this section, we first look at the importance of institutions in economics (Section 1.1.1). We then consider challenges in the literature (Section 1.1.2). The third subsection (1.1.3) looks at relevant material from the social simulation literature, notably *cognitive emergence* and *immersion*, both of which are defined below.

### 1.1.1 The Importance of Economic Institutions

There is now a substantial body of research showing that institutions are fundamentally important in economic activity. One prominent (though contentious) example is [North and Thomas \(1973\)](#) who state “the factors we have listed (innovation, economies of scale, education, capital accumulation, etc.) are not causes of growth; they *are* growth” (p. 2, emphasis included). North and Thomas go on to argue that factor accumulations and innovation are *proximate* causes of growth and that differences in economic performance are due mainly to institutions.

The literature that supports this point can be broadly divided into empirical and theoretical. Both are now voluminous. Douglass North in particular published a considerable amount of empirical work that places institutions at the heart of economic behaviour<sup>3</sup>. At the same time, there has been a substantial amount of work concerned with the theory of institutions, notably since the mid-1970s<sup>4</sup>.

---

<sup>3</sup>See [Hodgson \(2017\)](#) for a summary of North’s work.

<sup>4</sup>It is worth noting that Institutional Economics originally appeared at the turn of the 20th Century, including the work of Thorstein Veblen, John R. Commons, Wesley Mitchell, and Clarence Ayres. The field appeared to decline (but not disappear) after the 1930s before re-emerging in the 1970s.

Related to this, a search of articles in the Web of Science database<sup>5</sup> using the phrase ‘Institutional Economics’ (IE) indicated that 262 related articles were published in 1980, 503 in 1990, 1,084 in 2000, 3,347 in 2010, and 9,634 in 2020. Related to this, the Journal of Institutional Economics was first published in 2005; and the World Interdisciplinary Network for Institutional Research (WINIR) was launched in 2012.

Furthermore, the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel has been awarded to several researchers focused on institutional research, including Ronald Coase (1991), Douglass North (1993), Robert Fogel (1993), Elinor Ostrom (2009), and Oliver Williamson (2009).

The importance of institutions in economics can perhaps be most easily understood by referring to property rights and markets, possibly the two most important institutions in economies based on free markets<sup>6</sup>. These institutions appear to have been crucial for the economic success of Western countries over the past two centuries.

If we look more specifically at organic institutions, which include conventions and social norms, we can appreciate that these have been pervasive in human activities for possibly millions of years. For example, hunting parties typically include rules that coordinate groups in order to maximize the likelihood of success<sup>7</sup>.

In an equivalent way, modern corporations typically include workers who adopt a wide variety of conventions and norms.

### **Present Day Challenges**

As an economy changes, it throws up new institutional challenges. Arguably, an important role of IE is to help inform or enable new institutions that will meet these challenges, for the betterment of human welfare. This is not only about policy prescriptions for the state, it can also mean advice for corporations and civil society groups.

One of the features of Complexity Economics (CE) is the recognition that economies change non-ergodically (North, 2005). This is discussed further in Chapter 2 but here we can summarise this feature as recognising that economies re-pattern themselves over time in an open-ended process of change.

Non-ergodic change can give rise to new collective action challenges (below we describe two types that are described in Schultz, 2001). In some scenarios it might be helpful to catalyse an enabling environment which helps organic institutions emerge. In others, some type of legal rule (which is known to people and known to be enforced) might

---

<sup>5</sup>Dated 6 February 2023.

<sup>6</sup>One could argue that both are enabled by language so perhaps this is more important.

<sup>7</sup>The oldest archaeological evidence of Homo Sapiens and their ancestors hunting dates to approximately 2 million years ago.

be beneficial. Having a firm grip on how institutions emerge and evolve is essential for providing useful advice in these areas.

A good example of a present-day institutional challenge concerns the use of cryptocurrencies. This is a particularly relevant example given the important role played by the institution of money in modern economies and because money is a recurring theme in this thesis. Also, money is viewed as one of the ‘paradigm cases’ of spontaneous order.

To date, many cryptocurrencies have been created by people in civil society as an alternative to traditional currencies that are issued (or backed) by the state<sup>8</sup>. They are typically based on blockchain - or equivalent - technologies.

These new currencies have thrown up a number of challenges. Most notably, there have been cases of fraud (such as individuals selling cryptocurrencies<sup>9</sup> that are immediately worthless), theft via hacking, and price volatility (because of the lack of an effective peg to the real economy)<sup>10</sup>.

How we deal with these challenges depends on how we frame the subject (the ‘quality problem’ of money, which is discussed in more detail in Chapter 3, is helpful here). If we align cryptocurrency challenges with Schultz’s (2001) first type of scenario (coordination situations), then the problem is one of coordinating behaviour and we would expect any outcome to be self-sustaining.

However, if the ‘problem space’ appears more aligned with Schultz’s (2001) second type of scenario (collective action situations) then we recognise that these currency systems provide opportunities for people to cheat others. How we proceed will depend on the nature of this cheating.

There has been an enormous amount of progress made in our understanding of the nature of institutions as well as a substantial amount of empirical work. This can help institutional economists to advise in general terms but the lack of a consensus understanding of how organic institutions emerge (and the potential role of legal rules) is an important problem in the face of new institutional challenges. Contributing to this literature, in order to help solve real-world problems, is the ultimate goal of this thesis.

### 1.1.2 Challenges in the Literature

Here we focus on the ‘gap’ in the literature concerning institutional emergence.

---

<sup>8</sup>At the time of writing, some central banks (e.g., The Bank of England) are exploring digital currencies similar to cryptocurrencies but these would be exchangeable for fiat money issued by the same central banks.

<sup>9</sup>in exchange for traditional currencies.

<sup>10</sup>We can add the enormous energy and environmental costs of operating Bitcoin.



A particularly helpful paper is [Hodgson \(2002a\)](#), which is entitled *The Evolution of Institutions: An Agenda for Future Theoretical Research*. This article contains a number of important arguments regarding research concerned with the emergence of new - and the evolution of existing - organic institutions.

Some good progress regarding organic institutional emergence has been made since the publication of this paper, e.g., [Hodgson and Knudsen's \(2004\)](#) simple traffic convention model. However, the need for the research it promotes appears equally relevant at the time of writing.

Before we look more closely at this paper, however, let us briefly consider the literature concerned with spontaneous order. Chapters 3 and 4 below contain detailed discussions of this (mostly theoretical) literature.

The earliest known reference to spontaneous order was from the fourth century BC Chinese Philosopher Zhuang Zhou who observed that “good order results spontaneously when things are let alone” ([Hamowy, 1987](#), p. 6). Several Enlightenment figures, including David Hume, Adam Ferguson, and Adam Smith were focused on spontaneous order *avant la lettre*; and, more recently, Friedrich von Hayek wrote extensively about it.

Chapter 4 is focused on Hayek's framing of spontaneous order. Hayek is generally regarded as “the most famous and systematic spontaneous order theorist” ([Luban, 2020](#), p. 69). We discuss in that chapter how, despite the impressive ontology on which it is based<sup>11</sup>, Hayek's framing of social evolution has been heavily criticised, both due to its conceptual weaknesses and its limited ability to explain empirical observations.

In what follows we focus on five points made in [Hodgson \(2002a\)](#), which helped guide the computational research reported in this thesis. Moreover, these points also heavily influenced the research questions stated in the next section.

### **Institution-Free World and Infinite Regress**

Geoffrey Hodgson notes an important criticism of some New Institutional and game theoretic explanations of institutional emergence, which was provided by Alexander Field ([1979, 1981, 1984](#)). Field argues that “game theory must be used to explain the emergence of some institutions, but to do so it has to assume a significant number of rules and constraints at the outset.” ([Hodgson, 2002a](#), p. 114).

The problem Hodgson points to is that pre-existing institutions have to be explained, which gives rise to an infinite regress problem.

---

<sup>11</sup>We also discuss how this ontology is broadly consistent with CE.

This problem, however, does not necessarily negate attempts to theorise about or model the emergence of new institutions. Rather, the crucial question is whether the mechanisms that explain how a new institution emerges have any explanatory power for how any assumed institutions came about.

In Chapter 5 we discuss Field's points in more detail. The approach taken in this thesis is to fully accept that new institutions can emerge in part due to pre-existing, enabling institutions. This is linked to one of the features of CE discussed in Chapter 2: openness to stratified ontologies. For now this can be understood as accepting that emergent structures (which represents a single 'stratum') within a population can give rise to further emergent structures (additional strata). Such stratified ontologies are not uncommon in the complexity sciences.

Furthermore, the models developed for this thesis allow space for (but do not intentionally design) new social structures (including institutions) that can emerge from previously emerged social structures. Indeed, simulations based on the second model show four emergent strata: (i) agents learn to defend their resources from attempts at theft; (ii) this then (in part) enables property rights to emerge; (iii) efficient markets are then enabled by property rights; and (iv) agents specialise as a result of being confident of trading (because of the market).

The question remains, however, as to whether the mechanisms identified in institutional emergence in (ii) and (iii) above are relevant for understanding prior institutions. This is discussed in the Conclusion: we cannot be absolutely certain but there are some features common to both (ii) and (iii) which might explain these pre-existing institutions.

In his paper, Hodgson also promotes the idea of "a more evolutionary and open-ended framework of analysis. Instead of focusing on just two points in time - the given starting point and the evolved outcome - the next step is to develop an evolutionary approach, in which the emphasis is on the ongoing process of change." (Hodgson, 2002a, p. 116)

The approach taken in the models developed for this thesis is to employ CE as an overarching ontology and agent-based modelling. Both of these are compatible with Hodgson's open-ended framework and they give rise to the four strata noted above, which seem consistent with Hodgson's open-endedness.

### **Malleable Preferences / Mental Models**

As Hodgson (2002a) states, "most economists recognise that preferences are malleable in the real world [but] they have often taken the assumption of fixed preferences as a reasonable, simplifying assumption." (p. 116). In this paper, Hodgson argues the case for treating preferences as malleable and hypothesises that such malleability plays a role in institutional emergence.

This emphasis sits neatly alongside a general acknowledgment of mental model flexibility in the complexity sciences, including CE. Within IE, the importance of mental models is explored in [Denzau and North \(1994\)](#) where the process of co-adaptation of agents' mental models is discussed<sup>12</sup>. Furthermore, examples of computational research that feature such co-adaptation in CE include [Arthur \(1994\)](#), [Kirman and Vriend \(2000\)](#), and [Vriend \(2002\)](#). An example from within IE is [Hodgson and Knudsen \(2004\)](#), which is analysed in detail in Chapter 5.

The approach taken in this thesis is to build mental model flexibility into both the theory and the computational models. Most notably, in Chapter 2 we discuss the idea of 'semi-permeable agents' who exchange information across their 'boundaries', resulting in mutual influence and co-adaptation. Moreover, the agents in the two computational models developed for this thesis are endowed with mental models that can change as a result of reinforcement learning and habituation (both of these are defined in Section 1.4 below).

The simulation results of both models do indeed show, consistent with the results of [Hodgson and Knudsen's \(2004\)](#) simple traffic convention model, that the co-adaptation of the agents' mental models does play a role in organic institutional emergence.

## Habits

When we consider the enactment of well-established institutions in our every day lives, it is clear that habits play an important role. For example, when beginning a car journey in the UK, in general we drive on the left side of the road without consciously / deliberately choosing that side. We generally do not think about the Highway Code, nor the possibility of punishment if we are caught driving on the right<sup>13</sup>.

There is a now a large body of research concerned with the role played by habits in human cognition. This is emphasised by pragmatist philosophers, notably Charles Sanders Peirce, William James, George Herbert Mead and John Dewey; and also by the original institutionalists, especially Thorstein Veblen<sup>14</sup> and Walton Hamilton. Other notable proponents of habits within human cognition include Hayek, Gary Becker and, within IE, Hodgson.

The source of contention among a number of these researchers is less about the importance of habits than it is about the precise mechanisms at play. For example, [Becker \(1992\)](#) frames habits as "sequentially correlated behaviour" ([Hodgson, 2004b](#), p. 653)

---

<sup>12</sup>[Denzau and North \(1994\)](#) use the phrase 'co-evolution' to refer to agent-to-agent adaptations. This thesis uses 'co-adaptation' instead, and uses 'co-evolution' to refer to the simultaneous evolution of two or more groups of agents (consistent with [Janzen, 1980](#)).

<sup>13</sup>We can say the same thing about institutions like money and language: in most transactions in the UK we do not think about which currency to use, and the same is true of speaking English.

<sup>14</sup>Veblen even included the term in his definition of institutions, which he viewed as "settled habits of thought common to the generality of men." ([Veblen, 1919](#), p. 626).

which mitigates information search costs (see also [Stigler and Becker, 1977](#)); whereas Hodgson emphasises our propensity to repeat previous thoughts or action (discussed further below).

Looking again at [Hodgson's \(2002a\) \*Agenda\*](#), he writes that a “pressing issue for future research is the extent to which these mechanisms of habituation play a role in different cases of institutional evolution.” ([Hodgson, 2002a](#), p. 117). This thesis explores whether habituation plays a role in institutional emergence by including both reinforcement learning and habituation in the agents’ mental models, and then observing the effects of both.

### **Role of the State and Legal Rules**

Some researchers in IE refer to ‘formal institutions’ and ‘informal institutions’ as if the two are fully distinguishable, e.g., [Williamson \(2000\)](#) writes about “institutions (formal and informal)” (p. 597). It is a tempting distinction, that we should think of social norms and conventions as a different category to legal rules like laws and regulations. However, we find that on closer inspection this distinction is confusing. Most notably, laws do not exist ‘outside’ of the human mind - they influence mental models in a way that blurs the above neat distinction.

This raises a number of questions, one of which is “whether, and if so in what circumstances, the state or other powerful organisations can facilitate the emergence and stability of other institutions.” ([Hodgson, 2002a](#), p. 113). This is an important question for the thesis: the two models developed focus on the endogenous emergence of organic institutions but, using the second model, we can also question whether legal rules might enable such institutions when they do not emerge endogenously.

To that end, four experiments are set out in [Chapter 12](#), each of which starts with a set of conditions when it is known that property rights do not emerge endogenously. We find that in all of these, legal rules can indeed facilitate the emergence of property rights.

Moreover, when the results of these experiments are analysed, we can understand the impact of the legal rules on the agents’ mental models in the context of the mechanisms identified when property rights emerged spontaneously. We see that this is mainly about the impact of legal rules on reinforcement learning.

### **A Paucity of Agent-Based Models**

The main point to note here is the lack of agent-based models focused on institutional emergence. This was true in 2002, when “in the twenty years since the first mass production of the cheap microcomputer, very few agent-based computer simulations

exhibiting the emergence of an institution along Mengerian lines have been published.” (Hodgson, 2002a, p. 112). It also appears to be broadly true at the time of writing.

The models developed for this thesis are attempts to add to the literature here but not for its own sake: as Gräbner (2016) argues, Agent-Based Models (ABMs) appear well suited to institutional research. Included in this thesis is a discussion of how ABMs are sufficiently flexible to enable a pragmatic, bottom-up focus on mechanisms that give rise to organic institutions.

In Chapter 5 we describe and evaluate a few notable exceptions to this lack of ABMs. The most important are those of Hodgson and Knudsen (2004) and the four EMergence In the Loop (EMIL) models. The latter were reported in 2010 as part of the EMIL Project, a 3-year, multi-million Euro project funded by the European Union. In that chapter we also look at some game theoretic models as well as various models of monetary emergence, e.g., Marimon et al. (1990).

While the computational models critiqued in Chapter 5 have contributed to the literature, there remains a remarkable lack of such models (relative to, say, the enormous body of work based on game theory). Three fairly recent papers (co-) authored by Claudius Gräbner<sup>15</sup> confirm this point: several ABMs in IE are discussed but the only model concerned with the emergence of organic institutions referred to is Hodgson and Knudsen (2004).

### 1.1.3 Social Simulation Literature: Emergence and Immergence

While in this thesis we focus on economic institutions, we should note some significant advances in our understanding of social norm emergence within the social simulation (and related) literatures. Most notably, the EMIL Project Report was published in 2010 and this has been followed by a number of related papers, e.g., Andrighetto et al (2010), Andrighetto, Campennì and Conte (2010), Castellani (2010), and Conte et al (2013).

This work goes beyond the ‘Beliefs, Desires, and Intentions’ (BDI) and ‘Beliefs, Obligations, Intentions, and Desires’ (BOID) architectures, and is heavily influenced by Conte and Castelfranchi’s (1995a) idea of *cognitive emergence*.

The body of work concerned with social norms in the social simulation literature is enormous. A detailed evaluation of this work is beyond the scope of this thesis; however, two components of the ‘norm emergence’ strand are particularly important for our purposes (one conceptual and the other related to models). The conceptual component

---

<sup>15</sup>Gräbner and Kapeller (2015), and Gräbner (2016, 2018).

is the idea that social norms simultaneously emerge among agents and *immerge* “in the agents’ minds” (Andrighetto et al, 2007, p. 11). This concept of immergence is defined below.

The modelling component is the set of four computational models contained in the EMIL Project Report: the simulation results indicate that norms ‘immerge’ via reinforcement learning<sup>16</sup>. These models are evaluated in Chapter 5.

Cognitive emergence (Conte and Castelfranchi, 1995a) occurs “when agents become aware, through a given ‘conceptualisation’, of a certain ‘objective’ pre-cognitive (unknown and non-deliberative) phenomenon that is influencing their results and outcomes, and then, indirectly, their action.” (Castelfranchi, 1998, p. 27, emphasis removed).

Gilbert (2002) helpfully distinguishes between two different types of cognitive emergence. The first is when some phenomenon arises in the minds of agents while not being recognised as a distinct concept. By contrast, ‘second order’ cognitive emergence occurs when agents explicitly identify some emergent phenomenon, like a ‘club’ or ‘society’. This recognition of a new cognitive phenomenon can then influence the agents’ behaviour, e.g., by treating people ‘inside’ and ‘outside’ of some club differently<sup>17</sup>.

The notion of ‘immergence’ is closely related to Gilbert’s (2002) first type of emergence. Immergence is defined here as “the process through which the macro-level emerging structure or global result ‘feedbacks’ [sic] into the micro-level re-shaping the ‘elementary’ behaviours.” (Castelfranchi, 1998, p. 26). We will see below that this closely resembles the idea of downward causation / effects in, e.g., Sperry (1969), Campbell (1974), Hodgson and Knudsen (2004), and Hodgson (2011).

In this thesis we are interested in Gilbert’s (2002) first type of emergence described above, which is closely associated with our definition of ‘immergence’. The second form of emergence seems fascinating and important in a general sense but it is outside of the scope of this thesis.

---

<sup>16</sup>The Report indicates that norms emerge via reinforcement learning, imitation, and “normative learning” (Lotzmann and Mohring, 2010, p. 98). A closer look at the second of these, however, indicates that it is essentially an indirect form of reinforcement learning (relating to other agents’ experiences). The third concerns agents learning about pre-existing norms, which is more about norm diffusion than emergence. These issues are discussed further in Chapter 5.

<sup>17</sup>Related to this, in the simulations based on the two models developed for this thesis, the agents might meet to transact at the same location in a geographic space without understanding the concept of a ‘market’; and they might choose not to steal each other’s resources without knowing what ‘property rights’ are. We can think of ‘markets’ and ‘property rights’ here as resulting from reification, which adds a new ‘cognitive category’ to the agents’ mental models.

## Cognitive Institutional Economics

Before we move on to discuss the research questions, we should note a strand of the IE literature that seems highly compatible with the above discussions of cognitive emergence and immergence. It also appears to sit comfortably alongside Hodgson's (2002a) *Agenda*.

Cognitive Institutional Economics (CIE) "investigates the cognitive processes underlying the genesis and evolution of economic and social institutions" (Ambrosino, Fontana, and Gigante, 2018, p. 769). Papers in this field include Rizzello and Turvani (2000); Rizzello and Turvani (2002); Ambrosino (2016); and Ambrosino, Fontana, and Gigante (2018).

Consistent with our interest in real world institutional problems, Rizzello and Turvani (2000) argue that attempts "to translate theory into suggestions for the design of institutions and their reform have generally yielded poor results. Underlying these difficulties is a lack of interest in and a poor understanding of the institutional behavior of the individuals." (p. 177). Furthermore, "we have no theory of the interplay between the individual and the creation of norms." (ibid).

Rizzello and Turvani (2000) argue for a better understanding of institutions as "the expression of the capabilities of the mind, which are not innate but develop and are organized in connection with other individuals" (ibid).

It is worth noting also that this strand of research is "characterized by the dynamic and expanding forces of evolving complex systems." (Ambrosino, Fontana, and Gigante, 2018, p. 780).

However, one of the curious features of this literature is that it has not developed any formal models to explore its terrain.

Overall, the aims of this literature sit well with the orientation of the research presented in this thesis. With this in mind, let us look more closely at the research questions addressed by the thesis.

## 1.2 Research Questions

Here we first look at the research questions and then discuss why CE and ABMs are suited to answering these question. There are two questions:

1. Can organic institutions emerge spontaneously across a population while also immerging within individuals' mental models via reasoning, learning, and habituation? and
2. Can 'liberal legislation' catalyse institutional emergence when it does not occur endogenously?

These questions are discussed and analysed in Chapter 6 so here we discuss them briefly.

The first point to note is that the hypotheses that correspond to these research questions can be easily inferred from them, that: (i) organic institutions emerge spontaneously while also immersing within mental models via reasoning, learning, and habituation; and (ii) liberal legislation (as legal rules in the experiments below) can catalyse institutional emergence when this does not arise endogenously.

The **first research question** focuses on the spontaneous emergence of organic institutions. The question asked is can they emerge because the presupposition is that they do so only when the environment is sufficiently enabling.

The question also refers to both emergence and ‘immergence’. In the discussion of the main results in Chapter 6, the former term is aligned with emergence ‘outside’ of the agents’ boundaries but which results from their actions. The word ‘immergence’ was defined above.

Interestingly, the simulation results presented below suggest that the emergent and immergent processes observed are the ‘process’ counterparts to Hodgson’s (2006a) analogy of a Klein bottle for institutions, for which “the subjective ‘inside’ is simultaneously the objective ‘outside’” (p. 8). This is discussed further in later chapters.

In answering the first research question we focus on the combination of reasoning, learning, and habituation. The CE framing set out in Chapter 2 includes an emphasis on uncertainty and in the computational models set out below, the agents attempt to reason under conditions of uncertainty. This draws on work in the complexity sciences and CE, notably Holland et al (1986) and Arthur (1994).

The reference to ‘learning’ in the first research question is specifically about reinforcement learning, which is defined in Section 1.4.6 below. Does this type of learning play a role in institutional emergence? This follows Erev and Roth (1998) but we should note that this process is well now established in various fields, including the complexity sciences and artificial intelligence.

Inclusion of habituation in the first research question follows from the discussion of habits in the previous section.

We should note that the combination of reinforcement learning and habituation in the agents’ mental models extends the research of Hodgson and Knudsen (2004) and the EMIL Project. Hodgson and Knudsen’s (2004) simple traffic convention model includes agents whose mental models contain instincts and habituation. Hodgson and Knudson conclude that habituation plays a role in institutional emergence. In this paper, the



authors mention that their model could be usefully extended to include forms of reinforcement learning. In effect, the models developed for this thesis meet this challenge, with a focus on property rights and markets.

We will see in Chapter 5 that institutions emerge in the simulations based on the four EMIL models via reinforcement learning (habituation was absent). Two problems with the reported simulation results are identified in Chapter 5: (i) it is difficult to identify the precise mechanisms at play; and (ii) there was almost no attempt to explore the parameter spaces of these models. The two models created for this thesis are partly an attempt to develop this EMIL research, to explore how habituation might work alongside reinforcement learning, and to better identify the mechanisms of emergence and immergence.

The final point to make concerning the first research question is that the focus on reinforcement learning and habituation is not meant to suggest that these are the only mechanisms that play a role in organic institutional emergence. For example, [Elsenbrich and Gilbert \(2014\)](#) discuss how conformity, obedience, and compliance all seem to influence pro-social behaviour. As mentioned in the Conclusion, these could be added to the models developed for this thesis to explore their impact on institutional emergence.

The **second research question** is focused on the issue noted in the previous section: can legal rules help catalyse or enable institutional emergence?

‘Liberal legislation’ is referred to in this second question, which is a term used by Hayek to refer to laws that help enable and support an economy based on free markets. [Vanberg \(1994a\)](#) notes that Hayek’s reference here is related to his defence of classical liberalism (his focus during much of his career). As Vanberg notes, ‘liberal legislation’ should not be viewed as direct intervention in the economy that has the aim of bringing about specific outcomes; rather, it is supportive of markets.

As a final note, it is perhaps worth mentioning that there is overlap between the first and second questions. In fact, we could subsume the second question in the first by framing legal rules as part of enabling environments. However, while a single research question has value, the subject of legal rules seems sufficiently distinct to warrant a second research question that is focused on this subject.

### 1.2.1 Complexity and Computational Modelling

In this sub-section we briefly discuss why CE and ABMs are employed in this thesis to address the above research questions. This subject is considered in more detail in Chapter 2.

The complexity sciences have proven to be well suited to investigating emergent phenomena across multiple fields over the past few decades (e.g., [Kauffman, 1993](#); [Cowan, Pines and Meltzer, 1999](#); [Rosser, 2009](#)). In fact, the affinity between emergence and the complexity sciences is such that the former is viewed by many as the *sine qua non* of the latter (e.g., [Mitchell, 2011](#)) even if the first use of ‘emergence’ preceded the complexity sciences by about a century ([Hodgson, 2000c](#)).

It is important to emphasise, also, that CE (the application of the complexity sciences to economics) now appears mature enough to help us answer the above research questions. This maturation is important because, as discussed in [Chapter 2](#), we need to be careful when mapping from a field that has mostly arisen from the natural sciences into the social sciences. Complexity Economics, which is defined in [Chapter 2](#), helps us with this mapping.

On the value of the complexity sciences in researching economic institutions, the following quote from Hodgson is worth noting:

Although the possibility of general principles of complexity is sometimes exaggerated, the impact of the new ideas on complexity is profound. Not only is the common obsession with precise prediction confounded; the whole tradition in science of attempting to reduce each phenomenon to its component parts is placed into question. A reconstructed institutional economics can usefully build on this literature. ([Hodgson, 2004a](#), p. 408)

Turning to agent-based modelling, there is also now a large body of work in which ABMs have been used to investigate complex phenomena (e.g., [Railsback and Grimm, 2011](#); [Miller and Page, 2007](#); [Elsenbrioch and Gilbert, 2014](#)). This work indicates that ABMs are likely to be well suited to the modelling of organic institutional emergence, a point that is supported by the successful use of an ABM in [Hodgson and Knudsen \(2004\)](#).

This point is also supported by a number of the arguments made in [Gräbner \(2016\)](#), which carefully evaluates ABMs in the context of institutional research more generally. [Gräbner](#) argues that ABMs are a good fit with “the core principles of institutionalist methodology” (p. 252), noting, in particular, that ABMs allow us to model:

- mutually interdependent systems, including ‘downward effects’;
- the dynamics of economic systems, including self-organising processes, stating that “[t]his is exactly what ABMs were invented for.” (p. 248);
- the structures and networks we typically see in economic systems;

- the economy as an evolutionary system<sup>18</sup>, and agency within agents' mental models; and
- agents' "habits, heuristics, and cognitive abilities." (p. 253).

One additional point worth emphasising is that Gräbner (2016) mentions how ABMs are particularly helpful in identifying causal mechanisms, including how different initial conditions within a complex system can generate different results. This theme of causal mechanisms is discussed further in Chapter 5 where some of the limitations of game theoretic models are set out.

We can add to Gräbner's (2016) list of points the idea that both Complexity Economics and ABMs can help us better understand the principle of *formation* in the economy (Arthur, 2013), including institutional emergence. Arthur (2013) argues that orthodox economics has historically focused on the *allocation* of resources and has paid far less attention to how an economy structures and restructures itself over time. W. Brian Arthur refers to the examples of "institutions, arrangements, and technological innovations" (p. 1). Complexity Economics and ABMs both appear well suited to the study of institutional emergence as economic *formation*.

Finally, we should note that Gräbner (2016) identifies three risks with using ABMs:

1. the temptation for researchers "to take a constructionist-instrumental standpoint" (p. 255), which would lead them to "not try to describe ... reality accurately but consider their theories to be mere instruments replicating observed data." (ibid);
2. adopting an "[i]mplicit focus on predictive power" (p. 256), in line "with Friedman's methodological instrumentalism." (ibid);
3. "[o]verparameterisation and decreased transparency." (p. 256). This is the risk of a researcher "adding variables, processes, and methods until one gets ... the patterns one wishes to explain." (p. 256).

These risks, and how they were mitigated in the ABMs developed for this thesis, are discussed in more detail in Chapter 6.

### 1.3 Executive Summary

In this section we provide a summary of the thesis to help with orientation.

Before proceeding, we should note that the assumed audience for this thesis includes researchers interested in IE and CE. Importantly, however, the thesis has been written

---

<sup>18</sup>In the sense of 'generalised Darwinism', as defined in Chapter 2.

so it is also accessible to people interested in self-organisation in the complexity sciences more generally. The models and simulation results ought to be of interest to that group too.

**Chapter 2 - Complexity Science and Complexity Economics** develops the ontological foundations on which the rest of the thesis is based.

It starts by describing the complexity sciences. A distinction is made between *defining features* and *conceptual features*. The latter include concepts that are associated with the field (e.g., emergence) but do not define it.

An argument is then made that economies fit the definition of complex systems. However, as mentioned above, we must be careful when mapping from a domain that has mostly grown out of the natural sciences, to the social science of economics. Eleven principles are outlined that represent this mapping and which also help to define the ontology of a complex economic system, e.g., the use of a stratified ontology where emergent properties are formed ‘from below’.

**Chapter 3 - A History of Spontaneous Order** sets out a summary of the history of this term.

This chapter looks specifically at the framing of the concept of spontaneous order prior to the Enlightenment, by various thinkers during the Enlightenment, and then post-Enlightenment. The chapter finishes by describing the historical influences on Hayek.

**Chapter 4 - Hayek’s Spontaneous Order** focuses specifically on Hayek’s work, including his theory of cultural evolution.

As mentioned above, Hayek’s framing of spontaneous order is generally viewed as the most extensive and sophisticated. We discuss his general ontology (including his theory of mind, his emphasis on dispersed knowledge, and a paper he wrote on ‘complexity’); his reference to ‘catallaxy’ and the ‘market order’; his theory of cultural evolution, and various criticisms of this, particularly those of John Gray, James Buchanan, and Victor Vanberg.

In addition, we discuss [Ullmann-Margalit’s \(1978\) \*invisible-hand explanations\*](#) as background to understanding Hayek but also because her paper dovetails neatly into the two computational models developed for this thesis. We also focus on Vanberg’s discussions of Hayek’s ‘liberal legislation’ to contextualise the second research question and to motivate the experiments discussed in [Chapter 12](#).

**Chapter 5 - Models of Organic Institutional Emergence** evaluates models that are related to the first research question.

The first model discussed is [Hodgson and Knudsen's \(2004\)](#) simple traffic convention. This model and the simulation results are extremely helpful; however, there is one important criticism (which is recognised by the authors): the habituation variable in the agents' mental models can also be interpreted as a proxy for reinforcement learning. This criticism does not negate Hodgson and Knudsen's conclusions but it leaves open the question of whether the traffic convention observed emerged because of habituation or reinforcement learning (or both).

Models of institutional emergence / immergence from the Social Simulations literature are also discussed. The most significant models for our purposes are the four EMIL Project models. These are critiqued in detail: they indicate that organic institutions emerge via reinforcement learning. Some of the criticisms of these models were mentioned above.

Game theoretic models are also discussed, including how such models are used to define institutions either as equilibria or as correlating devices that achieve equilibria (research concerned with rules-as-equilibria is also considered). The main criticisms of these models include: (i) those taken from [Field \(1979, 1981, 1984\)](#), as discussed above; and (ii) questions about the assumption of utility maximization, which means that they cannot and do not help identify the mechanisms by which institutions emerge in the real world. This criticism follows [Hodgson \(2012\)](#) in particular. In this thesis we are interested in what these mechanisms are.

Finally, this chapter looks in some detail at the strand of literature focused on the emergence of money. The papers discussed include [Kiyotaki and Wright \(1989\)](#), [Marimon et al. \(1990\)](#), [Brown \(1996\)](#), [Duffy and Ochs \(1999\)](#), and [Duffy \(2001\)](#). One of the fascinating characteristics of this literature is the mixture of game theory, computational models, and empirical evidence. Most notably, the empirical studies highlight how the Nash equilibria identified by [Kiyotaki and Wright \(1989\)](#) are not observed among live subjects.

**Chapter 6 - Models: Rationale, Design and Results** bridges the theoretical material and models discussed in the previous four chapters and the computational models and results set out in chapters 7 to 12.

The chapter begins by rationalising and decomposing the two research questions, drawing heavily on the previous four chapters. We then discuss various factors that influenced the design of the two computational models presented in chapters 7 and 9; and then an overview of the models is provided.

The final section of this chapter discusses the main results of the simulations presented

in subsequent chapters. The computational research proceeded in an iterative and exploratory way and the main results are presented in lieu of some pre-conceived hypothesis. Also, the models and simulations are relatively large so these results help with general orientation.

**Chapter 7 - Market Emergence Model** describes the first computational model. We develop a rudimentary economic system in which agents forage for resources and then attempt to find each other (within a geographic space) to trade. In this model, agents are assumed to respect each other's property, i.e., the agents do not steal from each other.

This chapter also sets out the results of five sets of simulations: three 'null' experiments; one that uses the 'default parameters' (where mental models change via reinforcement learning only); and then a set that adds habituation to the default set. The main results are: (i) we observe that the agents' mental models co-adapt in such a way that markets emerge, i.e., single locations where all the agents congregate in order to trade; and (ii) habituation can obstruct 'immersion' but it also ensures, over time, that these locations are firmly embedded in the agents' mental models.

Furthermore, the agents begin life as generalists foraging for two resources but they specialise after the market emerges, i.e., specialisation is contingent on the existence of a market. Productivity then increases, as does the total population.

Finally, we compare the single emergent markets with [Hodgson's \(2006a\)](#) definition of institutions (see Section 1.4.1 below): the markets appear to be a close fit with this definition.

**Chapter 8 - Market Emergence Model: Exploring the Parameter Space** includes a summary of a range of experiments that adjust each of the parameters in the first model. Detailed results and analysis are set out in [Appendix B](#).

It is commonplace to explore a parameter space to understand the robustness of some results; however, here, the results are more helpful than this. They make it clear that the market institution only emerges when the environment is sufficiently enabling. Markets do not *necessarily* emerge but they can under specific conditions.

**Chapter 9 - Property Rights Model** presents the second model, which is a variation of the first. Now, agents can steal from each other. This is done by giving agents a propensity to steal others' resources and a propensity to defend their own resources. These propensities are metrics which vary via reinforcement learning and habituation.

**Chapter 10 - Property Rights Model: Simulations and Discussion** presents the results of simulations based on the second model, including four 'null' experiments; the new default parameter set (based on reinforcement learning); and a set that includes

habituation. In the default simulations we find that *defence of property* emerges among the agents but this becomes obsolete after *property rights* emerge. The latter is sustained across the population until the end of the simulations. Also, this organic institution enables the emergence of efficient markets and then specialisation, i.e., the results of the first model's default simulations are replicated.

When habituation is added to the default parameter set, we find on the whole it catalyses the emergence of property rights. Also, the results indicate that reinforcement learning is essential if property rights are to emerge but, like the first model, habituation helps to embed the institution of property rights in the agents' mental models.

Finally, we compare the property rights observed with Hodgson's (2006a) definition of institutions (set out below): they appear comfortably consistent with this definition.

**Chapter 11 - Property Rights Model: Exploration of the Parameter Space** includes a summary of a range of experiments that adjust each of the parameters in the second model. Detailed results and analysis can be found in Appendix C. As with the first model, the results indicate that the institution of property rights only emerges if the conditions are sufficiently enabling.

**Chapter 12 - Liberal Legislation** is focused on the second research question.

Legal rules are applied to the three experiments set out in Appendix D in which property rights do not emerge. In addition, they are applied in a fourth scenario, when we know from exploring the parameter space that institutions do not emerge endogenously (when the 'cost of fighting' is below a specific threshold).

In all four experiments we identify legal rules that enable the emergence of property rights.

Furthermore, in all four experiments we explore whether corruption undermines the efficacy of legal rules. We find, perhaps unsurprisingly, that corruption does indeed have a detrimental impact on the emergence of institutions.

**Chapter 13 - Conclusion** summarises the thesis.

In this chapter we first address the two research questions. In discussing the first we identify the mechanisms observed in the emergence (and immergence) of property rights and markets. We find: (i) there are mechanisms common to both institutions; but (ii) for the second model an additional mechanism is observed, that of selection within a generalized Darwinian process. The common features are set out in the description of a generalised framework: we use this to discuss emergence/immergence of both organic institutions and to highlight idiosyncrasies.

In discussing the second research question, we highlight how legal rules influence the mechanisms observed within the above generalised framework.

In addition, the chapter revisits the question of an ‘institution-free state of nature’, discussed above; sets out the lessons learned from the process of research which preceded this thesis; and, finally, potential future research is discussed.

**Appendix A - Interaction Typologies** considers, briefly, different typologies of interaction types. It looks at on [Schotter’s \(1981\)](#) framing, which developed [Ullmann-Margalit \(1977\)](#), and that of [Schultz \(2001\)](#). As discussed in [Section 1.4.9](#) below, this thesis adopts [Schultz’s \(2001\)](#) typology.

**Appendix B** contains detailed discussions of the exploration of the parameter space of the first model (summarised in [Chapter 8](#)).

**Appendix C** explores the parameter space of the second model. A summary of this work is included in [Chapter 11](#).

**Appendix D - Property Rights Model: Experimentation** sets out three experiments that are based on the second model. These three scenarios are designed to be used in the ‘liberal legislation’ experiments of [Chapter 12](#). The first experiment assumes agents never defend their resources when challenged by other agents, and the last two explore different forms of ‘power’ (used to determine the outcome of conflicts between agents). In versions of all three of these experiments, property rights never emerge endogenously.

## 1.4 Assumptions and Definitions

The research conducted for this thesis can be split in two dimensions. The first is between computational modelling and theoretical work (approximately 50% of research time was spent on each); and within the theoretical content, there is a split between the three overlapping subjects of IE (incorporating spontaneous order), the complexity sciences (including CE), and published computational models. Here, the research time was split approximately 40-40-20 (respectively).

Moreover, the study of institutions spans multiple literatures, including sociology, psychology, philosophy, and economics.

The tension between the limited research time available and the enormity of the literature was dealt with by narrowing the scope of the research in specific ways. One of these was to use specific definitions of various concepts that are important in this thesis: these are set out below.



Each of these concepts is located in a large literature. In what follows we specify the definition used and, where appropriate, briefly discuss aspects of the related literature.

### 1.4.1 Institutions

Gräbner and Ghorbani (2019) provide a helpful summary of the various definitions of institutions used across different literatures, identifying a total of eleven<sup>19</sup>.

Other useful overviews of the large literature concerned with defining institutions include Hodgson (2015), which discusses the difference between institutions-as-rules, institutions-as-equilibria, and rules-in-equilibrium (Hindriks and Guala, 2015); and Hodgson (2019), which discusses definitions of institutions in the context of taxonomic definitions in general.

The focus of the models developed for the thesis is partly on the cognitive mechanisms of institutional emergence: the origination of new rules is a part of this so a definition of institutions as rules (rather than equilibria) seems more appropriate. Furthermore, being “outcomes of individual interactions, equilibria are secondary to the relational framework that generates their possibility.” (Hodgson, 2015, p. 500). This thesis focuses on the ‘primary’ mechanisms by which organic institutions emerge.

As stated above, in this thesis we adopt Hodgson’s definition of institutions as “durable systems of established and embedded social rules that structure social interactions.” (Hodgson, 2006a, p. 13). This definition is commonly cited in the IE literature and sits comfortably in the institutions-as-rules category.

Looking more closely at this definition:

- ‘durable’ means that there are processes in place that maintain the institution over time, including the idea that the institution “can usefully create stable expectations of the behavior of others” (Hodgson, 2006a, p. 2);
- ‘systems’ points to the idea that institutions can include features that are cognitive and external (to the agent), including ‘artefacts’;
- ‘established’ and ‘embedded’ convey a sense of the institution being manifested in the population in a stable way;
- ‘social’ helps to distinguish these types of rule from personal rules (which are about choices made when there is no interaction, e.g., tying one’s own shoelaces);

---

<sup>19</sup>The authors also provide a synthesis definition which is not too far from that used in this thesis: institutions are “codifiable systems of social structures (in particular norms and rules) that lead to inclinations for people to act in specific ways.” (Gräbner and Ghorbani, 2019, p. 1)

- ‘rules’ is an important word in this definition and is discussed in more detail below; and
- ‘structure social interactions’ refers to the organisation of the parts of an institution system (including artefacts) and the agents in the interaction.

Let us now focus on the word ‘rule’ in this definition:

The term *rule* is broadly understood as a socially transmitted and customary normative injunction or immanently normative disposition, that in circumstance  $X$  do  $Y$ . (Hodgson, 2006a, p. 3)

The idea that a rule is *socially transmitted* “means that the replication of such rules depends upon a developed social culture and some use of language ... they depend upon contingent social structures” (ibid).

The definition requires that a rule either be a “customary normative injunction” or an “immanently normative disposition”. The normative component of both requires alternatives to option  $Y$ , e.g.,  $Y^*$ . If  $Y$  was the only option available to an agent then there would be no need for a norm.

The former term points to some type of social pressure that is directing of agents in some way (to do  $Y$  and not  $Y^*$ ) whereas the latter appears to be more about cognitive tendencies to behave in a particular way. In addition, “*immanently normative* requires that if the rule is scrutinized or contested, then normative issues will emerge” (ibid).

In terms of the *types* of rules, Hodgson (2006a) refers to three: social norms, social conventions, and legal rules.

One distinction between these three is that legal rules are typically formalised, i.e., written down and enforced by designated third parties; whereas social norms and conventions are typically neither written down nor enforced by a designated third party (although they might still be enforced in some way, e.g., by those who observe a transgression).

## 1.4.2 Organic Institutions

These are unplanned institutions. In this thesis we mostly use ‘organic’ but the two words refer to the same thing.

This idea of organic institutions is linked to Carl Menger, notably Menger [1890] (1981) and Menger [1883] (1985). His account of monetary emergence is an archetypal example which is discussed in more detail in Chapter 3.

These types of institution fall within Ferguson’s definition of spontaneous order stated above: they result from human action in that they emerge endogenously in a population but they are not intentionally designed.

There are other ‘classical’ types of spontaneous order. Those most often referred to in the literature are language, money, and the order brought about by markets. Some, notably Hayek, claim that English common law is also an example of spontaneous order but this is contested (and is discussed in more detail in Chapter 4).

Moreover, organic institutions stand in contrast to planned institutions (legal rules are the most obvious example of this). Here, there is intent behind the creation of the rule.

### 1.4.3 Order

We use [Ferguson’s \(1767\)](#) definition of spontaneous order in this thesis; however, we should also be clear about what is meant by ‘order’. [Luban \(2020\)](#) summarises well the problems we face when attempting to do this:

‘Order’ is an ambiguous term. It can refer to any system exhibiting regularities of any kind; alternately, it can be restricted to only those systems exhibiting certain normatively desirable regularities. A social order, for instance, might be defined either by the mere existence of ‘predictable patterns of behavior,’ or more strongly by the prevalence of ‘cooperative behavior’ ([Elster, 1989, 1](#)). Others draw a different distinction between ‘normative order’ and ‘factual order’ that depends on how order is produced: whether by subjective consensus, or by unintended consequences occurring behind the backs of actors themselves ([Parsons, 1968, 91–2](#)). ([Luban, 2020, p. 69-70, footnote 5 removed](#))

There are a variety of definitions of order we might use but there is no consensus in the literature regarding which is preferable. Here we adopt Hayek’s definition, stated below, as a working definition because it seems more compatible with this thesis than others. There are two parts to this. First, Hayek’s definition makes no normative commitment, i.e., it is more generalised than those focused on ‘good’ outcomes. We will see that in some of the simulation results presented below that ‘bad’ institutions emerge under certain circumstances and ‘good’ institutions emerge under others.

The second reason is that Hayek’s definition includes the important concept of expectations, which fits with many other researchers’ references, e.g., [North \(1990\)](#), [Ostrom \(1991\)](#), [Aoki \(2001\)](#), [Hodgson \(2006a\)](#), and [Gräbner and Ghorbani \(2019\)](#). We observe in the simulations reported below that ‘good’ and ‘bad’ order are both associated with expectations fulfilment.

Hayek's definition of order is:

a state of affairs in which a multiplicity of elements of various kinds are so related to each other that we may learn from our acquaintance with some ... part of the whole to form correct expectations concerning the rest, or at least expectations which have a good chance of proving correct. (Hayek, 1973, p. 36, emphasis excluded)

### Spontaneous Order Enabled by Planning

As stated above, in this thesis we adopt Ferguson's definition of spontaneous order, which includes the words "not the execution of any human design". For the second research question, it is important to understand how spontaneous order can result from planning, which might appear counter-intuitive. The point is that forms of planning can *enable* order to emerge spontaneously, including organic institutions, rather than directly *causing* it.

In this context, Hayek (1973) writes "it is at least conceivable that the formation of a spontaneous order relies entirely on rules that were deliberately made." (p. 45). This is perhaps surprising given his vociferous arguments against central planning as 'constructivist rationalism', e.g., in Hayek (1945, 1988). However, Hayek's words "relies ... on" does not mean "deliberately caused by", which is an important distinction.

Similarly, Boehm (1994) writes that it "is certainly feasible that a spontaneous order *may arise* from deliberately designed rules." (p. 299, emphasis added). Consequently, Hodgson (1994) argues that this "suggests that there is no complete bifurcation between organic and pragmatic institutions: even spontaneous orders may have designed elements" (p. 584).

The 'liberal legislation' experiments of Chapter 12 help us understand how forms of planning (specifically, legal rules) can enable subsequent organic institutions as spontaneous order. Legal rules are imposed on the agents, and these enable the emergence of property rights (this is planned). However, markets subsequently emerge, which can be interpreted as spontaneous order that was enabled by planning.

#### 1.4.4 Legal Rules

These are included in the three types of rules mentioned in Hodgson (2006a), noted above. In the 'liberal legislation' experiments reported in Chapter 12, legal rules are applied to agent interactions so it is helpful to be clear what these are:

For laws to become rules in the sense discussed here, they have to become customary ... there are examples of laws that are widely ignored and have

not acquired the customary or dispositional status of a rule. Ignored laws are not rules. For new laws to become rules, they have to be enforced to the point that the avoidance or performance of the behavior in question becomes customary and acquires a normative status. (Hodgson, 2006a, p. 6)

The customary nature of legal rules means they are both codified and known to the agents in a population. Furthermore, these rules are enforced by a third party.

### 1.4.5 Habits

In this thesis we define a habit as:

...a largely non-deliberative and self-actuating propensity to engage in a previously adopted pattern of behavior. A habit is a form of self-sustaining, nonreflective behaviour that arises in repetitive situations. (Hodgson, 1998, p. 178)

In Hodgson and Knudsen's (2004) computational model, this definition is translated into code by the creation of a habituation variable. This variable quantifies past decisions and is applied within the agents' mental models as part of decision making.

The habituation experiments conducted using the two models developed for this thesis (sections 7.6 and 10.3) mimic both the above definition and Hodgson and Knudsen's (2004) approach.

Moreover, it is worth highlighting here that in the models developed for this thesis, changes in the agents' mental models (referred to as 'reconstitution') occur via two mechanisms: reinforcement learning and the process of habituation. We can think of the former as being sensitive to the agents' perception of the success or failure of its actions, relative to some goal(s); whereas the latter is inssensitive to success and failure.

A final point is to differentiate between the approach to habituation above and that of Becker (1992). Becker sees habits as arising from people repeating previously 'rational' decisions, i.e., habits are linked to a type of reasoning. The models developed below allow us to explore the combination of reasoning under conditions of uncertainty, reinforcement learning, and habituation that is independent of reasoning and learning.

### 1.4.6 Reinforcement Learning

This form of learning has been studied and used in a variety of literatures, notably in the complexity sciences, psychology, and artificial intelligence.

Here we adopt [Erev and Roth's \(1998\)](#) framing (see also [Roth and Erev, 1995](#)). In their 1998 paper, Erev and Roth write that reinforcement learning is linked to three principles. The first is [Thorndike's \(1898\) Law of Effect](#), which is when choices “that have led to good outcomes in the past are more likely to be repeated in the future.” ([Erev and Roth, 1998](#), p. 859).

The second principle is [Blackburn's \(1936\) power law of practice](#), which indicates that learning curves “tend to be steep initially, and then flatter.” ([Erev and Roth, 1998](#), p. 859).

The third principle, which is implicit in the first, is that “*choice behavior is probabilistic*. This is one of the basic assumptions of most mathematical learning theories proposed in psychology.” (ibid, emphasis included). Here, Erev and Roth cite, *inter alia*, [Estes \(1950\)](#).

In this thesis, the first and third principles are firmly employed in the reinforcement learning by the agents in both models. Also, the first principle is assumed to apply symmetrically, i.e., choices that have led to bad outcomes are less likely to be repeated in the future.

The second principle turns out to be irrelevant in the first model and problematic in the second (this is discussed further below). This principle is, however, applied to changes in the agents' foraging skills.

Note, also, that reinforcement learning in the models below occurs in the context of agents reasoning under conditions of uncertainty in a complex economic system. This should be viewed as very different to Becker's use of substantive rationality from which habits (supposedly) arise, e.g, in [Becker \(1992\)](#).

### 1.4.7 Property Rights

The following definition is taken from the Concise Encyclopedia of Economics:

A property right is the exclusive authority to determine how a resource is used, whether that resource is owned by government or by individuals. Society approves the uses selected by the holder of the property right with governmental administered force and with social ostracism. [Alchian \(2023\)](#)

In the second model described below (Chapter 9), agents have ‘propensities to steal’ that determine whether they attempt to steal from or trade with other agents. In the resulting simulations, we say that property rights are established when all of the agents' propensities to steal are below zero<sup>20</sup>.

<sup>20</sup>The precise mechanics of this are described in Chapter 9.

### 1.4.8 Downward and Upward Causation / Effects

In this thesis, *downward causation* “refers to possible effects of higher-level properties on lower-level components. The term ‘downward causation’ originates in psychology in the work of Sperry (1969).” (Hodgson and Knudsen, 2004, p. 38).

In Chapter 5 we discuss strong and weak forms of downward causation, following Sperry (1969) and Campbell (1974), respectively. These are referred to in Hodgson and Knudsen (2004).

Equivalently, *upward causation* refers to the effect a lower-level components has on a higher-level property.

We should note, also, Hodgson (2011) who proposes a change to the language here: “causation” seems inappropriate in light of the relationship between different levels in a stratified ontology. Higher level properties do not ‘cause’ lower-level components to act in certain ways. Hodgson’s suggestion is to replace ‘causation’ with ‘effects’. We generally follow this terminology in the thesis.

The final point to note is that while upward and downward effects are “largely unfamiliar to economists” (Hodgson and Knudsen, 2004, p. 38) in general, they sit comfortably within CE as described in Chapter 2.

### 1.4.9 Types of Interaction Challenges

Here we note and define two types of challenge observed when people interact, both of which are relevant for the models developed for the thesis. The literature can be somewhat unclear and incomplete when it comes to the types of interactions people face, e.g., Vanberg and Buchanan (1988) distinguish between “coordination” and “prisoners’ dilemma situations”.

In this thesis we adopt Schultz’s (2001) framing as it seems more comprehensive than others, e.g., Schotter (1981) (see Appendix A for a more detailed discussion).

As mentioned above, Schultz (2001) refers to two types of interaction: (1) coordination situations; and (2) collective action situations. The former includes situations where the agents in an interaction share some preferences over states<sup>21</sup> of the world and where a Nash equilibrium exists at the agents’ jointly preferred state<sup>22</sup>. This first category aligns with coordination games (both pure and non-pure) in game theory and ‘solutions’ are self-sustaining.

<sup>21</sup>Preferences need not be identical but there must be some commonality of ranking between all the agents in the interaction so that some outcomes are jointly preferred by the agents over others.

<sup>22</sup>Schultz uses different terminology, referring to whether strategies “coordinate” or “conflict” but the meaning of the former appears to be identical to that of a Nash equilibrium.

Schultz's collective action situations also includes overlapping preferences but a Nash equilibrium does not exist at the jointly preferred state. This includes the Prisoners' Dilemma but accommodates a much broader range of interaction types where an institution could help bring about a pareto superior outcome (when the agents would otherwise not bring this about). Unlike the first category, 'solutions' are not self-sustaining without mechanisms that help achieve them<sup>23</sup>.

Let us also be clear about nomenclature regarding 'problems' here. We will refer to problems in Schultz's first category as 'coordination problems' and to those in his second category as 'free-rider problems'. Strictly speaking, the origins of the latter lie with a particular class of collective action situation, that of public goods; however, we will use the term more generally, to apply to the wider category<sup>24</sup>. This appears consistent with other researchers in the literature, e.g., [Vanberg \(1986\)](#), and also the [Stanford Encyclopedia of Philosophy \(2023\)](#).

The final point to note here is that the first model described below (Chapter 7) broadly aligns with [Schultz's \(2001\)](#) first category and the second model (Chapter 9) gives rise to a wide range of interactions that fit into Schultz's second category, including the Prisoners' Dilemma.

## 1.5 Added Value

The value added by this thesis should be viewed in the context of various computational and theoretical foundations established by other researchers. As mentioned above, the most significant computational models are those of [Hodgson and Knudsen \(2004\)](#) and the four EMIL Project models, which are discussed in Chapter 5. The two models developed for this thesis extend this work, notably the reinforcement learning of the EMIL models and the emphasis on habituation in [Hodgson and Knudsen \(2004\)](#).

The theoretical parts of this research build on the notions of upward and downward effects, including reconstitution, discussed by [Hodgson and Knudsen \(2004\)](#) and in various papers written by Hodgson, e.g., [Hodgson \(2002c, 2003c, 2006a, 2006c, and 2011\)](#). The work also builds on CE, especially the work of W. Brian Arthur; and on the idea of immergence.

---

<sup>23</sup>Note that [Schultz's \(2001\)](#) framing is not exhaustive: he refers to interactions where no shared preferences exist as "moot" (p. 64).

<sup>24</sup>If collective action situations are defined by the lack of Nash equilibrium, this implies at least one agent would prefer to adopt a different strategy, which would be to their benefit at the cost of at least one other agent. In a more general sense, therefore, the agent would free-ride on others.



The claim here is that there are four ways in which the research adds value to the literature (with some overlap between them). The first two should be viewed as stronger claims than the third and fourth.

The **first value added point** is about the relevance of both reinforcement learning and habituation in the emergence of organic institutions (and is tied to the first research question).

Various experiments were conducted using both models but the main conclusions are:

- co-adaptive reinforcement learning is essential in the emergence of organic institutions in both models' simulation results (this supports the EMIL models' results); and
- while (co-adaptive) reinforcement learning is generally more important in the emergence of markets and property rights, over time habituation dominates agents' mental models such that by the end of simulations, these institutions simply look like habits (this is observed in experiments using both models when the habituation parameter is relatively small).

The **second value added point** is concerned with 'liberal legislation' (and is related to the second research question). The simulation results indicate that in situations when property rights do not emerge endogenously, legal rules can help to enable this emergence; however, this result is conditional on any corruption being overcome.

The mechanisms by which legal rules enable property rights must be understood in the context of mechanisms observed when organic institutions emerge endogenously. Most notably, reinforcement learning was essential to ensure agents co-adapted in a way that was consistent with the legal rule. This rule created an emergent feedback loop that would have been missing otherwise. This links directly to [Hodgson's \(2002a\)](#) query as to whether state organisations can facilitate the emergence of institutions, noted above.

The **third value added point** relates to the identification of a general framework that includes features that appear common to the simulations based on the two models developed for this thesis. These also appear common to the simulation results of [Hodgson and Knudsen \(2004\)](#) and those of the four EMIL models<sup>25</sup>.

Note that this general framework is presented here as a minor value added point because it extends work developed by other researchers.

The general framework, which is discussed in detail in the Conclusion, contains the following seven features:

---

<sup>25</sup>We should note that the correspondence between this framework and the EMIL simulation results is less clear because the descriptions of these results are extremely brief.

1. an emerging property, which is ‘external’ to the agent;
2. downward effects;
3. upward effects;
4. emergence via positive feedback;
5. reconstitution of mental models by reinforcement learning and habituation;
6. immergence as positive feedback related to reinforcement learning and habituation; and
7. an enabling environment.

The first five of these features are discussed in [Hodgson and Knudsen \(2004\)](#) and various other parts of Hodgson’s work (cited above). The sixth is related more to the work of [Conte and Castelfranchi \(1995a\)](#), the EMIL Project Report, and various papers discussing immergence, e.g., [Andrighetto, Campenni and Conte \(2010\)](#), [Castellani \(2010\)](#), and [Conte et al \(2013\)](#); and the seventh is fairly common in the complexity sciences.

Nonetheless, the value of this framework arises from two sources: the combining of these various features; and the coherence of these features when considered together. These points are discussed in more detail in the concluding chapter.

Two (related) points are worth emphasising: first, this framework should not be viewed as a completed general theory of organic institutional emergence. It should be viewed as a contribution that moves the literature further in the direction of such a theory.

The second point is that in the simulations results based on the second model we observe an additional (eighth) feature related to selection pressure within a generalized Darwinian process<sup>26</sup>. This is discussed further in the concluding chapter.

Note that in the simulation results of both models, the organic institutions observed are consistent with Ferguson’s spontaneous order, as stated above.

Finally, the **fourth value added point** is that the models developed for this thesis add to the body of ABMs in IE. Most significantly, the simulations showed relationships between property rights, markets, and the division of labour within a stratified ontology. Furthermore, the code will be made freely available online after the completion of the PhD programme<sup>27</sup>.

---

<sup>26</sup>This is when the principles of variation, selection, and inheritance / durability are considered in a general sense (discussed in more detail in Chapter 2).

<sup>27</sup>See [https://github.com/gregjfisher/PhD\\_Code\\_Final](https://github.com/gregjfisher/PhD_Code_Final).

## Chapter 2

# Complexity Science and Complexity Economics

I think the next century will be the century of complexity.

– Stephen Hawking (January 2000)

This chapter provides overviews of the complexity sciences and its manifestation in the field of economics, Complexity Economics (CE). These descriptions provide the conceptual foundations on which the rest of the thesis is built.

The first section below (2.1) sets out a number of ‘defining features’ of complex systems, which are the main concern of the complexity sciences. These features describe the general characteristics of such systems, including agents, interdependence, and time.

Following this, Section 2.2 outlines some of the ‘conceptual features’ of the complexity sciences. These differ from defining features because they are concepts and principles that have either arisen out of the study of complex systems or that preceded the subject but are emphasised by it, e.g., ‘emergence’.

This second section focuses on concepts that are the most relevant to the research presented here and are described in a manner that dovetails with the material discussed in later chapters. The section is not exhaustive: a sub-set of the concepts emphasised in the complexity sciences are included.

The third section (2.3) focuses on CE. We must be careful when mapping from an abstract conceptual framework that has arisen mostly in the natural sciences, to the social sciences “without a rigorous process of testing for appropriateness and relevance.” (Mitleton-Kelly, 2003, p. 25). For example, an important difference between complex

systems involving simple, myopic agents and those with humans is that agents in the former are often reactive whereas those in the latter have a greater capacity for imagination and anticipation.

This section describes the overall approach taken in CE and it lists eleven key characteristics of that perspective, e.g., the idea of a stratified ontology owing to emergent properties.

Each of these eleven characteristics are discussed and contrasted with the equivalent approach taken in Neoclassical Economics (NCE), which has been the dominant school of thought in economics for decades.

This third section ends by contrasting CE and Neoclassical theory in a particular way: by focusing on the different internal consistencies (between their micro and macro features) of the two approaches.

Section 2.4 considers and responds to some of the criticisms of the complexity sciences, including CE; and Section 2.5 concludes.

As a final point, this chapter refers to several researchers in the complexity sciences but it places a greater weight on the work of W. Brian Arthur. Arthur was one of the first economists to consider the implications of the study of complex systems not only for economics but for the social sciences more generally. He led the economics department of the Santa Fe Institute for several years after its inception and has been a leading thinker in this field since<sup>1</sup>.

## 2.1 The Complexity Sciences: Defining Features

When describing or defining the complexity sciences we are faced with three problems: (i) it is a relatively new subject; (ii) it is studied by researchers in different disciplines; and (iii) it has evolved since it emerged in *circa* the 1970s. These issues make the field somewhat difficult to define.

Moreover, Arthur (2013) describes the complexity sciences not as a distinct field or academic discipline but as a *movement* in academia. This seems apt because it is used by - and applied to - a wide variety of disciplines, e.g., quantum mechanics, chemistry, biology, evolution studies, and the social sciences.

---

<sup>1</sup>It is perhaps worth noting here that Hayek wrote a paper (Hayek, 1967) which is broadly compatible with the description of CE in this chapter. This paper is discussed in more detail in Chapter 4.

Despite the difficulties of defining the complexity sciences, what follows is a summary of what are typically viewed as the main features of complex systems at the time of writing. The defining features are organised in to three themes, which are discussed in more detail in the first three sub-sections below:

1. Multiple (heterogeneous) agents;
2. Interdependence; and
3. Time.

### 2.1.1 Multiple Heterogeneous Agents

Complex systems typically contain ‘agents’ as the primary focus of analysis, e.g., sub-atomic particles in quantum mechanics, planets within solar systems, stars in their constitutive galaxies, ants in ant colonies, and people in social systems. Moreover, complex systems typically contain multiple agents that are often viewed as heterogeneous, i.e., where each contains unique features<sup>2</sup>.

Furthermore, these agents can themselves be viewed as: (i) complex systems of nested complex systems; (ii) having semi-permeable boundaries; (iii) using internal models to make sense of their reality; and (iv) existing in broadly decentralised systems.

#### 2.1.1.1 Nested Complex Systems

It is easy to interpret the above description as implying agents in complex systems are ‘atomized’ (in the philosophical sense). In Greek Philosophy *atomism* was the idea that atoms are the most basic and indivisible particles that make up the whole universe.

Equivalently, in sociology, atomism refers to “the tendency for society to be made up of a collection of self-interested and largely self-sufficient individuals, operating as separate atoms.” (Heyward, 2015, Glossary).

Either of these descriptions of atomism would lead us to believe that agents: (i) were simple, single points of reference; and (ii) have impermeable boundaries. However, this thesis takes the view that agents can themselves be complex systems with semi-permeable boundaries that render them open to flows of matter, energy, and information.

The idea that agents are themselves complex systems is not difficult to imagine if we consider organisms like ants and people. The word ‘organism’ implies a system of organs and these can be viewed as complex because they fit the features described in this sub-section. Furthermore, organs themselves can be viewed as complex systems, and this

---

<sup>2</sup>It is possible to imagine complex systems of homogenous agents but, in general, agents in complex systems are viewed as heterogeneous.

can continue ‘downwards’ until we consider atoms as complex systems of subatomic particles. All of this presents a picture of complex systems of *nested* complex systems (this is related to the idea of a *stratified ontology* which is discussed further below).

When modelling (cognitively or formally) these systems, abstraction often requires simplification so agents might be presented in an abbreviated way. With complex systems of nested complex systems, this must be done carefully.

### 2.1.1.2 Semi-Permeable Boundaries

The atomistic view is also criticised here because of its association with impermeable boundaries: agents are viewed as isolated entities. [Holland \(2014\)](#) explores the issue of boundaries in complex systems, referring to the idea of ‘semi-permeable’ boundaries. For agents, this means their boundaries would be impermeable enough for an agent to be identifiable but permeable (open) enough for flows to occur between the agent and its environment (including other agents).

Consistent with [Holland \(2014\)](#), this thesis assumes that agents are semi-permeable in nature, notably with regard to flows of information. This is especially important in the context of agents’ mental models co-adapting to each other via interaction.

### 2.1.1.3 Internal Models

Related to information flows, a number of researchers in the complexity sciences (e.g., [Holland et al, 1986](#), and [Arthur, 1994](#)) have emphasised the idea that agents use internal models to make sense of their environments and to make decisions by processing information in a structured way.

This idea that agents use internal models is of course not unique to the complexity sciences, e.g., in orthodox economics, firms seek to maximize profit and people seek to maximize utility. Maximization algorithms like this, which are forms of substantive rationality, are equivalent to internal models where the agents are firms or people.

Internal models can be highly simplistic in nature, e.g., cellular automata that switch between 0 and 1 depending on the state of their neighbours; or they can be highly complex, e.g., containing large neural networks. Indeed, we can think of neural networks as complex systems themselves - this brings us back to the idea of nested complex systems.

It is helpful to distinguish between simpler internal models that are generally *reactive* to information and events (information is collated and then some decision made); and sophisticated internal models that include *anticipation* of potential events, i.e., the imagination of future scenarios. The cognitive ‘depth’ of human agents leads to particular

challenges in human social systems, e.g., mutual contingency in which people try to anticipate the actions of others who are, in turn, anticipating them.

Finally, to be clear about nomenclature, in this thesis the phrase “internal models” is used to refer to all forms of agents, including non-biological types like atoms; whereas the phrase “mental models” is used for biological organisms, including humans.

#### 2.1.1.4 Decentralised Control / Bottom-Up Focus

Arthur (2013) referred to the economy as a “parallel system of concurrent behavior” (p. 2). This description can be applied to complex systems in general: the idea that there is no master controller and that agents make decisions synchronously. Control of the system - if there is such a thing - is devolved to the agents.

This does not imply, however, that agents are equal with respect to the influence they have vis-à-vis other agents or the system as a whole. This relates to the concept of power, which is defined here as the ability to influence future events.

Furthermore, this emphasis on local behaviour does not mean that patterns and structures do not emerge within the system. Indeed, this thesis focuses on structures that immerse within the agents’ mental models.

Given this focus on devolved behaviour, researchers in the complexity sciences tend to focus their attention at the local level when trying to understand complex systems. The ‘bottom-up’ approach to making sense of complex systems is often emphasised in the complexity sciences, e.g., Epstein (2012) refers to “growing” phenomena of interest via computational modelling.

### 2.1.2 Interdependence

If the agents in complex systems are equivalent to (semi-permeable) nodes in a network, here we focus on the links in that network, i.e., the connections between the agents.

In complex systems, agents typically have relationships with other agents, e.g., parents and children, planets in the same solar system, and electrons in the same atom. If we imagine each agent as having a set of bilateral relationships, we can appreciate that multiple agents will exist in a larger network of ties.

Clearly, this picture of a network with nodes and links is highly abstract - agents and how they relate to each other will vary considerably between types of complex system. For example, the relationship between planets in the same solar system will be mostly gravitational in nature; the relationship between a predator and its prey will be one of hunting and avoidance; and the relationship between parents and children is often emotional, with children being initially dependent on their parents.

Regarding the relations between agents, the nature of mental models is important here also. When agents have simple (non-anticipatory) mental models, the ties between the agents might simply reflect one agent affecting another on a systematic basis. The effected agent might then respond in some way that then affects the agents it is linked to. Here, interdependence will be *reactive* in nature (the term *reactive agents* is used here to denote the actors in such networks). For an example of this type of phenomena and how change can propagate across entire networks, see [Watts \(2002\)](#).

By contrast, sophisticated mental models might include perceptions of other agents, e.g., expectations of other agents' behaviour in particular circumstances. This awareness will help agents to anticipate the actions of others (the term *anticipatory agents* is used here to denote such actors). Therefore, we can see that with human agents, agent interdependence can be anticipatory in nature and not only reactive.

### 2.1.3 Time

The third defining feature of complex systems is that they are dynamic. There are a number of related concepts worth mentioning here: adaptation; evolution and co-evolution; and path-dependence.

In complex systems, adaptation is the characteristic that an agent will adjust to its environment, e.g., by reacting to others' actions. An example would be some prey avoiding a predator.

The concepts of evolution and co-evolution are discussed in more detail in [Section 2.3.3.5](#) below. Evolution is viewed as including the processes of variation, inheritance, and selection; and co-evolution is interpreted as reciprocal changes between two or more groups of agents.

In terms of history, an important concept in the complexity sciences is *path-dependence*. This is the idea that the current *state* of a system reflects its past, and that its current state will influence its future. Hence, the past, present and future of the system are inherently connected. See [David \(2007\)](#) for a detailed discussion of this concept and its use (and misuse) in economics.

The mapping of the past in to the present can occur in two broad ways, via: (i) agents' mental models (either because of memories or because knowledge / learning from previous experiences is carried forward in some way); and/or (ii) the relationships between the agents. Both of these can be understood as historical patterns brought forward from the past in to the present.



We can infer from this that some complex systems will not exhibit path dependence. This will be true if agents have no memories and/or they do not learn; and/or if any form of relationship between the agents is not carried forward over time.

## 2.2 The Complexity Sciences: Conceptual Features

The previous section looked at the three main defining features of complex systems: agents, interdependence, and time. In this section we focus on analytical features. As mentioned previously, these are features that are emphasised by the complexity sciences but either pre-date the movement or have arisen from the study of such systems. Furthermore, the list includes those features that are most relevant to this thesis. An open-source and more extensive list can be found on-line at [Santa Fe Institute \(2023\)](#).

### 2.2.1 Feedback

Feedback is defined in The Oxford English Dictionary (OED) as “The modification, adjustment, or control of a process or system (as a social situation or a biological mechanism) by a result or effect of the process.” (OED, 2021, Definition 2).

In the complexity sciences, there are generally two types of feedback: positive and negative.

Positive feedback is viewed as amplifying change, e.g., if a variable related to an agent (say, temperature) is increased then positive feedback will further increase that variable. By contrast, negative feedback is seen as moderating change, e.g., if the same variable increases then negative feedback will reduce it, and vice versa.

### 2.2.2 Non-Linearity

Non-linear is defined in the OED as “involving or possessing the property that the magnitude of an effect or output is not linearly or proportionally related to that of the cause or input.” (OED, 2021, Definition 1a).

With linear systems, the impact ( $\Delta y$ ) of any perturbation ( $\Delta x$ ) will be proportional to the magnitude of the perturbation, i.e.,  $\Delta y = \beta \Delta x$  where  $\beta$  is a constant<sup>3</sup>. By contrast, non-linear systems will see perturbation having a disproportional (large or small) effect, e.g.,  $\Delta y = (\Delta x)^\beta$  when  $\beta \neq 1$ .

---

<sup>3</sup>and  $\beta \neq 0$ .

### 2.2.3 Space of Possibilities

In the complexity sciences (and elsewhere) it is common to consider the ‘state space’ of a system, which is the complete set of *potential* states the system could be in. For our purposes this will include every possible location and trajectory of every agent in the system and all of the possible configurations of their mental models.

If we consider that a specific agent will have limited knowledge of the system in which it finds itself, the ‘space of possibilities’ for this agent will be the sub-set of the state-space *it can imagine*<sup>4</sup>. Moreover, an agent can improve its limited knowledge of the wider system by exploring it, e.g., by moving around a geographic region and/or interacting with other agents. This exploration would also increase the agent’s range of states it can imagine, hence the phrase used by some complexity scientists, of agents “exploring the space of possibilities.”

This will be important in the computational simulations when agents explore their environment: do institutions emerge during this exploration?

The idea of exploring the space of possibilities is discussed in [Mitleton-Kelly \(2003\)](#). The following quote gives us a sense of the challenging localised dynamics faced by agents in complex systems:

Complexity suggests that to survive and thrive an entity needs to explore its space of possibilities and to generate variety. Complexity also suggests that the search for a single ‘optimum’ strategy may neither be possible nor desirable. Any strategy can only be optimum under certain conditions, and when those conditions change, the strategy may no longer be optimal. ([Mitleton-Kelly, 2003](#), p. 14)

This quote also helps us understand that complex systems might never reach an optimum, nor an equilibrium: agents might have to continuously explore the system they occupy, allowing their mental models to adapt, and adjusting strategies as other agents adjust to theirs.

### 2.2.4 Far-From-Equilibrium

This is the idea that complex systems could be forced away from equilibrium conditions, or might never experience these conditions at all. [Chan \(2001\)](#) describes how this phenomenon was explored in [Nicolis and Prigogine \(1989\)](#):

---

<sup>4</sup>It is possible the agent will not be aware of many parts of the system so it will only be able to imagine a localised part of the sub-set of the whole state-space.

In 1989, Nicolis and Prigogine showed that ... [t]he “far from equilibrium” phenomenon illustrates how systems that are forced to explore their space of possibilities will create different structures and new patterns of relationships. (Chan, 2001, p 6)

We will explore the ideas of non-equilibrium and dis-equilibrium further in Section 2.3.3.2 below when we discuss CE. These are also related to non-ergodicity, which is the subject of Section 2.3.3.6.

## 2.2.5 Emergence

As stated in the Introduction, emergence is defined by the Santa Fe Institute (2023) as a “process by which a system of interacting subunits acquires qualitatively new properties that cannot be understood as the simple addition of their individual contributions.”

This feature of complex systems is essential for differentiating the complexity sciences from reductionist strategies for making sense of such systems. Classic examples of emergence are the properties of water from its constituent atoms (like ‘wetness’): these properties are not reducible to hydrogen or oxygen.

Hodgson (2000c) includes a helpful discussion of the history of the term ‘emergence’, noting that it was first suggested by the philosopher George Lewes in 1875. We can appreciate, therefore, that this term precedes the complexity sciences by about a century so while it is associated with that field, it is not attributable to it.

In discussing non-reducibility, Hodgson (2000c) quotes Morgan (1932) who explains:

the hypothesis is that when certain items of “stuff,” say  $o$   $p$   $q$ , enter into some relational organization  $R$  in unity of “substance,” the whole  $R(o p q)$  has some “properties” which could not be deduced from prior knowledge of the properties of  $o$ ,  $p$ , and  $q$  taken severally. (Morgan, 1932 as quoted in Hodgson, 2000c, p. 65)

Furthermore, in terms of emergence, Hodgson (2000c) quotes Morgan (1925):

the emphasis is not on the unfolding of something already in being but on the outspringing of something that has hitherto not been in being. It is in this sense only that the noun may carry the adjective “emergent.” (Morgan, 1925 as quoted in Hodgson, 2000c, p. 66)

Hodgson (2000c) also notes that interest in emergence “re-emerged” after the late 1960s, which coincided (approximately) with the formation of the complexity sciences.

Consistent with this, [Axtell \(2007\)](#) notes “[t]here is a large and growing literature on the idea of emergence in physical, biological, and social systems (cf. [Haken \(1987\)](#), [Baas \(1994\)](#), [Morowitz \(1998\)](#), [Howitt and Clower \(2000\)](#), [Johnson \(2001\)](#), [Sawyer \(2001\)](#), [Morowitz \(2002\)](#), [Sawyer \(2002\)](#)).” (p. 111).

Finally, it is common among researchers to discuss emergence as if it were external to the agents, e.g., [Arthur \(2013\)](#) refers to traffic jams as forms of emergence among driver agents. The ensuing traffic problem is a result of agents reacting to each other’s actions but the emergent pattern can be thought of as essentially external to each agent.

However, it is also possible for emergence to involve the adjustment of mental models, i.e., for the emergent pattern to be internal to the agents. As mentioned in the Introduction, [Conte and Castelfranchi \(1995a\)](#) refer to ‘cognitive emergence’, which is related to ‘immergence’ in the social simulation literature. Once again, if we think of an agent (including its mental model) as a complex system, we can imagine a pattern of behaviour having emerged from the agent’s interaction with other agents. This idea of immergence being endogenous ‘within’ agents’ mental models is explored further in this thesis.

## 2.2.6 Self-Organisation

This concept is a type of emergence: self-organisation is the process in which patterns emerge spontaneously in a way that is helpful to the agents, e.g., by enhancing their likelihood of survival. An example would be the practice of wolves hunting in packs. This skill is not reducible to individual wolves because it requires coordination between them but, also, it is not achieved through some kind of centralised command and control process. It arises from agent-to-agent interaction. Furthermore, it enhances the ability of the pack to survive because hunting is more effective this way.

We should note here that [Sheehan and Wahrman \(2015\)](#) date the origins of this concept to the Mississippi and South Sea financial bubble of 1719-1720 so it clearly precedes the complexity sciences. Furthermore, as mentioned in the Introduction, the related (but not identical) idea of ‘spontaneous order’ pre-dates that of self-organisation by millennia<sup>5</sup> (even if the phrase was coined in the 20th Century).

The complexity scientist most known for promoting ‘self-organisation’ is Stuart Kauffman. [Kauffman \(1993\)](#) argues that evolution by natural selection is not the only ‘force’ that determines the likelihood of a species surviving. Kauffman writes “[i]t is this single force view which I believe to be inadequate, for it fails to notice, fails to stress, fails to incorporate the possibility that simple and complex systems exhibit order spontaneously.” (p. xiii).

---

<sup>5</sup>Recall from the Introduction that the earliest known reference is from the fourth century BC Chinese Philosopher Zhuang Zhou ([Hamowy, 1987](#), p. 6).

Another way of describing self-organisation is that it is the endogenous creation of new order in some system. Indeed, a useful way of understanding self-organisation is to contrast it with order imposed from outside of that system.

Other examples include swarming behaviour of animals such as birds, insects, and fish; and communication within human systems: language can only have emerged endogenously, it could not have been imposed from outside. [Luhmann \(1995\)](#) emphasises this point, noting also that a social system can only reproduce itself if there is communication between people.

Moreover, this thesis is concerned with the emergence of organic institutions as types of ‘spontaneous order’. This is closely related to the concept of self-organisation within complex systems, and is discussed in greater detail in chapters [3](#) and [4](#).

Finally, while the idea of order emerging spontaneously in complex systems is helpful, this thesis attempts to be more concrete in identifying mechanisms by which organic institutions emerge and immerse in complex economic systems.

### 2.2.7 Lock-In

This idea is related to the concept of path dependence, which was mentioned in Section [2.1.3](#) above. It refers to some pattern of behaviour that is maintained in some system through endogenous factors.

The classic (though contested) example of lock-in is the QWERTY keyboard. The narrative goes that this design was originally intended to slow typists down - engineers thought that mixing the letters in a relatively random way would achieve this task. However, typists became familiar (and efficient) with the new layout so manufacturers sold keyboards with this design. In turn, new users of this layout learned to type with the same design. Over time, the QWERTY keyboard became ‘locked in’.

[David \(2007\)](#) provides additional examples of what he believes are “sub-optimal equilibria” that have been “‘selected’ by a dynamic process” (p. 130). He refers to “640K lower memory in the IBM PC, AC vs. DC electrical current, light-water reactors, and VCR formats” (ibid, p. 136).

More broadly, [Foxon \(2002\)](#) discusses how in economics increasing returns to scale can ensure certain technologies can be locked in place, citing four ‘forces’ as developed in [Arthur \(1994b\)](#): scale economies, learning effects, adaptive expectations, and network economies. Foxon also mentions North’s references to institutional lock-in: the same four forces can sustain an institution within a population.

To understand the points of contention here, let us first differentiate between the lock-in of behavioural patterns in general and a sub-group of such patterns that lead to

sub-optimal outcomes. [Liebowitz and Margolis \(1990\)](#), for example, do not dispute the idea that QWERTY has become locked in; rather they dispute the inefficiency aspect, as argued by David. They state that “the continued use of Qwerty is efficient given current understanding of keyboard design” (p. 2); and that the “trap constituted by [this] obsolete standard may be quite fragile” (p. 21), i.e., the lock-in effect is relatively weak.

Clearly, the issue here is not so much about keyboard designs. Rather, there are important questions about whether pareto inferior outcomes can become locked in and, if so, how prevalent they are in economic systems. This is particularly important in free market economics because locked-in sub-optimal equilibria are examples of market failures that are durable by nature<sup>6</sup>.

This is not the place to answer the above question of the prevalence of sub-optimal ‘traps’. Instead, we can simply state that, from a complexity sciences point of view, we should keep an open mind to such phenomena. Moreover, [David’s \(2007\)](#) framing of the issues at play here is particularly helpful so let us briefly look at this.

David refers to positive feedback effects that move agents into a “trapping region - the basin of attraction that surrounds a locally (or globally) stable equilibrium.” ([David, 2007](#), p. 131). The problems of “escaping from lock-in of the system to a globally inferior (but locally stable) attractor are rooted in ‘pure’ coordination costs.” (ibid, p. 132)<sup>7</sup>. The sides of the attractor basin (steepness and length) correspond to the stability of the locked-in phenomenon and how difficult (or expensive) it is to escape. This gives rise to context-specific practical problems and questions about the net present value of the costs and benefits of changing to a globally preferred outcome, if this is even feasible.

However, it is important to consider that such “trapping regions” do not necessarily persist. In [Section 2.3.3.6](#) below we discuss how economies can be re-patterned over time, which means that following “‘exogenous [to the attractor basin] innovations’ (in the state of relevant knowledge, or in the regulatory institutional regime), the previous attractor(s) may be destroyed, freeing the system to endogenously begin to evolve some new configurations.” (ibid, p. 132). [David \(2007\)](#) gives the example of deregulation in the US’s telecommunications networks in the 1950s, which “formed new ‘attractive paths’ ... for the evolution of digital telecommunications technologies.” (p. 133).

In addition, David argues that by being aware of positive feedback effects that result in sub-optimal trapping regions, we can take pre-emptive, ameliorative actions if we observe them playing out. This can be linked to the discussion in the Introduction of

<sup>6</sup>This very much relates to [Luban’s \(2020\)](#) reference to spontaneous processes giving rise to positive and negative outcomes - this is discussed further in [Chapter 3](#).

<sup>7</sup>Note that this framing corresponds to [Schultz’s \(2001\)](#) coordination situations and not his collective action situations.

Institutional Economics being used to advise on present day institutional challenges such those related to cryptocurrencies.

The final point to note here is that [David \(2007\)](#) uses his framing of lock-in to describe the state of orthodox economics itself. “For many economists, their own sunk cost in mastering that discipline have produced a facility for reasoning that suppresses natural, human intuitions about historical causation.” (p. 138). He refers to this as “intellectual sunk cost hysteresis” (*ibid*), which is analogous to an attractor basin. This thesis can be interpreted as a contribution to attempts to move out of that trapping region.

### 2.2.8 State of Paradox

This is the idea that complex systems can exist in states that exhibit seemingly contradictory phenomena. [Chan \(2001\)](#) writes “[t]his reinforces the idea of bounded instability or the edge of chaos that is characterised by a state of paradox: stability and instability, competition and cooperation, order and disorder.” (p. 6).

For this thesis, the paradox that is perhaps most relevant is that of order and disorder (recall from the Introduction that we adopt Hayek’s definition of order<sup>8</sup> in this thesis).

A simple example of the co-existence of order and disorder is a professional football league. The rules of the league are typically codified and known to all the teams, and games tend to be agreed months in advance. Furthermore, the ‘rules of the game’ are also generally known by the players.

However, individual football games tend to exhibit a considerable amount of disorder (even if the rules are followed perfectly): notably, a team’s tactics, as well as those of individual players, will generally seek to confound the expectations of their opponents. Furthermore, the tactics of both teams often co-evolve during a game. Hence, in professional football leagues we tend to observe a co-existence of order and disorder.

Looking at economics, Hayek viewed ‘the market order’ as the coexistence of order and disorder. Certain phenomena, notably entrepreneurial behaviour and ‘liberal legislation’, would support a broadly ordered system. However, within this ‘order’, phenomena like new technologies and changing consumer preferences would create a degree of disorder. Hayek’s market order, as well as [Vanberg’s \(1994b\)](#) reference to “conditional evolution”, are discussed further in Chapter 4 below. These help add to our understanding of order and disorder as a state of paradox.

In the simulations based on the Agent-Based Models developed for this thesis, a rudimentary economic system starts from a state of disorder in which no markets have

---

<sup>8</sup>See Section [1.4.3](#).

emerged (and property rights do not exist in the case of the second model)<sup>9</sup>. New organic institutions (markets and property rights) represent new order; however, we find that the agents, while competing for resources when foraging, also act in a beneficial and collective way when trading resources (after appropriate organic institutions have emerged). Hence, order co-exists with disorder in these simulations.

### 2.2.9 Sensitivity to Initial Conditions / Predictability

The idea that the future path of a complex system is sensitive to its initial conditions was famously observed by Edward Lorenz when he developed a set of differential equations in the 1960s to approximate the state of - and changes in - the weather.

This concept is related to non-linearity: Lorenz noticed that infinitesimally small changes in initial conditions meant that, over time, a system exhibits significantly different outcomes. Lorenz first conducted his experiment in 1960 where the initial conditions of his state variables were measured to 6 decimal places. When he repeated the experiment in 1961, these initial conditions were measured to 3 decimal places (see [New York Times, 2008](#)). Lorenz had expected the results of the two experiments to be essentially the same; however, they were not. Over a sufficiently long enough time horizon, the results were significant different.

Lorenz's experiments were conducted using (deterministic) differential equations; however, the same effect can also be observed in any complex system that exhibits non-linearity (e.g., due to positive feedback effects).

This idea that outcomes are sensitive to initial conditions when non-linearity exists is one factor that warns us about the ability to predict outcomes. If we are unsure about initial conditions and / or the patterns that describe some system, how confident should we be that we can predict the outcome?

### 2.2.10 Symmetry Breaking

As the name suggests, symmetry breaking occurs when a type of equilibrium, balance, or symmetry in some system is changed or broken.

In this thesis we will use the following definition, which appears in [Castellani \(2010\)](#):

...the process by means of which the considered symmetry is broken and is therefore usually ascribed a “dynamic” character in the literature (in

---

<sup>9</sup>Note, however, that certain institutions like language are assumed. As stated in the Introduction, the focus of the models is on the emergence of *new* institutions: we do not assume that we start from an institution-free state of nature.



contrast with the “static” character attributed to a situation of symmetry).  
([Castellani, 2010](#), p. 321-322)

Now that we have discussed the defining features and some of the analytical features of complex systems, let us turn our attention to CE.

## 2.3 Complexity Economics

CE is the study of economic behaviour from a complexity science perspective. For books and papers written on this general approach, see, for example, [Rosser \(1999\)](#), [Colander \(2000\)](#), [Judd \(2006\)](#), [Teshfatsion \(2006\)](#), [Epstein \(2012\)](#), [Axtell \(2007\)](#), [Colander et al \(2009\)](#), [Kirman \(2011\)](#), [Colander and Kupers \(2012\)](#), [Farmer \(2012\)](#), [Gallegati and Kirman \(2012\)](#), [Arthur \(2013\)](#), and [Arthur \(2021\)](#).

In the first sub-section below we ask whether the complexity sciences are applicable to economics. The second sub-section discusses the overall approach taken in CE, focusing on eleven general principles (for each we contrast the approach taken in Neoclassical theory).

### 2.3.1 Should Economies be Viewed as Complex Systems?

One way of framing this question is to ask whether economic systems include the three defining features listed in [Section 2.1](#): multiple heterogeneous agents, interconnectivity, and time.

It should be clear from the discussions in that section that economies can indeed be described by these three features: people, firms, and other organisations are typically thought of as economic agents; they exist in networks of relationships; and, in general, they evolve and adapt to each other over time.

From a logical point of view, the statement that economies should be thought of as complex systems represents a syllogism: (i) systems with the three defining features are defined as complex systems (the major premise); (ii) economies include these defining features (the minor premise); therefore, (iii) economies are complex systems.

It would be a mistake, however, to apply the conceptual features of the complexity sciences to economics in an unquestioning way because the former emerged, broadly speaking, from the natural sciences. One must be cautious when mapping from one domain to another.

Most notably, the idea that human agents are different to agents in natural complex systems (atoms, stars, bees, etc.) was mentioned several times in [sections 2.1 and 2.2](#) above, especially the idea that “humans have the capacity to reflect and to make

deliberate choices and decisions among alternative paths of actions.” (Mitleton-Kelly, 2003, p. 25). Moreover, the fact that a ground-up view of complex systems is generally advocated by complexity scientists means that if human agents have features that are different to agents in natural systems then it is likely that whole systems of these agents will too: different patterns might emerge.

### 2.3.2 Dealing with Mapping and Other Challenges

There are a variety of ways in which this mapping challenge could be handled in a thesis but it would be sensible to consider any other related challenges at the same time.

One such challenge is to articulate clearly the nature of CE given this is an important part of the thesis. Another is to set out in detail some fundamental principles that will act as premises in arguments made later on. Yet another challenge is to compare and contrast a CE approach to that of Neoclassical theory.

A good solution to all of these challenges is to describe a list of principles that together make up what is defined here as CE. Some of these principles help to overcome the mapping challenge because they represent what the complexity sciences *mean* in the context of economics, e.g., being open to a stratified ontology. Others represent premises that will be used in subsequent arguments.

### 2.3.3 The Eleven Principles of Complexity Economics

In this sub-section we set out eleven general principles that will help us understand better what a CE approach means in this thesis.

These principles are:

1. Economies are computational in nature;
2. Economies exhibit non-equilibrium and dis-equilibrium features;
3. Economies include the processes of *formation* and *allocation*;
4. Economies have a stratified ontology, resulting from emergence;
5. CE is open to a ‘generalized Darwinian’ approach;
6. Economies are non-ergodic;
7. Uncertainty is an important feature in economic behaviour;
8. Agents use mental models to make decisions;
9. CE is open to inter-disciplinarity;

10. CE values realism and is sympathetic to instrumentalism; and
11. CE has an affinity with pragmatism.

Each principle is discussed below and then contrasted with the Neoclassical approach. For general criticisms of Neoclassical theory see, e.g., [Eichner and Kregel \(1975\)](#), [Nelson \(2001\)](#), [Stiglitz \(2001\)](#), [Blaug \(2002\)](#), [Arnsperger and Varoufakis \(2006\)](#), and [Colander et al \(2009\)](#).

Before proceeding, it is worth noting that Complexity Economics is a maturing field of study, which means that different interpretations and emphases are bound to exist. What follows is one interpretation although it is probably reasonable to say that the only possibly contentious principle listed below is the last, that related to pragmatism. The argument here is that this seems to be a natural extension of the bottom-up orientation of the complexity sciences, and it appears to be implicit in much of Arthur's work, e.g., [Arthur \(2013\)](#).

### 2.3.3.1 Economies are computational in nature

In response to the stable, equilibrium-focused world of Neoclassical theory, Arthur wrote:

A better way forward is to observe that in the economy, current circumstances form the conditions that will determine what comes next. The economy is a system whose elements are constantly updating their behavior based on the present situation. To state this in another way, formally, we can say that the economy is an ongoing *computation* — a vast, distributed, massively parallel, stochastic one. Viewed this way, the economy becomes a system that evolves procedurally in a series of events; it becomes algorithmic. [Emphasis included] ([Arthur, 2013](#), p. 6)

This quote raises the question of how we define computation, which is an open question even in the field of computational science.

[Denning \(2010\)](#) discusses the history of definitions of computation and offered a “Transformation of Representations” version, which appears compatible with the general usage of the term at the time of writing.

A representation is defined by Denning as “a pattern of symbols that stands for something” ([Denning, 2010](#), p. 7). His proposed definition of computation is of an *information process* driven by a representation, i.e., information is transformed in the act of computation.

In economics we can think about economic agents who receive information that is then processed (transformed) by their mental models. The output might be information also

(e.g., a signal to another agent) or actions performed by the agent (including forms of work). These agents need not be humans - they might also be machines, firms, trade unions, etc.

The idea that economies are parallel systems of concurrent computation is compatible with a complex systems approach because it fits well with the architecture of multiple agents interacting over time. In fact, computation provides a framework for describing how agents in the system process information, which is based on a mature body of work.

#### *Computation and Closed Form Solutions*

In discussing computation, [Arthur \(2013\)](#) refers to Alan Turing's halting problem theorem ([Turing, 1936](#)). The result of this theorem is that:

...there is no analytical method to decide in advance what a given algorithm will do. All we can do is follow the computation and see what it brings. Of course, with simple algorithms we can often see they will settle down to a given outcome. But algorithms don't have to be particularly complicated before we cannot decide their outcomes ([Wolfram, 2002](#), as quoted in [Arthur, 2013](#), p. 6).

This point has substantial implications for the analysis and modelling of computational processes, which Arthur argued included economies. It means that, in general, we must be sceptical about closed form solutions of such processes; rather we must "follow the computation and see what it brings" (ibid).

#### *Implications for Economics*

There are two broad implications worth highlighting here. First, when framing or modelling economic systems, we must account for their computational nature. If we fail to do this then we run the risk of a framing error.

Indeed, this point seems particularly important for the economics profession which has become highly mathematized in recent decades. A mathematical framing of computational systems represents information compression, which Turing warned against.

Second, the halting problem has significant implications for our understanding of how agents process information, including questions about 'rationality'. We must ask ourselves how agents process information and achieve any goals if the environment in which they exist is not 'solvable' in any way.

*Comparison with Neoclassical Theory*

Historically, NCE has not treated economic systems as computational in nature. There is a substantial literature looking at the formal approach taken by Neoclassical theory, e.g., [Arthur \(2013\)](#) discusses this and how it relates to equilibrium conditions.

[Arthur \(2013\)](#) argues that for certain features of the economy, equilibrium-like conditions might well exist. Furthermore, “[w]e can often do much useful pre-analysis of the qualitative properties of nonequilibrium systems” (ibid, p. 6). As a result, the standard analytical techniques in Neoclassical theory can be useful. However, for

highly interconnected systems, equilibrium and closed-form solutions are not the default outcomes; if they exist they require justification. And computation for such systems should not be regarded as the avoidance of analytical thinking; rigorously speaking, it may be completely necessary. ([Arthur, 2013](#), p. 6)

Here, Arthur is arguing that if the economy is a computational process then computational modelling is a suitable technique for analysing such systems. The standard mathematical techniques of Neoclassical should be viewed as approximations that are useful only under certain conditions. To reinforce this point, Arthur writes “computation ... allows us to see phenomena that equilibrium mathematics does not.” (ibid, p. 8).

This point should also lead us to query the assumption of substantive rationality in computational systems: historically, NCE has been oriented around agents solving a utility maximization problem under known conditions. The halting problem means that in computational systems, this type of maximization is questionable.

### **2.3.3.2 Economies exhibit non-equilibrium and dis-equilibrium features**

The computational view of economics described above helps us appreciate how the economy, or parts of it, can exist in states of non- or dis-equilibrium. This perspective “shows us an economy perpetually inventing itself, perpetually creating possibilities for exploitation, perpetually open to response. An economy that is not dead, static, timeless, and perfect, but one that is alive, ever-changing, organic, and full of messy vitality.” ([Arthur, 2013](#), p. 19).

This dis-equilibrium view was emphasised in Farmer’s (appropriately entitled) paper, ‘Economics Needs to Treat the Economy as a Complex System’:

...in many situations there is no unique equilibrium. When there are multiple equilibria it may be difficult to predict which agents will converge to; in other circumstances they may fail to converge to any equilibrium at all. (Farmer, 2012, p. 8)

### *Comparison with Neoclassical Theory*

Arnsperger and Varoufakis (2006) argue that NCE was comprised of three underlying axioms. One of these was ‘methodological equilibration’, which is the idea that markets (and economies more generally) tend to equilibrium.

There is an important question here about why Neoclassical theory has taken this approach. One explanation, shared by many of NCE’s critics, is due to the analytical techniques available to economists for most of the 20th Century that for the most part were mathematical in nature. Economists looking to formally determine how economies worked, who were restricted to these analytical techniques, built models orientated around equilibrium outcomes and microeconomic behaviour that was consistent with equilibrium. This approach has been taken in various parts of Neoclassical theory, including the theory of the firm, international trade, and financial markets (Arthur, 2013, p. 2).

It should be clear, however, that the development of computational approaches and the acceleration in processing power since the 1970s, both mean that these restrictive mathematical techniques are no longer the only tools available to economists. Computational modelling of economies offers a way of increasing our understanding of non- and dis-equilibrium phenomena.

### **2.3.3.3 Economies include the processes of *formation* and *allocation***

Economics is often defined as the study of the *allocation* of scarce resources. Moreover, Neoclassical theory has focused heavily on this issue to the extent that Arthur (2013) writes that of the two great problems in economics, “[t]he allocation problem is well understood and highly mathematized.” (p. 17).

The second major problem in economics Arthur refers to is about *formation* within the economy. It is tempting to think of this in a conventional way, as being about stocks of assets, including capital, labour, and natural resources. However, Arthur argues that the issue is more about how the economy structures and re-structures itself over time. Formation is about:

how an economy emerges in the first place, and grows and changes structurally over time. This is represented by the ideas about innovation, economic development, structural change, and the role of history, institutions, and governance in the economy. (Arthur, 2013, p. 17)

This reference to ‘economic patterning’ mimics the language used by other complexity scientists. It is often used to refer to the relationships between parts of systems: the network of relationships in an organisation, circuit board designs, how the various parts and processes of an organism work to sustain its life, etc. (e.g. Boulton, 2019).

Moreover, if we stand back and consider the economy as both formed / re-forming patterns *and* allocation, this helps us visualise the economy as both a stock of networked (or patterned) resources *and* as a flow of resources across that network.

When it comes to understanding what types of pattern exist in the economy, Arthur (2013) refers to “institutions, arrangements, and technological innovations.” (p. 1). Institutions include social norms; arrangements can refer to a variety of phenomena, including management structures in organisations; and technological innovations are about how materials and energy are arranged to achieve tasks. Of these three patterns, Arthur’s own research has been focused on technology, e.g., Arthur (2009).

#### *Comparison with Neoclassical Theory*

Arthur’s view is that Neoclassical theory has focused on the allocation problem at the expense of understanding formation.

If we consider the formal modelling technology on which Neoclassical theory has relied, this is not surprising. Formation is not simply about the accumulation of stocks of resources, which can be handled by differential equations; it is about two important issues: the *creation* of networks of relationships and how matter, energy, and information are *processed* across those networks.

The standard analytical techniques of Neoclassical theory, which are largely based on classical physics, are not well suited to modelling patterned phenomena. However, they can be modelled computationally, including via agent based modelling. This approach has been used for decades in the complexity sciences to model patterns of relationships and how resources and information can be processed across such networks.

#### **2.3.3.4 Economies have a stratified ontology**

This principle arises from the concept of emergence, which was discussed in Section 2.2.5.

The reference to ‘stratified’ comes from the idea that system-wide phenomena can emerge from the interaction of the system’s parts. These system-wide phenomena are framed as if they are ‘above’ the parts / agents in some hierarchy, hence ‘stratified’.

An important point that distinguishes the complexity sciences, which are open to emergent properties, and reductionist perspectives (‘reductionism’ is defined below) is the view of the former that accords a different ontological status to emergent properties. For complexity scientists, emergent properties are ontologically ‘new’ and the properties of these phenomena are not reducible to the parts / agents from which they arose. Furthermore, the ontological ‘newness’ of emergent properties allows for ‘downward effects’ whereby these properties act on the very agents that gave rise to them.

A classic way of describing this stratification in the complexity sciences is to point out that atoms emerge from the interaction of sub-atomic particles; chemicals emerge from the interaction of atoms; biological entities emerge from the interaction of chemicals; and so on, until the properties of galactic superclusters emerge from the interaction of their constituent galaxies.

Researchers who take a reductionist approach to analysis can also recognise a type of stratification in the sense groups of agents can give rise to system-wide properties. The difference, however, is that these properties are not seen as ontologically ‘new’ - they are reducible to the parts / agents of the system.

For the purpose of this thesis, we are most interested in how social phenomena - like organic institutions - emerge from the behaviour of ‘individual’ agents. In addition, we look at how the division of labour can be viewed as an emergent phenomenon.

One point worth emphasising is that the claim here is not that all phenomena must be explained through the concept of emergence. There are a wide variety of phenomena that are explainable through reductionism, e.g., if we wish to understand the trajectories of balls on a snooker table; or if we wish to model a satellite in orbit around the Earth. A stratified ontology means we are open to the idea of emergence; it does not mean we reject all reductionist explanations.

### *Reductionism*

In this thesis we adopt a definition of reductionism as the “practice of describing or explaining a complex (esp. mental, social, or biological) phenomenon in terms of relatively simple or fundamental concepts, especially when this is said to provide a sufficient description or explanation...” (Oxford English Dictionary, 2023).

This is a relatively simple definition which is sufficient for our purposes here. The subject is, however, much deeper than this definition implies and it is an important topic in the



philosophy of science. See, for example, the discussions of ontological, methodological, and epistemic reductionism in the [Stanford Encyclopedia of Philosophy \(2023\)](#).

### *Comparison with Neoclassical Theory*

Neoclassical economists have historically taken a reductionist stance while also acknowledging that “social variables, not attached to particular individuals, are essential in studying the economy or any other social system” ([Arrow, 1994](#), p. 8).

[Arnsperger and Varoufakis \(2006\)](#) point to NCE’s reductionist leanings in the first of the three axioms they associate with Neoclassical theory:

To the neoclassical economist ... individual agents ... are to be studied, like the watchmaker’s cogs and wheels, independently of the social whole their actions help bring about. ([Arnsperger and Varoufakis, 2006](#), p. 8)

This axiom is referred to as *Methodological Individualism* in [Arnsperger and Varoufakis’](#) paper but it also corresponds to the definition of reductionism stated above.

NCE has followed a reductionist strategy in a number of ways. First, the interaction of agents in Neoclassical models is typically very limited; and, second, the mental models of the agents have historically been restricted to substantive rationality and exogenous preferences. These two features are typically combined within a mechanistic framing (and mathematics) that exist at the same ontological ‘level’: there are no emergent properties and the agents behave in a deterministic way.

We should ask, in the context of the quote from [Arrow \(1994\)](#) above, how Neoclassical theory explains social structures if, as Kenneth Arrow states, “such categories are in fact used in economic analysis all the time.” (p. 1). In this paper, Arrow uses the examples of prices in general equilibrium analysis and the rules assumed in game theoretic analyses to support his point.

Game theory, which [Arrow \(1994\)](#) describes as “[t]he current formulation of methodical individualism” (p. 4), has been used by Neoclassical economists to model how organic institutions arise in economics. [Hodgson \(2007\)](#) points to [Field \(1979, 1981, 1984\)](#) as providing crucial arguments for why such explanations are insufficient. This important subject is discussed further in [Chapter 5](#) where criticisms of game theoretic models are discussed.

### 2.3.3.5 Complexity Economics is open to ‘generalized Darwinism’

In this thesis, ‘Darwinian’ is defined by the three features of *variation*, *selection*, and *durability*<sup>10</sup>.

There is a substantial literature concerned with whether Darwin’s theory of evolution by natural selection applies to social systems, and a comprehensive survey of it is beyond the scope of this thesis. Nonetheless, an important question in this literature is whether or not to interpret the above three principles in a very narrow sense, where variation is due to random genetic mutation; selection is about the survival of an organism in some environment; and durability is viewed as inheritance from parent to offspring. If we accept this interpretation then we would clearly reject the use of Darwinian evolution, so defined, for explaining social phenomena.

However, these three principles can be generalized in two broad ways (consistent with “generalized Darwinism” as discussed in [Hodgson and Knudsen, 2006](#)): regarding abstraction in the meaning of the three principles, and the ‘level’ of any Darwinian explanation.

In terms of the first generalization, the three principles can be interpreted as follows:

- variation includes adaptation on behalf of some agent for whatever reason, including via reinforcement learning or conscious deliberation;
- selection refers to “a process of sifting and preservation of fortuitous adaptations.” ([Hodgson, 2003a](#), p. 89); and
- durability is related to the concept of path dependence, which ensures “that much of the pattern and variety is passed on from one period to the next.” ([Hodgson, 2003a](#), p. 89).

Consistent with this higher level of abstraction, [Hodgson and Knudsen \(2006\)](#) write “[i]t is not that social evolution is analogous to evolution in the natural world; it is that at a high level of abstraction, social and biological evolution share these general principles.” (p. 14).

These three principles are discussed in more detail when Hayek’s theory of cultural evolution is considered in Chapter 4; and in the discussion of simulations based on the second model in Chapter 10.

---

<sup>10</sup>Different versions of these three principles exist. For example, [Hodgson \(2002a\)](#) and [Hodgson and Knudsen \(2006\)](#) use variation, selection, and inheritance; and [Stoelhorst \(2007\)](#) uses “retention” in place of inheritance because it “takes us further away from the genetic overtones of the term inheritance.” (Endnote 1, p. 251). Durability is used here for the same reason and because it is consistent with [Hodgson’s \(1991\)](#) use in the context of Hayek’s theory of cultural evolution (discussed further in Chapter 4).

The second dimension in which the three principles of Darwinism can be generalized relates to the stratified ontologies discussed above. Consistent with this, [Hodgson \(2002b\)](#) writes that it “is possible that some of the reaction against ‘biological analogies’ is grounded on a mistaken view that theories operate on the one level only.” (p. 273). Instead, “Darwinism biology invokes multi-levelled explanations, in which the theory of natural selection is the over-arching and organising theoretical framework.” (ibid). This includes the notion of “group selection”, which is relevant when we discuss Hayek (his theory of cultural evolution hinges on this principle).

This thesis is open to the idea of generalized Darwinism, which includes both the abstract interpretation of Darwin’s three principles and multi-level explanations.

However, as was the case with emergence above, this use of a Darwinian perspective does not mean that every aspect of an economic system should be explained by these forces. Rather, our analysis is *open* to framing agent behaviour and social change in this way. Also, as [Hodgson and Knudsen \(2006\)](#) argue, generalized Darwinian theories should be viewed as necessary but not sufficient for explaining certain phenomena: we also require a coherent explanation of them.

#### *Comparison with Neoclassical Theory*

In addressing this issue, it is important to distinguish between free market economics and Neoclassical theory.

The broader approach to Darwinian evolution as outlined above is compatible with the idea that firms operate in competitive free markets. Firms could be viewed as equivalent to organisms in an ecosphere; they might intentionally adapt to changes in consumer demand; these changes might be sustained for some time; and a firm will ‘win’ relative to other firms if its adaptation is commercially successful. Firms that do not adapt will ‘lose’ and might therefore fail.

However, Neoclassical theory is not the same as free market economics. Broadly speaking, NCE’s emphasis on stable, equilibrium conditions means it sits awkwardly with generalised Darwinism. As mentioned above, Arthur referred to the Neoclassical approach as “dead, static, timeless, and perfect.” ([Arthur, 2013](#), p. 19)<sup>11</sup>.

Furthermore, we might also mention how the Austrian School of Economics, instituted by Menger, supports a broadly free-market view of economics while not accepting many of the premises of Neoclassical theory. Hayek<sup>12</sup> is discussed in some detail in Chapter 4, including a focus on what he meant by ‘the market order’.

---

<sup>11</sup>Related to this, Veblen wrote a paper that asked (and was entitled) “Why is Economics Not an Evolutionary Science?” ([Veblen, 1898](#)).

<sup>12</sup>Hayek is generally viewed as an economist from the Austrian school (as well as an economist who was born in Austria!). However, as we discuss briefly in chapters 3 and 4, he differed from Menger and von Mises (two of the most prominent Austrian economists) in certain respects.

Moreover, it is curious that a number of economists associated with the Neoclassical school had affiliations with socialism. For example, Kenneth Arrow states in [Klein \(2013\)](#) that he had been a socialist in his youth and sustained these views until he entered graduate school, after which “his work was shaped by a deeply leftist sensibility” (p. 268). Furthermore, [Klein and Daza \(2013\)](#) refer to the “market socialists of Abba Lerner and Oskar Lange”, two economists who took a broadly neoclassical stance in their work. [Lange and Taylor \(1938\)](#) went as far as to equate Marxist Economics with Neoclassical theory.

### *Co-Evolution and Co-Adaptation*

In the context of ‘evolution’, it is helpful to be clear about the meaning of co-evolution in this thesis as well as the similar notion of co-adaptation.

[Janzen \(1980\)](#) writes that co-evolution “may be usefully defined as an evolutionary change in a trait of the individuals in one population in response to a trait of the individuals of a second population, followed by an evolutionary response by the second population to the change in the first.” (p. 611).

Janzen’s definition pertains to multiple groups of organisms evolving in response to each other.

Equivalently, let us refer to co-adaptation as the adaptation of two or more agents to each other. More specifically, let us define co-adaptation as “adaptation involving reciprocal changes in two or more agents that affect their interactions.”

The word “reciprocal” is used in this definition to emphasise the iterative nature of co-adaptation.

In the complexity sciences it is common to refer to agent-to-agent co-adaptation as ‘co-evolution’. However, it seems preferable, in a general sense, to distinguish between these two processes in order to minimise confusion and to sustain two different mechanisms in our lexicons.

### **2.3.3.6 Economies are non-ergodic**

As mentioned in [Davidson \(1996\)](#), “[e]rgodic theory was explicitly expounded as part of the development of the theory of stochastic processes ... In a broader sense, however, ergodicity means the presumption of a preprogrammed stable, conservative system where the past, present, and future reality are predetermined whether the system is stochastic or not.” (p. 480).

[North \(2005\)](#) takes an equivalent view, stating explicitly that the economy is non-ergodic in nature:

An ergodic economy is one in which the fundamental underlying structure of the economy is constant and therefore timeless. But the world we live in is non-ergodic - a world of continuous novel change; and comprehending the world that is evolving entails new theory, or at least modifications of that which we possess. (North, 2005, p. 16)

North's reference to the economy being non-ergodic can be linked to Arthur's (2013) emphasis on how the economy is a combination of micro-economic behaviour and pattern formation and re-formation.

It would be a mistake, however, to believe that complex systems are always re-patterning themselves in an unstable way. Such systems can exhibit epochs of stability, with any re-patterning occurring infrequently.

#### *Comparison with Neoclassical Theory*

North's view was that Neoclassical theory is strictly ergodic in nature, stating "the ergodic hypothesis is implicit in much of current economic theory." (North, 2005, p. 19). For North, this meant that "the ergodic hypothesis is a-historical" (ibid), i.e., it runs counter to empirical evidence.

#### **2.3.3.7 Uncertainty is an important feature in economic behaviour**

In economics, the traditional way of describing uncertainty is to refer to Frank Knight's distinction between risk and uncertainty in Knight (1921):

For Knight, risk was a condition in which it was possible to derive a probability distribution of outcomes so that one could insure against such a condition. Uncertainty according to Knight was a condition in which no such probability distribution existed. Theorizing under the condition of uncertainty therefore was not possible, according to eminent theorists such as Kenneth Arrow (1951) and Robert Solow (1981). (North, 2005, p. 13)

This distinction between risk and uncertainty is useful but the idea that the latter is concerned only with a lack of probability distribution is not sufficient if we wish to examine *how* uncertainty arises in complex social systems. For example, is a distribution knowable in principle but we just do not know what it is now? Or is the distribution not knowable in principle?

We can also consider different 'sources' of uncertainty. In his research, North referred to three information-based sources: data, knowledge, and limited human cognition. We might also add a fourth, which is when agents' decisions are mutually contingent and an

infinite regress problem exists as a result of the agents trying to anticipate each other's actions.

In this thesis we take the view that CE is open to all forms of uncertainty that exist in the economy.

#### *Comparison with Neoclassical Theory*

In the quote above, [North \(2005\)](#) refers to Arrow and Solow who believed that theorizing “under the condition of uncertainty ... was not possible.” (p. 13). Furthermore, regarding uncertainty, North wrote:

Economists have themselves displayed a good deal of ambiguity on the subject, largely proceeding as though uncertainty was an unusual condition and therefore the usual condition, certainty, could warrant the elegant mathematical modelling that characterizes formal economics. But uncertainty is not an unusual condition; it has been the underlying condition responsible for the evolving structure of human organization throughout history and pre-history. ([North, 2005](#), p. 14)

While North refers to ‘Economists’ here, his critique appeared to be directed at Neoclassical economists.

Note that North's reference to “elegant mathematical modelling” points to a conflation about theorizing and the use of mathematics in Neoclassical economics. In general, mathematical techniques can incorporate risk (including stochasticity and probability) but they are not compatible with uncertainty, which might explain why Arrow and Solow thought uncertainty prohibited any theorizing.

However, theorizing does not necessarily have to involve mathematics: it can be conceptual in nature and it can also involve computational models. As Arthur noted, “[t]he objective, we should remember, is not necessarily to formulate equations or to arrive at necessary conditions. The objective, as it is with all theory, is to obtain general insights.” ([Arthur, 2013](#), p. 7).

#### **2.3.3.8 Agents use mental models to make decisions**

The complexity sciences generally focus on localised behaviour by agents in complex systems, which means there is typically a focus on how agents receive and process information and make decisions. As a result, complexity scientists are interested in agents' mental models (this was discussed in Section 2.1.1 above). CE adopts the same approach and in this thesis we focus on people as agents, i.e., we are interested in human mental models.

Furthermore, the economic ontology we are interested in includes uncertainty, hence in this thesis we take a strong interest in human mental models working under conditions of uncertainty. Moreover, North viewed the mitigation of uncertainty as one consequence of institutions, which further motivates our interest in mental models under conditions of uncertainty.

In Neoclassical theory, the standard assumption is that agents use substantive rationality to make decisions. The problem we face is that such reasoning is not suitable under conditions of uncertainty:

To the degree that outcomes are unknowable, the decision problems they pose are not well defined. It follows that rationality — pure deductive rationality — is not well-defined either, for the simple reason that there cannot be a logical solution to a problem that is not logically defined. It follows that in such situations deductive rationality is not just a bad assumption; it cannot exist. (Arthur, 2013, p. 4)

This raises the question of how human agents reason under conditions of uncertainty. For Arthur, the answer is that people use pattern recognition and simple models to help them make decisions when uncertainty is prevalent:

Indeed, as Shackle (1991) puts it, “The future is imagined by each man for himself and this process of the imagination is a vital part of the process of decision.” One way to model this is to suppose economic agents form individual beliefs (possibly several) or hypotheses — internal models — about the situation they are in and continually update these, which means they constantly adapt or discard and replace the actions or strategies based on these as they explore. They proceed in other words by induction (Holland et al (1986); Sargent (1993); Arthur (1994)). (Arthur, 2013, p. 4)

It is important to note here that the word ‘induction’ has a variety of meanings across different literatures. Arthur (1994) very clearly aligns this word with the broad notion of pattern recognition, which is different to (for example) induction as ‘making generalizations from the particular’ and from making inferences from data alone. Moreover, Arthur’s meaning appears consistent with Holland et al’s (1986) use of the term.

Given the different meanings of ‘induction’ which exist, this thesis will generally avoid using the term and will refer instead to pattern recognition when discussing Arthur and Holland et al’s research.

*Comparison with Neoclassical Theory*

The approach taken in Neoclassical theory was mentioned briefly above: the default assumption is that agents use substantive rationality to make decisions. This can be viewed as a form of mental model.

Here it is useful to distinguish between two different forms of reasoning: *substantive rationality* and *procedural rationality*, following [Simon \(1976\)](#).

Substantive rationality is used “when it is appropriate to the achievement of given goals within the limits imposed by given conditions and constraints” (ibid, p. 66). This approach is suitable when agents have fixed preferences and a goal like utility maximization: substantive rationality deduces the agents’ choices that achieve their goals. In this type of challenge there is no uncertainty.

By contrast, “[b]ehaviour is procedurally rational when it is the outcome of appropriate deliberation. Its procedural rationality depends on the process that generated it.” (ibid, p. 67). Herb Simon noted that when psychologists use the term ‘rational’, they are typically referring to procedural rationality. Note also that this type of rationality can work under conditions of uncertainty.

The mental models used in the computational models developed for this thesis fit what Simon meant by procedural rationality. The agents use simple models to make decisions under conditions of uncertainty.

Simon used the assumptions of Neoclassical economics to show how the field achieved analytical reductionism:

... the assumptions of utility or profit maximization, on the one hand, and the assumption of substantive rationality, on the other, freed economics from any dependence upon psychology. ([Simon, 1976](#), p. 66)

[Arnsperger and Varoufakis \(2006\)](#) referred to substantive rationality as *methodological instrumentalism*. This was the second of the three axioms of Neoclassical theory they identified.

### 2.3.3.9 Complexity Economics is open to inter-disciplinarity

One rationale for an inter-disciplinarity strategy is its link to stratified ontologies, which were discussed in Section 2.3.3.4 above. With this in mind, we find that, unlike Neoclassical theory, complexity economists take an interest in psychology, e.g., when understanding agents’ mental models. The interest would not stop there: neuroscience and the cognitive sciences would also be of interest in making sense of how people process information, including under conditions of uncertainty.



Equivalently, if economics is (or should be) about identifying social phenomena like institutions, it would be natural to look at how sociologists think about such things. In the context of a stratified ontology, we might consider the rationale for looking at psychology, cognitive science, neuroscience, and sociology as ‘vertical’ in nature.

However, we can think of inter-disciplinarity in the ‘horizontal’ sense also. As mentioned previously, the complexity sciences grew out of the natural sciences, and this has included a range of abstract concepts that relate to complex systems in different disciplines. As emphasised above, we should be careful mapping between domains but abstract concepts like self-organisation can be useful even if they merely provide inspiration to ask different questions.

Furthermore, other disciplines can be used to help us understand how other researchers have looked at the same - or similar - phenomena. For example, the agency / structure debates in philosophy and sociology appear related to how economists have looked at institutions in economics.

#### *Comparison with Neoclassical Theory*

This has been covered in earlier sections: Neoclassical theory is essentially reductionist in nature. Recall the quote from Simon above that showed how utility maximization and substantive rationality have been used to “free” economics from any dependency on psychology. As a result, NCE has historically taken relatively little interest in other fields of study.

#### **2.3.3.10 Complexity Economics values realism and is sympathetic to instrumentalism**

We argue here that in complex, non-ergodic systems, instrumentalism has a place but theories that emphasise realism are (i) necessary if we want to bring about change in the real world; and (ii) more likely to prove accurate when faced with novel change than theories that do not.

The debate concerning instrumentalism and realism is extensive within the philosophy of science literature and a detailed evaluation of the many arguments is beyond the scope of this thesis. This debate extends as far back as Dewey who developed instrumentalism in the late 19th Century / early 20th Century as part of his pragmatist philosophy<sup>13</sup>. The [Stanford Encyclopedia of Philosophy \(2023\)](#) contains a more detailed discussion of these origins, under the entry “John Dewey”.

It is, nonetheless, helpful to understand some parts of this debate and to consider them in the context of the complexity sciences.

---

<sup>13</sup>Pragmatism is discussed further below.

Let us start the discussion by defining what we mean by instrumentalism and realism.

The **instrumentalist thesis** states that theories should be properly considered as just tools or instruments for making observable predictions. Thus, the question of whether a predictive theory is true or false or whether its theoretical terms refer is of little moment for the pure instrumentalist. The usefulness of a theory is determined, therefore, by its predictive scope or range of applicability. (Keita, 1983, p. 79-80, emphasis added, footnote 2 removed)

By contrast:

The **realist thesis** ... states that science aims at giving a true picture of the world, and that the theoretical terms of successful and well-confirmed theories which purport to refer to existent entities actually refer to existent entities. Furthermore, acceptance of a scientific theory implies that what the theory states is true. (Keita, 1983, p. 80, emphasis added, footnote 3 removed)

### *Instrumentalism*

Dewey's instrumentalism originated in part from his emphasis on each individual's limited range of experiences in the world, e.g., in Dewey (1925), which is entitled *Experience and Nature*.

We can add to this two further limitations, due to cognition and observation. The former emphasises the distance between restricted human cognition and reality<sup>14</sup>. The latter recognises that many phenomena are beyond our ability to observe, e.g., atoms prior to the invention of powerful microscopes; and neurological processes prior to Magnetic Resonance Imaging.

We can contrast these limitations with a reality that is, from a complexity science point of view, fine-grained<sup>15</sup> and non-ergodic.

This tension between myopic, subjective agents and a complex reality means that simplifications are inevitable in human cognition, e.g., narratives, metaphors, reifications, heuristics, and abstract concepts are all simplified forms of sense making. Furthermore, in his earlier work, Dewey argued that a broadly Darwinian process (which appears to

<sup>14</sup>Heiner (1983), for example, discusses this gap in the context of an agent's *competence* versus the *difficulty* of selecting between preferred alternatives (his 'CD gap')

<sup>15</sup>Related to this, Gell-Mann and Hartle (2007) refer to coarse-grained cognition in contrast to a fine-grained reality.

fit with the generalised Darwinism described above) exists within the mind, resulting in the selection of those simplifications that prove the most accurate<sup>16</sup>.

Given the discussions in previous sections of the subjectivity of mental models, non-ergodicity, generalised Darwinism (and the discussion of pragmatism below), we can state that, overall, CE should be sympathetic to instrumentalism as described by Dewey. Furthermore, it should be clear that the combination of simplification and the value of prediction is as relevant in scientific research as it is in our daily lives. All theories are heuristics to some degree.

Before we discuss realism in more detail, it is helpful to frame the difference between instrumentalism and realism via a black box analogy.

Instrumentalism means that the aim of theories is to generate accurate predictions: the patterns within (or components of) the theory, which is like a black box, do not have to correlate with the patterns of the target domain; only the theory's predictions do. We can contrast this instrumentalism with realism by adjusting the black box analogy: the patterns inside the box should describe (thereby correlating with) reality<sup>17</sup>.

### *Realism*

Here we look at two arguments in support of realism in economics (and the social sciences more generally).

The first argument is that if a theory is used to bring about change in real world human social contexts, then it is strongly preferable for the corresponding theory to be realistic. The problem with instrumentalism in this regard is mixing causation with correlated patterns.

Consider game theoretic explanations for organic institutional emergence, including one group that frames such institutions as equilibria in games, e.g., Calvert (1995), which is discussed in Chapter 5. The equilibrium 'output' of such games appears to correlate, approximately, with organic institutions we see in the real world; however, the mechanisms assumed within Calvert's (1995) equilibria include substantive rationality as a mental model in addition to a range of game-related constraints<sup>18</sup>.

---

<sup>16</sup>This appears to overlap with both Hayek's theory of mind and Holland's classifier systems, which are discussed in Chapter 4.

<sup>17</sup>This raises the question of whether it is possible to directly describe reality. Addressing this question in detail is beyond the scope of this thesis; however, let us agree that we can recognise *degrees* of realism by common sense, e.g., describing a 'lion' as a cat is more realistic than referring to it as a 'planet'.

<sup>18</sup>In the case of Calvert (1995) the problem is more serious because his model only demonstrates the conditions under which equilibria exist. He is clear that his model is not concerned with the mechanisms of emergence but this does not stop him from declaring that "*Institution* is just a name we give to certain kinds of equilibria." (p. 74, emphasis included).

Now consider cryptocurrencies, which were mentioned in the Introduction, as a live example of an institutional ‘problem space’. A theory of institutions as equilibria is surely of no use to people managing and developing cryptocurrencies who want pragmatic institutional solutions to the various ‘quality’ problems identified in the Introduction. What is required instead is a theory that is grounded in realistic mechanisms that explain how institutions emerge, to help guide change in that domain.

We can reflect further on the problem here in the context of the black box metaphor above: attempting to bring about change in a social context is like tinkering inside the black box.

If the box contains reasonably realistic mechanisms, we stand some chance of predicting the likely (and desired) outcome but if the mechanisms are non-realistic then we run a risk related to trying to use correlated relationships in a causal manner. This resembles Goodhart’s law, which states that “when an empirical regularity starts to be exploited as a basis for economic policy, it is liable to break down” ([Black, Hashimzade, and Myles, 2009](#)).

Attempting to bring about change in the real world requires real causal mechanisms.

The sensible response to the above argument is that non-realistic theories should not be used in this way. Unfortunately, however, the use of such theories and models (notably those employing substantive rationality) is widespread in Western public policy. See [Colander et al \(2009\)](#) for a discussion of this point. We might note, also, that Goodhart’s law originated precisely because of the attempt to exploit correlated relationships to determine some output (inflation).

The second argument in favour of realistic theories is that, from a complexity science point of view, the value of instrumentalism is limited by the problem of prediction in complex systems. This probably applies less to realistic theories (provided, of course, they are accurate to some degree).

Instrumentalism is compatible with ergodic systems (as defined by North above) because the patterns of such systems do not change, which makes prediction reliable<sup>19</sup>. We can imagine a competitive, scientific, process in which the best predictors are selected over time.

A fundamental problem arises with instrumentalism when a system changes in a novel way, which is to say it is re-patterned in an unpredictable manner. This highlights an inherent tension between pre-conceived non-realistic theories and novel change: the former are meant to predict outcomes and the latter is an unpredictable structural

---

<sup>19</sup>Note that non-ergodic systems are not necessarily static: patterns can also be inter-temporal such that these systems change, predictably, over time. Furthermore, non-realistic theories can handle such dynamics: systems of differential equations are a case in point.

change. Note the problem is not so much about the inability of a non-realistic theory to predict novel change; it is about the resulting re-patterning and (likely) decline in predictive power. The implication is that non-ergodicity leads to non-realistic theories having a shelf-life.

A good example of the problem here is provided by standard Value At Risk models in finance, which are used to manage risk in asset portfolios. These models typically use historical data to derive a variance-covariance matrix between the prices changes of the assets in a portfolio. The efficient markets hypothesis leads to an assumption that this matrix is approximately accurate and can be used to choose how much risk a portfolio is exposed to. The problem, however, is that in times of stress, these backward-looking matrices no longer reflect reality and portfolio values can fall much more than was previously deemed likely<sup>20</sup>. Put another way, the patterns assumed inside the VAR ‘black box’ are rendered obsolete by novel changes in the system.

This problem of novel change also applies to realistic theories. However, the proposition here is that such theories are likely to be more robust in the face of novel change. This means they are likely to perform better vis-à-vis predictions in non-ergodic systems than their non-realistic counterparts and, also, they are probably easier to update after some structural change.

Whether this is true, however, will depend on various factors, notably the realism of the pre-existing theory and the nature (and magnitude) of the novel change. A highly realistic theory which is rendered slightly less realistic by some minor change is likely to perform better (and is more easily updated) than a less realistic theory faced with radical change.

The problem mentioned above, that reality might not be sufficiently observable and knowable, is an important qualification here. An ability to observe reality hampers the building of such realistic theories.

To summarise the above discussions, we can say that CE is sympathetic to instrumentalism but, ultimately, there are important reasons for theories to be realistic, notably their use to bring about change in social systems.

Finally, we should note the relevance of the above discussions for the models developed for this thesis, which were designed with realism in mind. As mentioned in the Introduction, the ultimate aim of these models is to help develop our understanding of organic institutional emergence (and how legal rules relate) with a view to being helpful in real-world problem spaces.

---

<sup>20</sup>During the 2007-8 financial crisis, some risk managers talked about asset price volatility that should not have been observed over the lifetime of the known universe.

*Comparison with Neoclassical Theory*

It is not possible to firmly align Neoclassical theory with either instrumentalism or realism: the relevance of either approach differs between researchers and individual theories. Indeed, it is hard to imagine even the most hardened advocate of Neoclassical theory rejecting realism on its own merit.

However, it is probably reasonable to state that Milton Friedman's famous perspective (which is firmly associated with instrumentalism) has been influential among Neoclassical (and other) economists. The classic text is [Friedman \(1953\)](#) where he uses the example of billiard players to make his case:

Consider the problem of predicting the shots made by an expert billiard player. It seems not at all unreasonable that excellent predictions would be yielded by the hypothesis that the billiard player made his shots as if he knew the complicated mathematical formulas that would give the optimum directions of travel, could estimate accurately by eye the angles, etc., describing the location of the balls, could make lightning calculations from the formulas, and could then make the balls travel in the direction indicated by the formulas. Our confidence in this hypothesis is not based on the belief that billiard players, even expert ones, can or do go through the process described; it derives rather from the belief that, unless in some way or other they were capable of reaching essentially the same result, they would not in fact be expert billiard players. ([Friedman, 1953](#), p. 21)

From a complexity science perspective, we can interpret Friedman's argument as proposing a reductionist strategy for making sense of systems that are simple, mechanistic, and ergodic. Such a strategy would be reasonable for these types of system because they can be broken down into constituent parts (players, table, balls, etc.) with known, predictable relationships. Furthermore, as-if assumptions can be used to approximate how players play the game provided one's aim in understanding is not subverted by these approximations, e.g., simply knowing the trajectory of the balls.

However, given the defining features of complex systems in Section 2.1 above, Friedman's description of billiards is a terrible metaphor for economic systems<sup>21</sup>, i.e., the premises on which his analysis is based are flawed. Imagine instead if billiards is played by many players, each adapting to each other's style of play, when the balls, table, rules, and aims were also evolving, all simultaneously.

---

<sup>21</sup>Archibald, Simon, and Samuelson do not hold their punches when they write that "[t]he expressed purpose of Friedman's principle of unreality (or as-if hypothesis) is to save Classical theory in the face of the patent invalidity of the assumption that people have the cognitive capacity to find a maximum" ([Archibald, Simon, and Samuelson, 1963](#), p. 230).

The distance between complex economic systems and Friedman's billiards analogy tells us that instrumentalism should be used in economics with care. This is particularly true in light of the arguments in support of realism set out above.

### 2.3.3.11 Pragmatism

The final principle of CE is an association with pragmatist philosophy as depicted by Charles S. Peirce (e.g., Peirce, 1905), William James (e.g., James, 1890), and Dewey (e.g., Dewey, 1938).

There is an enormous literature on pragmatism versus Cartesian philosophy and positivism; and these approaches span many academic domains. What follows are concise summaries of pragmatism and positivism and short discussions of their relationship to economics.

The underlying assumptions of pragmatism include the ideas that perception is subjective, that reality is socially constructed, and that observers are an integral part of what is being observed. Furthermore, each mind is an assemblage of habits and beliefs; and the idiosyncrasies of personal histories means that minds are heterogenous.

The consequences of these assumptions can be divided into: (i) the implications for our understanding of phenomena; and (ii) the process of research. Regarding the former, our understanding of reality has to be aware of context and also reflexive in that observers should attempt to be aware of their impact on a situation. Multi-perspectivalism is generally valued in light of the heterogeneity of minds.

In general, pragmatists accept that the process of research is influenced by human interests and that in practice it is a human act: within literatures there is a social construction of understanding. Furthermore, qualitative and quantitative modes of research are both valued, and Peirce originated the notion that ideas are developed through *abduction*, which is "the process of forming an explicit hypothesis." (Peirce, 1903, p. 216).

There is a comfortable affinity between CE and pragmatism because they have similar ontologies. Most notably, Dewey's pragmatism emphasises a grounded-ness to understanding reality which fits neatly with the bottom-up orientation of the complexity sciences. Related, Blumer (1969) writes that no "theorizing, however ingenious, and no observance of scientific protocol, however meticulous, are substitutes for developing a familiarity with what is actually going on in the sphere of life under study." (p. 39). Also relevant is Commons' (1934) reference to pragmatism in economics as "the scientific investigation of economic relations of citizens to citizens" (p. 157).

Furthermore, the social construction of institutions (including, most importantly, language) and the subjective nature of sense-making both point to a reality which is perceived differently by agents.

Importantly, an affinity with pragmatism should not be mistaken for a wholly incorrect interpretation of this paradigm as anti-scientific, notably as devaluing verification of hypotheses. In fact, by appreciating a range of qualitative information such as interviews, it can be argued that pragmatists have a greater appreciation of verification than positivists.

Pragmatism is frequently contrasted with positivism, which contains a Cartesian duality: the separation of an independent observer and an objective, external reality. In terms of helping us understand phenomena, in principle this reality is comprehensible in the same way to all observers, i.e., a homogenous understanding of the truth is available. There is generally an emphasis on data and quantification given the ambiguity of qualitative information.

In terms of the process of science, the aim of the positivist is to find the objective truth. Moreover, Comte's positivism emphasised simplification, reductionism, and the attainment of laws that demonstrate a causal relationship between phenomena.

#### *Comparison with Neoclassical Theory*

NCE is closely associated with positivism (e.g., [Katouzian, 1980](#); [Bromley, 2006](#); [Aligica, 2013](#)). [Bromley \(2006\)](#) argues that the problem with this approach is that it is narrow: the “‘explanation’ of a phenomenon is always bound to the limits set up by axioms and assumptions” ([Aligica, 2013](#), p. 183). Given the association of NCE with “rationality, self-interest, and utility maximization” (ibid), explanations of economic behaviour are constrained by these phenomena.

### **2.3.4 Neoclassical Economics as Restricted Complexity Economics**

Given the descriptions in Section [2.3.3](#) above, we can interpret NCE - approximately speaking - as a special case of CE. Put another way, if we constrain our approaches in the eleven categories in particular ways, we arrive at CE:

1. Agents in Neoclassical theory use a restricted form of mental processing - utility maximization - whereas CE is open to other forms of cognition, including Simon's bounded rationality.
2. Neoclassical economics is focused on equilibrium systems whereas CE is open to non- and dis-equilibrium phenomena, in addition to equilibrium.
3. Whereas CE considers the formation of, and allocation within, the economy, Neoclassical theory is focused on the latter.



4. CE is open to a stratified ontology in addition to reductionist explanations, whereas NCE is reductionist in nature.
5. In Neoclassical theory, the economy does not evolve: it is static, or “dead”. In CE, the economy evolves over time in a way that is open to generalized Darwinism, although CE recognises that objects and patterns might be stable for periods of time.
6. Included in our understanding of CE is the idea that economies are non-ergodic, i.e., the idea the economy re-patterns itself. By contrast, Neoclassical theory treats the economy as ergodic, i.e., only involving fixed patterns.
7. CE is open to uncertainty (and various types and degrees of certainty) whereas Neoclassical theory rejects uncertainty and embraces risk (as defined above in the quote from [North, 2005](#)).
8. As will be discussed further below, the mental models mostly used in Neoclassical theory (substantive rationality) are restricted to create stable macro patterns (equilibrium conditions) for which they are also an appropriate response. In CE, mental models can be broader than this, e.g., pattern-based reasoning under conditions of uncertainty.
9. Neoclassical economics is less open to inter-disciplinarity than CE because of its reductionist stance.
10. Neoclassical theory appears to emphasise instrumentalism more than realism, approximately speaking, whereas CE is comfortable with both.
11. Positivism should not be viewed as a subset of pragmatism because the two do not share underlying assumptions of reality. However, pragmatists can incorporate much of the positivists practical ways of doing research, e.g., hypothesis testing and the use of quantitative data.

Consistent with these comments, when contrasting it with Neoclassical theory, [Arthur \(2013\)](#) argues that “Complexity economics, we can say, is economics done in a more general way.” (p. 19).

### 2.3.5 Choice of Framework in Economics

A question we must address is: which framework should we use for economic research?

From a number of the discussions above, it is tempting to argue that CE contains a more realistic description of how the economy works than Neoclassical theory so the former should be preferred over the latter.

However, this would be naïve. A general appeal to realism is not a sufficient argument for choosing CE because a Neoclassical economist might argue that abstractions are always necessary in economics. The question is then: which abstractions?

We can, however, appeal to realism in light of the arguments developed in Section 2.3.3.10 above. For this thesis, the most important of these is the aim of using economic research to improve real-world situations. If we assume this as a goal then it supports the use of CE, which values realism more than Neoclassical theory has historically done.

In addition, there is a great deal NCE misses in terms of how economies work. [Arthur \(2013\)](#) summarises this in the following way:

Equilibrium of course will remain a useful first-order approximation, useful for situations in economics that are well-defined, rationalizable, and reasonably static, but it can no longer claim to be the center of economics. Moving steadily to the center is an economics that can handle interactions more generally, that can recognize nonequilibrium phenomena, that can deal with novelty, formation and change. ([Arthur, 2013](#), p. 19)

Nonetheless, even if we assume the goal of using research to improve real-world situations, there are two ways in which NCE can still be useful. First, for some research questions, Neoclassical theory might be sufficient. For example, if we are interested in stable, equilibrium conditions and if substantive rationality were a reasonable approximation of human cognition given the subject matter then Neoclassical theory would be reasonable.

Second, this approach might provide a useful first-order approximation even if the economy behaved in a more complex way. [Arthur \(2013\)](#) makes this point, noting that “under the undeniable force of gravity an approximately equilibrium sea level has first-order validity.” (p. 12). However, “in the ocean the interesting things happen ... on the surface where ... the boats are.” (ibid). Put another way, first-order approximations can be a useful starting point before moving to something that better respects the complexity of the phenomena being studied.

[Bowles and Gintis \(2011\)](#) make a related point concerning analytical solutions for “highly complex selection processes operating at two levels - individual and group” (p. 125). Samuel Bowles and Herbert Gintis argue that this type of dynamic is not amenable to mathematical analysis but it is accessible by agent-based models. This point is very relevant for this thesis given the combination of individual and group-level phenomena observed in the simulations reported below.

The point being emphasised here is that it is not necessary to reject the Neoclassical approach entirely: it is a mature body of work that still has some use. However, there

are reasonable questions about its use in real-world contexts, and many phenomena lie outside its field of vision. In private correspondence, Arthur wrote that CE would not be necessary:

if you don't care about change, structural change, development, innovation, technology, evolution, formation, how institutions emerge, systems being gamed, etc.

[CE] is about the economy in formation, and that's at least as large a set of problems as the economy in equilibrium. (Arthur, 2019, personal communication, dated 13 December 2019<sup>22</sup>)

### 2.3.5.1 Rationale for Using Complexity Economics to Study Organic Institutions

The discussions above provide useful context for rationalizing the use of CE to explore organic institutions in this thesis. There are three points worth noting here.

First, as Hodgson (2002a) argues, and Field (2007) implies, there is a gap in the literature concerned with the *origination* of organic institutions. CE is well suited to addressing this - both conceptually and formally - because it is concerned with the emergence of macro patterns from the bottom up, as well as the relationship between such patterns and agents' mental models. Such a framing requires a stratified ontology and there is also a clear link to *economic formation*. Both of these are included in CE.

Second, North emphasised that institutions help to mitigate uncertainty for agents. Uncertainty is also a part of the ontology of CE, which is not the case for Neoclassical theory.

Third, CE includes formal modelling technologies (including agent-based modelling) that allow for the exploration of emergent institutions, as well as the role of regulations and laws, within whole economic systems. Furthermore, this technology allows for experimentation that is not possible in the real world, e.g., analysing the impact of corruption on the efficacy of legal rules.

Overall, CE appears very well suited to exploring the emergence of organic institutions, which is the overarching aim of this thesis.

### 2.3.6 Internal Consistencies: Complexity Economics and Neoclassical Theory

One of the most important and valuable features of NCE is the internal consistency it demonstrates between its micro and macroeconomic approaches. Arthur (2013) writes:

---

<sup>22</sup>Printed with permission.

economics early in its history ... asked ... what behaviors (actions, strategies, expectations) would be upheld by - would be consistent with - the aggregate patterns they caused. It asked in other words what patterns would call for no changes in micro-behavior, and would therefore be in stasis, or equilibrium. (Arthur, 2013, p. 2)

The framework created by Neoclassical theory is one of agents maximizing their utility via deductive reasoning (substantive rationality) within a macro environment that contains no uncertainty. Importantly, the agents' combined behaviour generates a macroeconomy in equilibrium *for which each agent's choices are appropriate*. Arthur (2013) notes this is true of general equilibrium theory, classical game theory, and rational expectations economics (p. 2). We can say, therefore, that NCE demonstrates a healthy internal consistency between its micro and macro theories.

The overarching ontology of CE is different to Neoclassical theory in all of the ways listed in Section 2.3.3 above. Arthur (2013) writes:

Complexity ... asks how individual behaviors might react to the pattern they together create, and how that pattern would alter itself as a result. This is often a difficult question; we are asking how a process is created from the purposed actions of multiple agents. (Arthur, 2013, p. 2)

However, even though the ontologies differ, it appears CE also includes an internal consistency between micro behaviour and macro patterns, i.e., it provides a coherence equivalent to that seen in Neoclassical theory. We can characterise this in the following way:

- Agents operate under conditions of uncertainty and use mental models to make sense of their environment and to make decisions;
- Micro interactions between agents give rise to patterns (which might or might not be stable); and
- The nature, adaptation and co-adaptation of other agents' mental models, and the emergence of macro patterns, all contribute to the uncertainty that the mental models are an appropriate response to.

If we contrast the two approaches, however, the coherence of Neoclassical theory appears much more fragile than that of CE. If we change the rationality assumption of NCE (even very slightly) the equilibrium outcome would most likely fall apart and it is unlikely that this changed behaviour assumption would be an appropriate response to whatever macro phenomena were then seen. By contrast, CE is open to different mental models being

suitable to varying environments and it is open to non- and dis-equilibrium states. In fact, we would expect the agents themselves to adapt their mental models to whatever macro environment they face via a continuous process of pattern recognition.

Finally, it should be clear that the internal consistency of CE is compatible with stable and unstable economic systems (in the latter, patterns are ephemeral and outcomes volatile). Understanding the patterns Arthur referred to (including institutions), their stability, and their relationship with agents' cognitive processes are important challenges that CE appears better suited to than Neoclassical theory.

## 2.4 Criticisms of The Complexity Sciences & Complexity Economics

So far this chapter has promoted CE as more useful than NCE for the study of organic institutions. Here we turn our attention to criticisms of the complexity sciences and CE.

There appear to be five main criticisms (following [Horgan, 1995](#); [Horgan, 1997](#); and [Rosser, 1999](#)):

1. The complexity sciences have not delivered relative to the hype or the promise of / desire for a unified theory of everything.
2. Complexity science is the latest version of three previous fads: cybernetics, catastrophe theory, and chaos theory.
3. Complexity science is too loosely defined - there is no consensus definition of what it is.
4. Evaluation via empirical evidence is very difficult if we treat the economy as a complex system.
5. Complexity science is not necessary in economics: orthodox economics does a good enough job.

The first criticism is articulated in [Horgan \(1995\)](#), repeated in [Horgan \(1997\)](#), and referred to in [Rosser \(1999\)](#). Given Horgan's criticism first appeared in 1995 when the complexity sciences and CE were in their infancy<sup>23</sup>, it seems reasonable to ask how such expectations could have been met in a relatively short space of time (versus approximately 200 years of largely reductionist science since the Enlightenment).

---

<sup>23</sup>For example, the Santa Fe Institute was founded in 1984 and its economics department in 1988.

Furthermore, this is a relative criticism (delivery versus hype / promise) so it could equally be levelled at those responsible for raising expectations. Some complexity scientists themselves are probably guilty of this, which is perhaps not surprising given they are promoting a new subject in an overcrowded funding landscape. However, popular books like “Complexity” (Waldrop, 1993), while interesting and informative, have probably contributed to the hype<sup>24</sup>.

The second criticism was also set out in Horgan (1995) and referred to in Rosser (1999). The premises of this criticism are: (i) complexity science is essentially the same as cybernetics, catastrophe theory, and chaos theory; and (ii) these subjects were correctly found out to be fads by the academic community. The first of these premises is easily rejected by comparing descriptions of the four fields: there are a number of commonalities but they are not identical to each other.

The second premise can be challenged by arguing that cybernetics, catastrophe theory and chaos theory have made useful contributions to research, including ideas like the sensitivity of results to initial conditions. If these fields are to be viewed as fads then it is more likely because they delivered less than people expected, which would make the second criticism identical to the first.

Overall, the criticism that the complexity sciences are the latest manifestation of previous incarnations that proved to be fads appears very weak.

The third criticism is not unreasonable, e.g., Seth Lloyd identified over 45 definitions of ‘Complexity’ (listed in Horgan, 1997, p. 303, footnote 11). However, the complexity sciences represent a relatively new field / movement that is being used across multiple subjects, which means that, perhaps, this should be expected.

Furthermore, if we align the complexity sciences with multiperspectivalism then this diversity of definitions could also be viewed as a strength rather than a weakness (related to this, Cilliers, 1998 looks at the overlap between the complexity sciences and postmodernism).

The fourth criticism, regarding the problem of empirical evidence in the complexity sciences, relates to the discussions in Section 2.3.3.10 above. Rosser (1999) notes:

More generally, some argue that complexity implies a need to seriously rethink the nature of empirical testing in economics. Such a project threatens to lead into deeply philosophical issues such as induction versus deduction, objectivism versus subjectivism, and other difficult conundrums. (Rosser, 1999, p. 185)

---

<sup>24</sup>The back cover of Waldrop (1993) begins “In a rented convent in Santa Fe, a revolution has been brewing...”.

It is difficult to know precisely what Rosser means by this quote but the implication appears to be that delving into the philosophical issues he mentions would necessarily be a bad thing.

One counter-argument to Rosser's point is simply that orthodox economics appears unable to explain a range of empirically-observed phenomena, including the formation of the economy and organic institutional emergence.

More generally, the approach taken in this thesis is that the economy does resemble a complex system; that the complexity sciences would be useful in developing economic research (if done carefully); that empirical testing does require a "serious rethink" as Rosser put it; and that if this opens up "deeply philosophical issues" then so be it if this is a reasonable consequence of framing the economy as a complex system. A preference not to open up particular philosophical arguments is not a good rationale for not treating the economy as a complex system.

The fifth criticism, that the complexity sciences are not necessary in economic research, was discussed in Section 2.3.5 above so it will not be repeated here. As mentioned in that section, the complexity sciences deal much better with a variety of phenomena, including structural change, innovation, technology, formation, and organic institutional emergence, than orthodox economics.

Overall, it is telling that none of the five criticisms discussed above are conceptual in nature: in fact, it is noteworthy that critics of CE have failed to criticise through reasoned arguments about abstract concepts. Indeed, we can think of CE as providing a range of new concepts (like those listed in Section 2.2) *and* being more generalised than Neoclassical theory. In a sense, it appears *pareto superior* to NCE.

## 2.5 Chapter Summary

This chapter provided some of the foundations for the rest of the thesis by:

- articulating an overview of the complexity sciences (defining and analytical features), with an emphasis on the concepts most useful to this thesis;
- highlighting the potential problem of mapping these concepts, which mostly arose from the natural sciences, on to the social sciences;
- setting out eleven principles that are in part intended to help overcome this mapping problem and which also help define what is meant by CE in this thesis;
- discussing why a CE framing is more suited to understanding organic institutional emergence than Neoclassical theory; and

- evaluating the criticisms of the complexity sciences and CE, noting that all of these criticisms appear weak and that none of them include reasoned arguments about abstract concepts.



## Chapter 3

# Spontaneous Order - Rationale and History

To understand our civilisation, one must appreciate that the extended order resulted not from human design or intention but spontaneously.

– *The Fatal Conceit* by Friedrich Hayek (mostly)

The aim of this chapter and the next is to evaluate parts of the spontaneous order literature that are relevant to the thesis, and to highlight key issues and open questions within this literature that have a bearing on the research questions. This chapter focuses on the history of spontaneous order and the next chapter looks more closely at Hayek's framing, which is considered by many to be the most detailed and coherent available (Barry, 1982; Gray, 1998; Caldwell, 2004; and Luban, 2020).

While the term 'spontaneous order' was "coined by Michael Polanyi (1941, 1951)" (Luban, 2020, p. 68), the idea has a long and important history *avant la lettre*. Some of this history - that which is pertinent to the thesis - is the main subject of this chapter. Section 3.1 below introduces the term 'spontaneous order' and discusses the rationale for studying it. The second section (3.2) provides a 'potted history' of the term. Section 3.3 discusses specific historical influences on Hayek, whose version of spontaneous order is the subject of the next chapter.

### 3.1 Introduction and Rationale

The phrase often associated with 'spontaneous order' in the literature was written by Ferguson in his *Essay on the History of Civil Society* (Ferguson, 1767, p. 205). As stated

in the Introduction, he writes about order that is “the result of human action, but not the execution of any human design”<sup>1</sup>. Moreover, soon after Ferguson wrote these words, Smith mentioned essentially the same idea in *The Wealth of Nations*, whereby “man is led to promote an end which was no part of his intention” (Smith, 1776, p. 456).

The idea that order at some group level is a result of human action but unplanned is at the heart of spontaneous order. Let us now consider the rationale for this literature in some detail.

Assuming for the moment that forms of spontaneous order do exist in social systems then probably the simplest and most obvious reason for studying it is that it adds to our understanding of social systems. This, in turn, might help us - individually or collectively - to make better decisions than otherwise.

We can appreciate the importance of understanding spontaneous order if we see it in the context of “the ‘natural’ versus ‘artificial’ dichotomy handed down from Greek philosophy” (Boehm, 1994, p. 297); it creates “a distinct third category allowing for the explication of [this type of order]...” (ibid, pp. 297-8). Put another way, if explanations of phenomena are divided solely in to natural events like the weather or those involving human design, then this might render us blind to forms of order which result from neither.

In looking more closely at the distinction between spontaneous and designed order, Ullmann-Margalit (1978) argues that there are certain biases in human cognition which lead us to mistake the former for the latter. She wrote of “the ‘artificer bias’, that leads us to postulate a designer whenever we encounter what looks like evidence of orderliness and pattered structure...” (p. 268). Examples of planned phenomena include “the artist creating works of beauty”, “the inventor and manufacturer of elaborate machines”, and the “the manifestation of coordinated activity” (ibid, p. 268-9).

Ullmann-Margalit’s (1978) paper, entitled *Invisible Hand Explanations*, will be considered in more detail in the next chapter<sup>2</sup>. Most importantly, this paper, which follows Nozick (1974), helpfully distinguishes between two different “moulds” or “types” of spontaneous order, which map on to the two computational models described later in the thesis.

---

<sup>1</sup>We will see below that while Ferguson’s quote is synonymous with spontaneous order, he is often misunderstood. In addition, we can begin to appreciate the long history of spontaneous order by noting that Ferguson’s statement includes a footnote which references “De Retz Memoirs” (of 1669). In turn, Cardinal De Retz associates this notion with Oliver Cromwell, concerning “the fixity of all men’s designs and the uncertainty of their destiny.” (Forster, 1860, p. 24).

<sup>2</sup>Ullmann-Margalit (1978) has been generally well received in the spontaneous order literature, e.g., Barry (1982) refers to it as “an important article” (Endnote 2, p. 53) and Vanberg (1986) refers to it as “an excellent essay” (p. 81).

Returning to the rationale for studying spontaneous order, and bringing to mind its potential value for explaining some types of collective decision making, it should be of no surprise to us that this notion is prevalent in political ideologies.

In this context, it seems reasonable to state that Hayek's writings on spontaneous order have had a significant influence on classical liberalism, libertarianism, and conservative modes of thought. Indeed, Hayek saw his work as a defence of classical liberalism in the context of the Soviet Union and a bias toward 'rational constructivism' in Western thought, notably after the Second World War. More specifically, Hayek's ideas were cited by the architects of the New Right in the late 1970s and early 1980s. For example, Margaret Thatcher wrote in her autobiography that "the most powerful critique of socialist planning and the socialist state which I read at this time [the late 1940s], and to which I have returned so often since [is] F.A. Hayek's *The Road to Serfdom*." (Thatcher, 1995, p. 50)<sup>3</sup>.

There are two factors which might lead us to conclude that spontaneous order is less relevant today than it was, say, 30 years ago. First, the New Right experiment of approximately 40 years ago is viewed by many (though not all) as a failure, e.g., in discussing some of the problems of Reagan and Thatcher's New Right experiment, Gray (1998) writes that "Hayek ... failed entirely to comprehend how unfettered markets can weaken social cohesion in liberal cultures." (p. 147).

The second factor is that Hayek's thinking has been heavily criticised in the academic literature, including by Buchanan and Vanberg in *Economics*, and John Gray in *Social Philosophy*<sup>4</sup>. These criticisms will be developed in more detail in the next chapter but, for now, a common theme of these and other writers is that Hayek failed to articulate a convincing mechanism by which order emerges spontaneously in social systems.

However, these two points are not sufficient to reject spontaneous order being an important topic in the social sciences.

On the first point, the UK's Conservative Party voted for a leader in September 2022 (Elizabeth Truss) whose policy platform was based on the New Right and ideas developed at the Institute for Economic Affairs (IEA), a think-tank heavily influenced by Hayek's

---

<sup>3</sup>Given this impact on conservative thought, it is perhaps worth noting that Hayek wrote a postscript to *The Constitution of Liberty* (Hayek, [1960] 2006) entitled "Why I am Not a Conservative" (pp. 343-356). This apparent inconsistency can be understood as part of a larger theme concerning the wide variety of interpretations of Hayek's work (discussed further in the next chapter).

<sup>4</sup>In the first edition of Gray (1998), published in 1986, Gray was broadly complimentary of Hayek's body of work, but in the postscript to the third edition, published in 1998, Gray was much more critical, stating that Hayek's work "demonstrates that a powerful twentieth century project - the Marxian project of replacing market processes by central planning - is unachievable. *It tells us little else.*" (emphasis added, Gray, 1998, p. 150).

work. This tells us that approximately 40 years after Thatcherism, Hayek remains influential.

On the second point, while Hayek's framing is viewed by many (e.g., [Gray, 1998](#); [Luban, 2020](#)) as the most coherent explanation of spontaneous order, his failings are not necessarily failings of spontaneous order *per se*. Hayek does not own this idea, he is merely associated with one interpretation of it.

Moreover, if the main criticism of Hayek is that he failed to provide a persuasive theory of spontaneous order, we can present this as a challenge: can such a theory be developed which fits Ferguson's statement about order that results from human action but not human design? This question is too large for a thesis to answer comprehensively but it does provide further motivation for the work here.

There is a final point to note, which is perhaps best captured by [Luban \(2020\)](#) in discussing the seductive nature of spontaneous order to many:

...spontaneity (like the related notion of 'nature') has a way of remaining in the background even when explicitly disavowed, and many of those hostile to Hayek's particular political program still hold onto some vision of what a genuinely spontaneous order would look like. ([Luban, 2020](#), p 69)

Given the above comments, where does this leave us in terms of our rationale for studying spontaneous order? First of all, a slight of hand was used above in that an assumption was made that spontaneous order exists. This allowed us to consider various factors in the literature where this assumption is prevalent as well as political history where it appeared to influence decision making. It should be clear, however, that in the absence of a clear and compelling mechanism by which order can emerge in an unplanned way, we ought to be open to it not existing at all or, more conservatively, that it is of less significance than the likes of Hume and Hayek believed.

Put another way, the proponents of spontaneous order appear to take it on faith that at least one mechanism must exist which we simply haven't discovered yet. This leaves us with a more fundamental challenge of whether a mechanism, more concrete than vague references to spontaneity or emergence, can be identified.

With these motivations in mind, let us now look at a concise, summarised history of spontaneous order.

## 3.2 A Short History of Spontaneous Order

The concept of 'spontaneous order' has a long history even if the phrase dates from the Twentieth Century. [Barry \(1982\)](#) notes that Hayek saw his work on spontaneous order

“...continues a long tradition. While acknowledging it is absurd even to speculate on the beginnings of a tradition, Hayek often refers to the original Spanish schoolmen [of the 16th Century] as the founders of the theory of spontaneous order.” (p. 12). Looking even further back, [Boehm \(1994\)](#) remarks there are “intimations present in Aristotle and Aquinas” (p. 296); and, as mentioned in the Introduction, [Hamowy \(1987\)](#) writes of the fourth century BC Chinese Philosopher Zhuang Zhou who observed that “good order results spontaneously when things are let alone” (p. 6).

What value does a historical analysis of spontaneous order have for this thesis? There are three related points to make. First, history is path dependent so it is helpful in general terms to know how we got to where we are. This path dependence manifests in various ways, including the impact of Hayek’s work on current political philosophies, as noted above. Second, almost all of the researchers named below have contributed to, or are associated with, some feature of the spontaneous order literature and it is worth looking at - and where necessary evaluating - each of these. Third, the fact many of the great names in philosophy and economics (e.g., Hume and Smith) have considered this topic over several centuries is indicative of how important it has been in social science research.

A number of useful histories of spontaneous order exist. The most noteworthy include [Barry \(1982\)](#), which is the most cited; [Hamowy \(1987\)](#), which focuses on spontaneous order during the Scottish Enlightenment; [Sheehan and Wahrman \(2015\)](#), which looks specifically at the origins of self-organisation from the beginning of the Eighteenth Century; and [Gray \(1988\)](#), [Gray \(1998\)](#), and [Caldwell \(2004\)](#) discuss the historical influences on Hayek. Finally, [Boehm \(1994\)](#) deserves an honourable mention - this paper includes some history within a neat 6-page summary of spontaneous order.

The summarised history and analysis in this section loosely follows [Barry’s \(1982\)](#) version but we add to this with others’ analyses, when helpful, along the way.

[Luban \(2020\)](#) provides a helpful warning when looking at the history of spontaneous order from a perspective that follows the beginnings of the Soviet Union. He notes that the world of the Scottish Enlightenment writers - who play an important role in this history - was very different:

Spontaneous order theory grew out of twentieth-century anticommunism, and the fight against state encroachment upon economic life more broadly. The dichotomy of state and market underlies the entire theory in ways that are far-reaching yet rarely made explicit - the market is bearer of spontaneity, the state as bearer of constructivist rationalism; the market as realm of peaceful competition, the state as realm of coercive force; the market as grown, the state as making and made. ([Luban, 2020](#), p 78)

We should therefore consider the history of spontaneous order prior to Polanyi and Hayek as very much *avant la lettre*, given the association between this concept and arguments against centralized planning.

### 3.2.1 Pre-Scottish Enlightenment

Barry's (1982) history starts with the School of Salamanca of the Sixteenth Century, notably Luis de Molina who Schumpeter and Hayek viewed as a precursor to spontaneous order. According to Schumpeter, Molina's natural law doctrine was concerned with events which would occur "if they were allowed to work themselves out without further disturbance" (Schumpeter, 1954, p. 112).

Here we see an important distinction between deliberately designed laws and 'other' types of law - in this context,, between positive and natural law. The latter "are not in force because they have been 'artificially' made by a specified human activity but because they stem from God, nature, or reason" (Kelsen, [1949] 2006, p. 392). Positive law "comes in, as compared with the law of nature: it is made by human will" (ibid). Spontaneous order, in Molina's framing, is aligned with natural law. Moreover, this distinction between natural and positive law appears correlated with (but not the same as) the distinction made between common and statute law.

Sir Matthew Hale was discussed by Barry as claiming that common law "possessed a greater inner wisdom and rationality than the anti-traditionalist and a priori theories of law precisely because it accommodated facts and circumstances unavailable to the unaided reason." (Barry, 1982, p. 15). In turn, Hale "inaugurated a tradition of jurisprudence which we normally associate with Adam Smith and Edmund Burke and, in the present day, Hayek" (ibid).

In this context, and being mindful of the bottom-up orientation of the complexity sciences, it is tempting to align common law with an organic, bottom up "grown law" perspective; and, further, to align statute law with top-down-ism. This would, however, be a significant over-simplification. For example, in the UK, white and green papers are used in the early stages of statute formation for reasons of consultation with interested parties, before bills are written and then debated in parliament. Equivalently, as will be discussed in the next chapter, English Common Law has been materially enabled by statute law in the past.

It is worth briefly mentioning Thomas Hobbes in this potted history, not as a proponent of spontaneous order but as providing an antithesis<sup>5</sup>. Hobbes' famous view is that human nature involves "a general inclination of all mankind a perpetual and restless desire of power after power, that ceaseth only in death" (Hobbes, [1651] 2017, chapter

<sup>5</sup>Hobbes' framing should not be considered as the only antithesis. Hayek, for example, was concerned with spontaneous order versus organizations (Boehm, 1994, p. 298).

13), leading to a life (in the absence of relevant enforced laws) which is “solitary, poor, nasty, brutish, and short.” (ibid). Order had to be achieved by a strong, sovereign figure (hence, a ‘Leviathan’). This is very different to the unplanned, emergent order articulated by Hayek and others.

Barry (1982) also discusses Bernard Mandeville who articulated “an outrageous demonstration of the social benefits that accrue from vicious and self-interested motivations” (p. 17). Mandeville was most famous for his *Parable of the Bees*, first published as a poem but later on with commentary. Although he shared some of Hobbes’ views of human nature, Mandeville appeared more aligned with spontaneous forms of order than a Leviathan state.

It is worth noting here that Sheehan and Wahrman (2015) write about the origins of the term ‘self-organisation’ following the Mississippi and South Sea financial bubble of 1719-1720, after which the “vision of Newtonian order - that consequences follow causes in linear and stable fashion - was disrupted.” (Dale, 2018, p. 928). Importantly in the context of this thesis, Sheehan and Wahrman (2015) write that:

contemporaries came up with a novel notion that, when left to their own devices, complex systems generate order immanently, without external direction, through self-organization (Sheehan and Wahrman, 2015, p. x). (Dale, 2018, p. 928)

This quote points to an important ambiguity in the meaning of spontaneous order, that *good order necessarily emerges*. This ambiguity will be addressed by the Agent-Based Models (ABMs) developed later in the thesis: we will see that beneficial organic institutions *can* emerge in complex economic systems but it is *not inevitable* (this point has been made by other researchers, e.g., Sugden, 1989; and Hodgson, 2006a). Furthermore, the simulation results presented below show: (i) that there must be an *enabling environment* for such institutions to emerge; and (ii) detrimental forms of organic institutions can emerge under certain conditions (we see this when the parameter space of the second model is explored and also when we run certain experiments with this model).

Next in Barry’s chronology is Josiah Tucker who saw a balance between the two extremes of unplanned order and a Leviathan state. While Tucker was a mercantilist, he saw a role for the state in providing the conditions “required for the operation of an otherwise self-regulating commercial machine.” (Barry, 1982, p. 20). This idea of the state providing an enabling institutional environment for an otherwise free market system precedes the ordoliberalism and the Constitutional Economics developed by Buchanan, Vanberg, and others. These researchers will be discussed in more detail in the next chapter but the point to note here is that Tucker alluded to a designed institutional framework co-existing with an economy based on free markets.

### 3.2.2 The Scottish Enlightenment

According to Barry (1982), the Scottish enlightenment writers (he listed Smith, Hume, Ferguson, Dugald Stewart, and Thomas Reid) were “largely successful in integrating all these significant hints at a doctrine of spontaneous order into a general social philosophy” (p. 21).

Ferguson and Smith are of particular interest because they dovetail with Hayek’s framing and because of the repeated use in the literature of Ferguson’s definition of spontaneous order (stated above). Also, Smith is viewed as the grandfather of *Laissez Faire* Western economics<sup>6</sup>. When we look more closely at Ferguson’s and Smith’s work, however, both were in different ways sceptical about whether spontaneous order necessarily gives rise to beneficial outcomes.

Luban (2020) argues that the Scottish Enlightenment writers in general had a more balanced view toward ‘good’ and ‘bad’ unintended consequences of human action (he refers to these as *fecundity* and *perversity*, respectively): “[t]hey are often held up as pioneers of spontaneous order theory [but] they might better be understood as critics *avant la lettre*.” (p. 69). Moreover, Daniel Luban argued this applied most significantly to Ferguson and Smith. His reading was that Ferguson tended to a view of human nature “closer to Machiavelli’s” (Luban, 2020, p. 78).

This is important because there is an argument that the more we expect human interaction to give rise to “perversity”, the greater is the case for planned order (Hobbes’ Leviathan being an extreme). This is a crude argument, however, because it assumes that ‘planning’ can overcome social problems: this is at the heart of the debate between Hayek and “rational” central planner.

Whyte (2019) adds another factor concerning Ferguson. She argues that his opinions should be viewed in the context of his Christian beliefs, stating that “Hayek struggles to obscure the providentialism underpinning the account of social order he derives from Adam Ferguson and the Scottish Enlightenment.” (p. 156). Put another way, his faith in a Christian god was part of Ferguson’s telos. This is significant because we can think of God within Ferguson’s framework as providing a role resembling that of a planning entity. Paradoxically, spontaneous order might be the result of some divine plan.

The implication is that belief in God might lead to an expectation that, ultimately, fecundity will win out against perversity. This is not the same as a rational constructivist’s faith in central planning but it is nonetheless a form of faith in the (net) positive outcomes of spontaneous processes. Hayek’s atheistic interpretation of spontaneous order contained no such faith, other than, perhaps, his interpretation of Darwinian forces working on institutions.

---

<sup>6</sup>Barry (1982) refers to him as “Newton of the social sciences” (p. 21).



Looking more closely at Smith, Barry's (1982) comments make Smith look much more like a classical liberal than, say, a *Laissez Faire* anarchist, which is to say Smith believed that a "system of natural liberty ... can only work in the context of a form of interventionism; that of the enforcement of the strict rules of justice." (p. 27). Like Hale, Smith believed that common law was preferable to statute law as balancing positive law and natural justice; however, statute law should supplement common law where necessary, including in the need to ensure judges are accountable. Moreover, Smith was clear that he thought free markets give rise to some detrimental outcomes, which justified "a sizeable list of ad hoc interventions" (Barry, 1982, p. 28), e.g., a state system of education.

These comments about Smith point to an important distinction in the spontaneous order literature, between a so-called free market economic system and the institutional framework (including jurisprudence). As will be discussed further in the next chapter, Hayek used the former as a metaphor for the latter, which raises the important question of whether this is appropriate<sup>7</sup>.

### 3.2.3 Post-Scottish Enlightenment

Barry (1982) writes that "[i]t is commonly thought that after Smith the theory of spontaneous order went into a decline until the rise of Austrian economics" (p. 28) in the guise of Carl Menger. However, he notes the French *laissez-faire* school (albeit they "contributed little in the way of original theory to economics", *ibid*, p. 29) and Herbert Spencer. The latter took a broadly Darwinian view but, like Hayek later on, faced the problem of institutions emerging which "embody anti-liberal values" (*ibid*, p. 30).

The theoretical foundations of spontaneous order appeared to mature further with the work of Carl Menger, notably in Menger [1890] (1981) and Menger [1883] (1985). Menger made a helpful distinction between *organic* (unplanned) and *pragmatic institutions* (planned): the former include money, language, markets, and law and the latter "are the product of human deliberation and will." (Barry, 1982, p. 32). Indeed, the institutions that emerge in the simulations reported later in the thesis can certainly be categorised as organic in nature<sup>8</sup>.

Menger is generally understood to have been the founder of the Austrian School of Economics (Gray, 1998, p. 16), which was set up originally in opposition to the German Historical School. The Austrians are viewed as an important school of economics in their own right but they also had an impact on Hayek (who was Austrian) as a young adult when he studied in Vienna.

---

<sup>7</sup>There is also the related point that markets should be viewed as institutions themselves which typically exist within a wider enabling institutional framework.

<sup>8</sup>In some experiments we will look at how legal rules can catalyse this emergence. This gives rise to a categorization problem: should we refer to these as organic or pragmatic institutions?

In Menger's work, by far the most significant example of an organic institution is money. This was the subject of Chapter 8 of Menger [1890] (1981), and he wrote a broader account of organic institutions in Part 4 of Menger [1883] (1985).

Menger's [1890] (1981) account of money is well worth considering in some detail because it is a quintessential example of what proponents mean by spontaneous order<sup>9</sup>, but also because there is an important criticism of this account which calls into question its spontaneous nature, in the form of the *quality* of money.

### Menger's Account of Monetary Emergence

For Menger, the origins of money start from a barter economy. Depicting a "smith of the Homeric age" (p. 259) wishing to exchange suits of armour for consumables like food and water as well as raw materials, the smith will be aware of the difficulty of bartering his wares (the 'double coincidence of wants' problem). Knowing, however, that there are certain products (like cattle, in the case of Ancient Greece) that are more in demand, the smith will see the value in exchanging his armour for cattle in order to subsequently 'buy' (exchange cattle for) consumables and raw materials. Over time, as

... each economizing individual becomes increasingly more aware of his economic interest, he is led by this *interest, without any agreement, without legislative compulsion, and even without regard to the public interest*, to give his commodities in exchange for other, more saleable, commodities, even if he does not need them for any immediate consumption purpose. [Emphasis included] (Menger, 1981, p 260)

Menger argues that cattle as money was appropriate for regions like Ancient Greece but other forms of 'money' could emerge in different cultures, with the characteristics of marketability and relatively low costs of maintenance and transport. Moreover, the state could contribute to the 'moneyness' of some commodity by accepting and making payments in that form but Menger was clear that at this 'cattle' stage, the state is not necessary in the emergence of money.

Menger's account then proceeds through time, noting that economic progress gave rise to precious metals being preferred as money in many regions of the world. These became more marketable (metals were used more in advanced agricultural equipment, crafts, and early forms of industry), had very low maintenance costs, and transportation costs were low relatively to their high market price.

It is in the move from precious metals to coins that Menger alludes to a possible role of the state, including, most importantly, a quality (or "fineness") challenge which resembles

---

<sup>9</sup>Note that game theoretic and computational models of monetary emergence are considered in Chapter 5.

Akerlof's (1970) information asymmetry problem. Ordinary people cannot ascertain the quality of a metal so there is a question of how this is achieved. To that end, Menger mentions the creation of coins or bars with a "stamp by a public official or some reliable person on a metal bar guaranteed ... its degree of fineness" (Menger, 1981, p. 282). Put another way, a credible source is required to guarantee quality. In the next page, Menger writes "[g]overnments have ... usually accepted the obligation of stamping the coins necessary for trade ... [b]ut they have so often and so greatly misused their power ..." (ibid, p. 283).

This quality problem is discussed in more detail in Hodgson (1992). In this paper, Hodgson argues that Menger understated the quality problem by "saying that money is likely to take the form of precious metals, and these are 'easily controlled as to their quality and weight' (Menger, 1892, p. 255)" (Hodgson, 1992, p. 403). Moreover, Hodgson references Mirowski (1990) in stating the quality of money problem means "the 'value' of money is continually under threat from many devices and stratagems, from coin-clipping to the modern expansion of debt." (ibid, p. 404). Hodgson also notes that, while different to forgery, the state's debasement of money by inflation can also be construed as an (inter-temporal) money quality problem.

In Menger's monetary emergence narrative, this quality argument was mentioned vis-à-vis coinage but there is an argument that it can be applied to his earlier forms of money like cattle: Akerlof-like quality problems hold for many resources. Hodgson (2002a) writes that "debasement is a potential problem at the inception of money, not merely its development stage." (p. 119).

This quality problem is an important challenge to Mengerian monetary theories because, ultimately, it changes money as an institution from being related to - in Schultz's (2001) terminology - a coordination scenario to a collective action scenario<sup>10</sup>. The difference between these two is crucial in the theory of spontaneous order and will be discussed further in the next chapter. For now, we can state that coordination situations are in a sense self-regulating (they have qualities akin to Nash equilibria in game theory) and any organic institutions that emerge can be categorized as conventions. By contrast, Schultz's (2001) collective action situations typically require some sort of mechanism for agents not to cheat / defect / free ride on the rest of the population. These mechanisms might be informal in nature, e.g., enforcement of a social norm, or more formal like a fine from a third party.

It is tempting to dismiss Menger's money emergence narrative as inadequate when we factor in the quality challenge. However, this would be an over-reaction: Menger's narrative is interlaced with many supportive empirical examples. Rather, a more balanced view is that "there are good reasons to assume that money will be - to use Menger's

---

<sup>10</sup>Section 1.4.9

terminology - partly a ‘pragmatic’ rather than a purely ‘organic’ institution.” (Hodgson, 1992, p. 408).

It is worth emphasising that this debate about whether money should be left to its own devices or if it requires some sort of “intervention” (from the state or otherwise) is an excellent exemplar in a number of the debates described in the next chapter. Most notably, Hayek’s prescribed policy that money be denationalised (Hayek, 1976a), where privately-issued currencies can compete against each other, appears founded on the assumption that money has no quality problem. Moreover, as discussed in the Introduction, this question is important today in the context of cryptocurrencies (of which there are now many): reports of fraud and theft via hacking suggest the quality argument is very much alive and relevant.

### Michael Polanyi

As mentioned above, Polanyi<sup>11</sup> coined the phrase ‘spontaneous order’. A detailed evaluation of his framing of this concept is beyond the scope of this thesis but a history of spontaneous order should contain at least a summary of his contributions.

In his work, Polanyi differentiated between synchronic (stable) and diachronic (dynamic) systems in a way that resembles the difference between ergodic and non-ergodic, as defined in Chapter 2. Within diachronic systems, Polanyi further differentiated between ‘corporate’ (hierarchical) and ‘spontaneous’ (horizontal) social orders. This distinction resembles Hayek’s differentiation between ‘economies’ and ‘catallaxy’, which are discussed in more detail in the next chapter.

Within his framing of spontaneous order, Polanyi developed his notion of ‘polycentricity’, which is the idea that certain social systems (he used the example of the scientific community) tend to have multiple decision centres but also a common goal (such as the pursuit of objective truth). Polanyi used this notion to argue that socialist planning is inappropriate for economic systems because it tries to force a polycentric system to be ‘monocentric’ (with only one decision centre).

In his work, Polanyi also argued that much knowledge in society is tacit. Hayek mimics Polanyi here - this is also discussed in more detail in the next chapter.

It is worth noting briefly, that Michael Polanyi’s older brother, Karl, was also a social scientist who was concerned with these matters. Karl was more critical of decentralised markets than his brother and he became more aligned with socialism and monocentricity.

---

<sup>11</sup>Polanyi was viewed by many as a polymath, having held professorial positions both in chemistry and the social sciences.

### 3.3 Historical Influences on Hayek

As mentioned previously, Hayek plays the leading role in the next chapter. Here we look briefly at the historical influences on his work, drawing mainly from [Gray \(1998\)](#) and [Caldwell \(2004\)](#).

[Gray \(1998\)](#) differentiated between two major groups who influenced Hayek's work, which we might call the Austrians and non-Austrians. Of the former, Gray mentioned Menger, Friedrich von Wieser, who was Hayek's teacher, and Ludwig von Mises, who was one of Hayek's doctoral supervisors. Of the non-Austrians, Gray listed five main influences: Immanuel Kant, Ernst Mach, Karl Popper, Ludwig Wittgenstein, and Michael Polanyi. A comprehensive evaluation of these historical figures is well beyond the scope of this thesis but there are a number of points worth making as context for later discussions, notably with respect to Hayek's theory of mind<sup>12</sup>. Let us start with the non-Austrians.

At the heart of Hayek's ontology lies an approach to knowledge which [Gray \(1998\)](#) describes as "Kantian in that it accords a very great measure of creative power, which is neither a receptacle for the passive absorption of fugitive sensations, nor yet a mirror in which the world's necessities are reflected." (p. 8). As we will see below, this Kantian scepticism regarding the nature of knowledge in social systems is fundamental to Hayek's framing of spontaneous order. He also sees knowledge as decentralized and largely tacit in nature - put together, all of this presents a powerful challenge to would-be social planners.

While Hayek's developing framework of ideas was broadly consistent with Kant, his relationship to Mach's positivist philosophy was in opposition. Hayek believed the "mind is thus the order prevailing in a particular part of the physical universe - that part of which is ourselves." ([Hayek, \[1952\] 2018](#), p. 290, as quoted in [Gray, 1998](#), p. 10).

Hayek's emphasis was not only on order emerging among individuals but it was also on forms of order arising in the human mind. This is relevant context for this thesis because the institutions that immerse and emerge in the simulations reported in later chapters involve both immanent, intra-agent patterning and concomitant inter-agent order.

[Caldwell \(2004\)](#) notes that during his student days in Vienna, Hayek wrote a paper elaborating his ideas which was viewed positively by Adolf Stöhr (who succeeded Mach) and the philosopher Alois Riehl. This paper formed the beginnings of *The Sensory Order* ([Hayek, \[1945\] 2018](#)), published 25 years later.

[Gray \(1998\)](#) argues that Hayek and Popper differed in a number of respects but there were affinities between Hayek's view on the growth of knowledge and Popper's work

---

<sup>12</sup>Hayek's main text here is *The Sensory Order* ([Hayek, \[1952\] 2018](#)).

corresponding to an ‘evolutionary epistemology’. Specifically, Gray writes that for Hayek “the human mind is itself an evolutionary product and that its structure is therefore variable and not constant.” (Gray, 1998, p. 11). Curiously, while Gray presented Popper as an influence on Hayek, it seems Hayek got to this position first: he noted that even though *The Sensory Order* was first published in 1952, it was composed in the 1920s whereas the idea of an evolutionary epistemology appeared in Popper’s work later on<sup>13</sup>.

There is relevance, again, for this thesis because of the overlap between Hayek’s theory of mind and the mental models assumed for the agents in the ABMs presented in later chapters. These computational models were in part inspired by a type of mental model described in Holland (1975) and Holland et al (1986), which are commonly referred to as *classifier systems*. These have been subsequently used in Arthur (1994), Kirman and Vriend (2000), Vriend (2002), and Marimon et al. (1990). A key characteristic of these models is that there is some type of evolutionary process concerning rules going on *within* the agents’ own mental models.

Turning to Wittgenstein<sup>14</sup>, whose “influence runs deep ... [in] Hayek’s system of ideas” (Gray, 1998, p. 13), Gray writes “[t]here are ... many evidences that Wittgenstein’s work reinforced Hayek’s conviction that the study of language is a necessary precondition of the study of human thought, and an indispensable prophylactic to the principal disorders of the intellect.” (p. 13).

Both Hayek and Wittgenstein believed that language has a fundamental role in enabling coherence within - or the order of - the mind. Language is not merely a means of communication: words are associated with cognitive patterns which play an important role when individuals attempt to make sense of their environment and when they reason.

Moreover, in the context of Schultz’s (2001) coordination and collective action situations noted above, language is probably the purest form of the former. This means, most simply, that once a group of individuals has associated a word with some phenomenon (say, “ceiling”), then it is unlikely any individual would want to change this designation. This is not to argue that language does not evolve but words in general are close to conventions in a pure coordination game<sup>15</sup>.

---

<sup>13</sup>Hayek and Popper were good friends for many years and we might expect them to have influenced each other’s thinking. Caldwell’s (2006) paper *Popper and Hayek: Who Influenced Whom?* argues, however, that neither influenced the other and the relationship was based on mutual support of each other’s arguments. The problem with this conclusion, however, is that we cannot know counterfactual histories: how would Hayek’s work have developed without Popper, and vice versa? Surely this is un-answerable.

<sup>14</sup>who was a distant cousin of Hayek’s.

<sup>15</sup>There are a number of other problems with this simple account, e.g., the lack of univocality with most words, and the coexistence in populations of multiple dialects. A detailed discussion of this, however, is beyond the scope of this thesis.

In addition to this emphasis on language, Hayek's theory of mind is concerned with social knowledge as embodied in rules which exist in the human mind "which in some cases are necessarily beyond articulation by us" (Gray, 1998, p. 14). Furthermore, "for Hayek ... all our knowledge is at the bottom practical or tacit knowledge: it consists, not in propositions or theories, but in skills and dispositions to act in a rule-governed fashion." (ibid).

Once again we see a considerable amount of overlap with Holland et al (1986) who emphasise both the rule-following nature of human mental models and the idea that some of these rules exist sub-consciously. In fact, given the overlap between Hayek's theory of mind (especially his book *The Sensory Order*) and Holland et al (1986) it is remarkable that the latter never mentions Hayek's work. This overlap, which was identified by Vanberg (2018), is discussed further in the next chapter.

Looking at the fifth and final influence on Hayek's system of ideas, Gray (1998) notes Polanyi's insight "that, since much of the knowledge we use is inarticulate, we always know more than we can ever say" (p. 15). Furthermore, this "gives a wholly new twist to the argument for liberty from human ignorance." (ibid). In the debate about the extent to which human knowledge can be formalized and communicated for the purposes of central planning, Polanyi's point is particularly significant and it is used by Hayek to reinforce his argument against rational constructivism.

Turning now to the Austrian School, Gray (1998) summarises the relationship by stating that "Hayek has followed and developed the Austrian School." (p. 16). Included within this is "Hayek's extension of Austrian subjectivism about value to the whole realm of social objects [which] in no way represents a deviation from the positions of his mentors, Menger and von Mises." (p. 17).

Menger had a considerable impact on Hayek's thought, both directly through his publications but also indirectly by the fact Menger established the Austrian school of economics. Also, Menger stimulated Hayek's interest in economics (Caldwell, 2004, p. 139) and the latter came to study under Menger's disciples (notably von Wieser who encouraged and supported Hayek). We can say, therefore, that Menger enabled the environment in which Hayek learnt economics, and catalysed his interest in the subject.

Caldwell (2004) describes how Hayek's relationship with von Mises was one of mentee and mentor (the latter was about 12 years ahead of Hayek in his academic career). Moreover, von Mises employed Hayek and the two formed a personal relationship which lasted for about 10 years, between Hayek obtaining his undergraduate degree and later joining the London School of Economics. Von Mises appears to have been significantly responsible for Hayek moving away from his earlier socialist beliefs to one of classical liberalism. Moreover, Caldwell characterized the intellectual relationship by quoting Hayek who

said that he agreed with many of von Mises' (liberal) conclusions but disagreed with the arguments which supported them.

Von Mises thought that “economic laws were deductions from a few axioms about human action” and that “all of economic science can be derived from a proper specification of the nature of human action.” (Gray, 1998, p. 17). Hayek took issue with this apodictic-deductive approach in von Mises' work, which focused on the scientific method as applied to social science, e.g., in Hayek [1964] (2014), which is discussed in the next chapter.

Now that we have considered the history of spontaneous order and the influences on Hayek's thinking, let us look in more detail at his approach to spontaneous order.



# Chapter 4

## Hayek's Spontaneous Order

Hayek is a puzzle.

– *Hayek's Challenge* by Bruce Caldwell

The main aim of this chapter is to describe and evaluate Hayek's framing of spontaneous order, especially his theory of cultural evolution, which is closely related to organic institutions.

Hayek's work takes centre stage because, as noted in the previous chapter, he developed the most systematic and sophisticated approach to spontaneous order. His framework includes various ontological 'foundations', featuring a theory of mind, a perspective on devolved, tacit knowledge, and - interestingly for this thesis - many of the principles of Complexity Economics (CE) discussed in Chapter 2. Moreover, Hayek looked at various forms of spontaneous order in multiple domains, including 'the market order', his own theory of cultural evolution, and English common law.

Section 4.1 below begins the chapter with a health warning regarding Hayek's work, which is voluminous and which seemed to evolve over the six decades in which he wrote. Section 4.2 looks in more detail at the ontology, or foundation stones, of Hayek's work, including the crucial 'knowledge problem' and how his ontology corresponds with CE.

In Section 4.3 we take a slight detour, looking in more detail at [Ullmann-Margalit \(1978\)](#). This paper is important because it provides further background to Hayek's theory of cultural evolution, and also because the two computational models described in chapters 7 and 9 loosely correspond to Ullmann-Margalit's two 'moulds' of invisible hand explanations.

With these 'background' sections out of the way, Section 4.4 focuses on Hayek's discussion of 'catallaxy' and the order he believed was achieved - spontaneously - in free market systems. This is important in its own right but it also provides background to Hayek's theory of cultural evolution, which is the subject of Section 4.5<sup>1</sup>. This theory is particularly relevant to the thesis because 'cultural evolution' incorporates organic institutions such as conventions, and social norms.

Section 4.6 focuses on how Hayek distinguished between common law and legislation in the context of his theory of cultural evolution; and Section 4.7 discusses how Vanberg (1994b) tries to reconcile various strands of Hayek's work via the idea of *conditional evolution*. These two sections are related to the second research question, concerning "liberal legislation".

Section 4.8 concludes.

There are two points to note before proceeding. First, the literatures concerned with the various matters discussed in this chapter are enormous: the aim is to thread a way through this terrain and to highlight issues pertinent to the thesis while not doing violence to that research.

The second point is to repeat, for convenience, Hayek's definition of order since it plays a central role in this chapter. For Hayek, order is:

a state of affairs in which a multiplicity of elements of various kinds are so related to each other that we may learn from our acquaintance with some spatial or temporal part of the whole to form correct expectations concerning the rest, or at least expectations which have a good chance of proving correct. (Hayek, 1973, p. 36, emphasis removed)

As mentioned in the Introduction, one of the curious features of this definition is that it contains no commitment to spontaneous order necessarily giving rise to beneficial social patterns, including organic institutions. This is discussed further in Section 4.5 below.

## 4.1 Interpreting Hayek

Before we look in detail at Hayek's work, we should say a few words about how it seemed to change over time, and how people have interpreted him in different ways.

---

<sup>1</sup>One of the important criticisms of Hayek's work is from Buchanan who argued that Hayek used spontaneous market order as an inappropriate analogy for the emergence of institutions (Buchanan, 1986).

Hayek wrote prolifically over six decades and covered a broad variety of subjects<sup>2</sup>, all of which gave rise to an enormous and diverse body of work. Furthermore, his opinions appeared to change, which is to be expected of someone writing over several decades.

However, one of the key challenges of interpreting Hayek is that he did a poor job of advertising when his opinions had changed and explaining why. We can contrast this with researchers like North (also a recipient of the Nobel prize in economics) who flagged changes in his thinking and who showed “characteristic intellectual honesty” (Hodgson, 2017, p. 5) through much of his career. This appears less true of Hayek, which adds to the difficulty of interpreting his work.

Related to this, Hayek seems to contradict himself over and above the understandable changes in his thinking. For example, Caldwell (2004) notes how Hayek argued that policies aimed at income redistribution violated the rule of law while he also endorsed welfare measures that redistributed income “[w]ithin the covers of the same book” (p. 5).

Gray (1998) believes the most fundamental tension in Hayek's work is:

... between a conservative attachment to inherited social forms and a liberal commitment to unending progress. Hayek's distance from anything resembling traditional conservatism emerges most starkly when he commends progress, while acknowledging that “Progress is movement for movement's sake” (Hayek, [1960] 2006, p. 41). (Gray, 1998, p. 154)

There is one final point worth mentioning about consistency, which concerns Hayek's final book, *The Fatal Conceit* (Hayek, 1988): there is a question about how much Hayek participated in its writing. He fell ill in 1985 and it is “not clear how much of the book should be attributed to Hayek and how much to [William] Bartley.” (Caldwell, 2004, p. 317). Bartley was officially the book's editor and not a co-author. This question of authenticity is important because it is viewed as Hayek's last book so it might be interpreted as his ‘final say’ in various matters.

If we take all of the above observations together, and the fact Hayek was a controversial figure, we can appreciate how different researchers have interpreted him in varied - and sometimes contradictory - ways. For example Vanberg (2011) refers mainly to Hayek's earlier work (which emphasised classical liberalism) to discuss how Hayek's “liberal planning” laws could be considered within an evolutionary framework<sup>3</sup>; whereas Whyte (2019) claims a very different interpretation of Hayek, who “is tethered to deeply conservative opposition in politics.” (p. 160).

---

<sup>2</sup>Caldwell - one of Hayek's key biographers - referred to him as a polymath (Caldwell, 2006).

<sup>3</sup>We will see in Section 4.7 that Vanberg has a more complete command over Hayek's work than this implies, to the extent he attempts to reconcile seemingly incompatible parts of Hayek's work in Vanberg (1994b).

Given the apparent contradictions in his work we might conclude that Hayek is not worth paying attention to. This would be a mistake, especially in light of the orientation of this thesis around CE and spontaneous order. Hayek published *The Theory of Complex Phenomena* in 1964<sup>4</sup> (Hayek, [1964] 2014), which preceded most of the work published on social complexity by decades<sup>5</sup>. Moreover, while his work was inter-disciplinary, Hayek's orientation was around economics, and he was awarded<sup>6</sup> the Nobel prize in economics in 1974. These factors, combined with his sophisticated theorizing on spontaneous order, make Hayek a very relevant researcher for this thesis.

## 4.2 Hayek's Ontology

In framing Hayek's work, it is helpful to distinguish between the ontological foundations of Hayek's body of thought and his theory of cultural evolution. In this section we focus on the former, discussing Hayek's theory of mind; his orientation around knowledge; instincts, habits, and rules; and, finally his paper on complexity (Hayek, 1967).

### 4.2.1 Hayek's Theory of Mind

We saw in the previous chapter that Hayek's theory of mind has broadly Kantian origins and he follows Wittgenstein concerning the transmission of practical knowledge via social rules. However, Hayek's framing has "some entirely original features ... which it would be hard for Kant or Wittgenstein to accept." (Gray, 1998, p. 21). Most notable of these is Hayek's view of rules which operate sub-consciously and which cannot be known by the conscious mind. Moreover, Hayek argued that a type of evolutionary mechanism operates within the mind which ensures selection and modification of its set or rules.

This framing of the mind has profound implications for social theory. Gray (1998) does not hold back when he writes that it "entails the bankruptcy of the Cartesian rationalist project and implies that the human mind can never be fully understood itself" (p. 24).

One of the implications of the evolutionary view of rules in the mind is that, if we accept the idea that every individual has a unique history, the sets of rules operating within each mind will be equally unique, i.e., this points to radical heterogeneity in people's mental models. Another implication is that new experiences will result in changed sets of rules, i.e., we should expect mental models to evolve.

---

<sup>4</sup>Versions of this paper were published in 1964 and 1967 but Hayek completed the original manuscript in December 1961 (Hayek [1964] 2014, p. 257).

<sup>5</sup>Notably, the Economics Programme of the Santa Fe Institute was not set up until 1988.

<sup>6</sup>jointly with Gunnar Myrdal.

Gray (1998) discusses two broad implications of Hayek's theory of mind. First, a single mind cannot 'know' other minds because of the rules / knowledge contained in the sub-conscious that sense-making is partly based on: "inquiry can only be reflexive and never transcendental" (p. 25). Second, sub-conscious rules in others' minds will ultimately be hidden from view.

Moreover, in discussing rationality, "for Hayek rational calculation is inherently interstitial or supervenient - it fills the gap in a code of rules, resolves episodes of cognitive dissonance and aids judgment in applying norms." (ibid, p. 52). This is some distance from the assumption commonly held in Neoclassical Economics of substantive rationality. For Hayek, 'rationality' supervenes on sub-conscious foundations that are made up of rules.

The final point to note in this sub-section, mentioned briefly in the previous chapter, is that Hayek's theory of mind bears a great deal of resemblance to the discussion of mental models in Holland et al (1986) and to Holland's work on 'classifier systems' in general. This overlap between Holland and Hayek has been noted by others. For example, Vanberg discusses it in Section 11 of his Introduction to *The Sensory Order* (Vanberg, 2018, pp. 78-86). Here, Vanberg quoted Mark Miller who writes that "[a]lthough Holland developed his ideas without knowledge of Hayek's work, the two scholars are wonderfully complementary." (Miller, 1996, p. 59).

Holland's approach to mental models heavily informs those assumed in the computational models developed later in the thesis.

## 4.2.2 Hayek's Knowledge Problem

Many researchers have noted that Hayek's orientation around knowledge in social systems is crucial to understanding his body of thought, e.g., Gray (1998) and Caldwell (2004). The former argues that Hayek's framework of spontaneous order contains three principles: (i) his "invisible hand thesis" (p. 33); (ii) a form of cultural evolution of rules; and (iii) his understanding of the role of knowledge in social systems. Hayek referred to the first two of these as his "twin pillars". We can understand Gray's third component, concerning knowledge, as foundational to Hayek's ontology. Indeed, Hayek deemed the *division of knowledge* in economics "at least as important as ... the division of labour." (Hayek, 1937, p. 49).

It is important to emphasise that we ought not to discuss Hayek's approach to knowledge independently of his theory of mind. Rules contained in the mind - whether we are conscious of them or not - are part of society's body of knowledge. Boehm (1994) writes that, for Hayek, "knowledge of the social world is primarily embodied in practices and skills rather than in theories and that a great deal of this practical knowledge cannot be articulated and communicated." (p. 300), i.e., it is mostly tacit.

Related to this, “Hayek is at great pains to point out that ... dispersed knowledge is precisely not theoretical or technical knowledge, but practical knowledge of concrete situations” (Gray, 1998, p. 36). This is not to argue that knowledge cannot be theoretical or technical: Hayek is emphasising that knowledge in social systems is predominantly devolved, informal, and practical. To understand this more clearly, Hayek writes:

The skipper who earns his living from using otherwise empty or half-filled journeys of tramp-steamers, or the estate agent whose whole knowledge is almost exclusively one of temporary opportunities, or the arbitrageur who gains from local differences in commodity prices - are all performing eminently useful functions based on special knowledge of circumstances or the fleeting moment not known to others. (Hayek, 1976b, p. 80)

From a whole systems point of view, we might ask, therefore, how do we: (i) make maximal use of this dispersed, devolved knowledge? and (ii) enable it to usefully develop / accumulate?

It was clear to Hayek that this *ontology of knowledge* means that central planning along the lines of socialism is futile: to achieve this, devolved, tacit knowledge would have to be formalized, collated, and centralised; and whole-system decisions would have to be made, transmitted, and enacted instantaneously (else some marginal change in local knowledge might require a different system-wide ‘solution’). The obstacles to this should be clear from the discussion of Hayek’s theory of mind and his description of the nature of knowledge. As Gray (1998) puts it, such “public planning cannot avoid yielding calculation chaos.” (p. 36) even if the practical problems of centralised knowledge-gathering and execution could be overcome.

By contrast, Hayek argued that market systems provide the mechanisms through which (one type of) spontaneous order can emerge which will maximize the use of dispersed knowledge. Such market systems are exemplars of spontaneous order and are discussed in more detail below.

### 4.2.3 Instincts, Habits and Rules

Rules play an important role in Institutional Economics, and they are central to the mental models used in the Agent-Based Models (ABMs) described later in the thesis. Here we will look briefly at how Hayek used this term, a criticism of this use, and how it relates to habits in Hayek’s framework.

Hayek’s reference to ‘rules’ was mentioned above. It is “a central concept in [his] mature theory of social evolution...” (Hodgson, 2006b, p. 16). For Hayek the meaning of ‘rule’ was that of a “regularity of the conduct of individuals ... irrespective of whether such a

rule is 'known' to the individuals ..." (Hayek, 1967, pp. 66-7, as cited in Hodgson, 2006b, p. 16).

Hodgson (2006b) argues that the problem with this definition is that it is too broad. "It commits Hayek, for example, to regard all regular bodily functions as resulting from the observances of rules." (Kley, 1994, p. 44, as cited in Hodgson, 2006b, p. 16). Moreover, Hodgson argue that the lack of "an adequate explanation of the origin and impetus behind rules [means] his explanation is insufficiently Darwinian." (ibid, p. 17).

This raises questions of Hayek's framing which are both important and unanswered, like "[w]hat are the mechanisms involved in ... the transformation of a rule into an act?" and "[w]hat sustains the rule and gives it some durability through time?" (Hodgson, 2006b, p. 17). Indeed, both these questions provide good motivations for the models described in later chapters: the agents in these ABMs use mental models which provide both a framing of the context in which they operate and also a means to make decisions, i.e., they use both perceived patterns and reasoning to make choices.

As noted in the Introduction, the models developed below also explore the process of habituation. Unfortunately, for Hayek the concept of habit "is not prominent because it is also subsumed within his overly extensive concept of a 'rule'." (Hodgson, 2006b, p. 18). From what Hayek wrote, we can discern that he did not recognise that "habit is a necessary foundation for conscious reflection" and that "his casual use of the term suggests a conception of habit as settled behavior, more than a propensity of disposition." (ibid).

The conclusion we draw here is that Hayek's reference to 'rules', and his understanding of habits, are somewhat ambiguous.

#### 4.2.4 Hayek on Complexity

Caldwell (2004) describes how Hayek's ideas about knowledge, exemplified by his article *The Use of Knowledge in Society* (Hayek, 1945), at first "convinced no one, or virtually no one" (Caldwell, 2004, p. 10).

This problem illustrates a key tension between Hayek and the economics orthodoxy throughout his career. He saw the role of the 'economy'<sup>7</sup> as coordinating dispersed knowledge whereas the consensus among economists was (and remains today) that economics is concerned with the allocation of scarce resources<sup>8</sup>. Indeed it is noteworthy

---

<sup>7</sup>Hayek's preferred term here is 'catallaxy' rather than 'economy'. This is discussed further below.

<sup>8</sup>These two are clearly not mutually exclusive. The issue is more one of emphasis in research.

that Hayek's 1974 Nobel citation did not refer to knowledge, which was at the heart of his ontology<sup>9</sup>.

When considering the poor reception of his theories of knowledge, [Caldwell \(2004\)](#) writes that Hayek encountered a problem that "everyone of his opponents claimed to be doing 'real' science." (p. 10). In a move which mimicked "Menger before him, he turned to the study of methodology to make his case." (ibid). [Caldwell \(2004\)](#) stresses that Hayek's "first love was not methodology" (ibid, p. 11) but he felt compelled to develop this side of his work to counter his critics. It is this which led him to consider the problem of "complexity".

As part of this methodological turn, Hayek was influenced by General Systems Theory (GST) and, related to this, in 1961 he gave a series of lectures entitled "A New Look at Economic Theory". In a letter to Popper the year before, Hayek writes that these lectures "began with an attempt to restate my views of the nature of economic theory, and the conception of higher level regularities which I then formed continues to occupy me and seems fruitful far beyond the field of economics. I suspect it is really what [von] Bertalanffy with his General Systems Theory was after..." ([Hayek, 1960](#), as cited in [Caldwell, 2004](#), p. 307).

Ludwig von Bertalanffy's most famous text in GST is [von Bertalanffy \(1968\)](#); however, he and Hayek were in communication well before that<sup>10</sup>. Moreover, von Bertalanffy published a number of books and papers on the subject of GST in German (Hayek's first language) prior to 1961.

In addition to von Bertalanffy, who Hayek "was, perhaps, closest to" ([Caldwell, 2004](#), p. 362), Hayek also "cited Warren Weaver on organized complexity, John von Neumann on the logic of automata, and Norbert Wiener on cybernetics." (ibid).

There are two points worth emphasising here. The first is Caldwell's references to von Bertalanffy, Weaver, von Neumann, and Wiener, all of whom are associated with the origins of the modern day complexity sciences<sup>11</sup>, e.g., all are referenced in [Castellani and Gerrits' \(2021\) \*Map of the complexity sciences\*](#)<sup>12</sup>. Second, Hayek saw his project in the early 1960s as going beyond - or developing - GST.

---

<sup>9</sup>The citation reads "for [Hayek's and Myrdal's] pioneering work in the theory of money and economic fluctuations and for their penetrating analysis of the interdependence of economic, social and institutional phenomena."

<sup>10</sup>[Caldwell \(2004\)](#) notes that "von Bertalanffy ... had provided comments on *The Sensory Order* [published in 1952] when it was in draft form." ([Caldwell, 2004](#), p. 362).

<sup>11</sup>From a CE point of view, the relationship Hayek (who was an economist) had with these giants of the complexity sciences is both fascinating and important.

<sup>12</sup>According to [Rosser \(1999\)](#), Hayek also "had significant communication with Ilya Prigogine and Hermann Haken." (p. 185).



In the same year as his “New Look” lectures, Hayek completed the first manuscript of *The Theory of Complex Phenomena*, in December 1961 (Hayek, [1964] 2014)<sup>13</sup>. It is worthwhile examining and evaluating parts of this paper in some detail, for four reasons: (i) in general it supports the CE orientation of this thesis (the overlap between this paper and the description of CE in Chapter 2 is substantial though far from total); (ii) related, the paper discusses a number of challenges in attempting to deal with social complexity via a scientific approach, which were not discussed in Chapter 2; (iii) it helps inform the mental models used in the ABMs presented in later chapters; and (iv) it provides important ontological foundations for understanding Hayek's theory of spontaneous order.

The rest of this sub-section will look at four of the themes in Hayek [1964] (2014):

1. pattern recognition;
2. simple versus complex systems;
3. incomplete data; and
4. the implications for theories of social structures.

### Pattern Recognition

Extending his theorizing about the mind, Hayek [1964] (2014) emphasises the importance of pattern recognition, which allows for the identification of “regularities of nature ... recognised intuitively by our senses ... without having to resort to intellectual operations.” (p. 258). This distinction between intuition and intellect obviously pre-dates Hayek's paper but it is worth noting that Hayek's discussion preceded Daniel Kahneman's (2012) related references to System 1 (intuitive) and System 2 (intellectual) information processing by about 50 years.

Delving deeper in to Hayek's thoughts on mental models and complexity, Vanberg (2018) discusses one of Hayek's unfinished papers *Within Systems About Systems* (pp. 361-381). Related to this, in his Introduction chapter Vanberg (2018) notes “that Hayek's social theory, revolving around the ‘twin ideas of evolution and spontaneous order,’ can be understood as an application of the ‘systems within systems’ concept as a general explanatory scheme.” (p. 81). This scheme fits neatly into the notion of nested complex systems discussed in Section 2.1.1.1 above as well as the main aim of this thesis, which is to explore the emergence of organic institutions.

---

<sup>13</sup>Caldwell (2004) states that “Hayek's writings about spontaneous orders share certain similarities with recent work in complexity theory, the study of self-organizing complex adaptive systems ... it was his publication of *The Sensory Order* [in 1952] that his interest in the ubiquity of such phenomena became apparent.” (p. 361)

The second way Hayek's discussion is helpful concerns pattern recognition in science, which is discussed further below under theories of social structures.

### Simple Versus Complex Systems

In discussing degrees of complexity, we see most clearly the affinity between Hayek's ontology and the complexity sciences. He distinguished between simple and the present day complex systems, noting that physicists often deal with simple, mechanical problems whereas the study of social systems appears inherently more complex (he was clear this does not mean that all natural phenomena are non-complex). Complexity involves "[t]he 'emergence' of 'new' patterns ... [and] means that this larger structure as a whole will possess certain general or abstract features which will recur independently of the particular values of the individual data..." (Hayek, [1964] 2014, p. 262). As stated previously, the idea of emergence, which is non-reducible to constituent parts is viewed by many researchers in the field as the *sine qua non* of the complexity sciences (e.g., Waldrop, 1993; and Mitchell, 2011).

On the related subject of 'wholes', Hayek discusses the need to draw "a 'partition boundary', [which] will be determined by the consideration of whether we can thus isolate recurrent patterns of coherent structures of a distinct kind" (Hayek, [1964] 2014, p. 262). Here, the discussion of boundaries resembles that seen in *Signals and Boundaries* (Holland, 2014). Furthermore, Hayek refers to structures in which "a complex pattern has produced properties which make self-maintaining the structure showing it" (ibid), which preceded Maturana and Varela's [1972] (1980) reference to 'autopoiesis' by a few years<sup>14</sup>.

One criticism that can be levelled at Hayek's discussion here is his attempt to quantify complexity as "[t]he minimum number of elements of which an instance of the pattern must consist in order to exhibit all the characteristic attributes of the class of patterns in question" (Hayek, [1964] 2014, p. 260). There is now a large literature devoted to measuring complexity (see Wiesner and Ladyman, 2019, pp. 12-15 for a good discussion) and in this context Hayek's definition looks naïve. Notably, his definition omits any consideration of the relationships between 'elements' within a complex system, which might include non-linear feedback loops. Hayek's definition would, for example, lead him to view a glass of water as vastly more complex than Lorenz's equations (which include only three variables).

Nonetheless, Hayek's attempt at quantifying complexity aside (this is far from resolved anyway), his analysis of 'simple' and 'complex' systems still holds if we assume a simple, qualitative distinction between them.

---

<sup>14</sup>We should note that Prigogine had discussed the related concept of 'dissipative structures' in the 1950s.

### Incomplete Data

Hayek's discussion here fits his emphasis on tacit knowledge being located locally among an enormous number of agents. Incomplete data gives rise to theories which are highly conjectural, with "small empirical content ... and [where] the possibility of falsifying [them] correspondingly small." (Hayek, [1964] 2014, p. 264). Moreover, "[t]he advance of science will thus have to proceed in two different directions: while it is certainly desirable to make our theories as falsifiable as possible, we must also push forward into fields where, as we advance, the degree of falsifiability necessarily decreases. This is the price we have to pay for an advance into the field of complex phenomena." (ibid).

In this discussion of the difficulties of theorizing in the context of data that is far from complete, it is worth mentioning the potential value of computational models.

Hayek was at pains to emphasise the limitations of human cognition relative to the vast amount of knowledge in social systems (notably in *The Sensory Order* and *The Fatal Conceit*), to which we can add his discussions of complexity in *The Theory of Complex Phenomena*. We can also add the problems inherent in theorizing about emergent phenomena, which Hayek [1964] (2014) mentions.

We might think of a computational model as a 'cognitive annex' whereby information concerning complex phenomena can be processed by computers in a way that is well beyond the capacities of the human mind. Computational modelling can achieve in seconds (accurately) what people would need far longer to do (less accurately). However, this is not only about speed: spending months manually computing a single simulation has an enormous opportunity cost which people will generally avoid, so ABMs allows us to deal in phenomena we would not otherwise consider (or observe).

Regarding modelling, it is perhaps worth noting that this is one of the weaker areas of Hayek [1964] (2014). He wrote that "[t]he systematic construction of ... new patterns is the business of mathematics" (p. 259). Given that in 1961 mathematics was the dominant form of modelling and that computation was far from ubiquitous, Hayek can be forgiven. It does not require a great leap of imagination to see the construction of new patterns as also being the business of computational modelling.

### Theories of Social Structures

In the context of *Pattern Prediction*, Hayek [1964] (2014) considers the problem that "scientific procedure demands that we should find a theory of sufficient simplicity to enable us to derive from it predictions of particular events." (p. 263). This is central to Hayek's criticism of his opponents who appeal to "science" in their work (naïvely in Hayek's opinion).

The appeal of simplicity is, Hayek notes, probably in part because of physicists' success in dealing with "simple" phenomena and developing equivalently simple laws. However, Hayek argues that "a simple theory of phenomena which are in their nature complex ... is probably merely of necessity false - at least without a specified *ceteris paribus* assumption, after the full statement of which the theory would no longer be simple." (ibid). We might add to this that some pattern might be observable under certain conditions but not others. This idea is significant for the simulation results presented in later chapters because we see that organic institutions emerge only when the conditions are sufficiently enabling.

Hayek also helpfully distinguishes between predicting the appearance of a class of some pattern versus the prediction of a specific instance. "The distinction [between class and instance] assumes ... greater importance when we turn from the relatively simple phenomena with which the natural science deal, to the more complex phenomena of life, of mind, and of society..." (Hayek, [1964] 2014, p. 260). Simpler systems, when sufficiently understood, allow for concrete (spot) predictions whereas in complex systems we might expect some general phenomenon to arise, the precise nature and timing of which will be uncertain.

Darwin's theory of evolution by natural selection is used in Hayek (1967) as an example of a theory which provides general but not specific predictions about the future because we cannot know ahead of time the exact nature of future mutations.

This discussion of class versus instance is relevant for this thesis because we will see in the simulations based on the first model that a market (as a class of phenomenon) emerges when the environment is sufficiently enabling but we cannot predict beforehand where and when the specific market will manifest.

Hayek applied his distinction between class and instance prediction to *theories of social structures*. Looking specifically at economics, he writes that "economic theory is confined to describing kinds of patterns which will appear if certain general conditions are satisfied, but can rarely if ever derive from this knowledge any predictions of specific phenomena." (Hayek, [1964] 2014, p. 270). He gave the example of simultaneous equations representing supply and demand curves, post-Walras, as operating *as if* we know all the parameters and variables of the model. But he quotes Pareto who stated that such equations are not meant to predict exact prices because it would be "absurd" to believe we could find all the relevant data.

This begs the question of what value is gained from general insights into behaviour if we cannot make concrete predictions. Hayek's answer to this question relates to some of the simulation results we see later in the thesis regarding an *enabling environment*: "Since the theory tells us under which general conditions a pattern of this sort will form itself, it will enable us to create such conditions and to observe whether a pattern of the kind

predicted will appear.” (ibid, p. 271). Indeed, we will see below that this also relates to Hayek’s earlier “classical liberal” work in which he saw the role of “liberal legislation” as providing the conditions in which ‘free’ markets can operate.

### Conclusion

Overall, when considering the tension between the complexity of social systems and the limitations of human cognition, Hayek writes that “[i]t is high time ... that we take our ignorance more seriously.” (ibid, p. 275). This appeal to humility is directed at what he refers elsewhere (disparagingly) to as ‘scientism’; it means that “we must get rid of [this] naïve superstition that the world must be so organized that it is possible by direct observation to discover simple regularities between all phenomena and that this is a necessary presupposition for the application of the scientific method.” (ibid).

Equivalently, Hayek discusses the idea of ‘laws’ in theoretical science, which in the usual conception of the word (a deterministic relationship between two or three phenomena) “has little application to the theory of complex phenomena.” (ibid, p. 276). This would suggest that the aim in social science is not the identification of laws but the creation of, and debate about, generalized theories.

## 4.3 Invisible Hand Explanations

Before we discuss Hayek’s ‘market order’ and his theory of cultural evolution in more detail, it is helpful to look at two different types of “invisible hand explanations”. These are attributable not to Hayek but to [Nozick \(1974\)](#) and [Ullmann-Margalit \(1978\)](#).

There are two reasons for this short detour. First, the distinction between [Ullmann-Margalit’s \(1978\)](#) two “moulds” of invisible hand explanations provides helpful background to criticisms of Hayek’s theory of cultural evolution, described in Section 4.5.

The second reason is that the two ABMs set out in chapters 7 and 9 below correlate strongly with [Ullmann-Margalit’s \(1978\)](#) two “moulds” of invisible hand explanations: the Market Emergence Model of Chapter 7 is approximately aligned with her “aggregate” mould; and the Property Rights Model of Chapter 9 is associated with her “functional-evolutionary” mould<sup>15</sup>.

Invisible hand explanations are related to various types of spontaneous order and not only to organic institutions. The term “invisible hand” was originally used in [Smith \[1759\] \(2011\)](#) and also [Smith \(1776\)](#). The latter use corresponds to the “market order” described in the next section.

<sup>15</sup>[Nozick \(1974\)](#) referred to “equilibrium” and “filtering processes” (p. 21), which [Ullmann-Margalit \(1978\)](#) developed.

Ullmann-Margalit (1978) provides five examples of invisible hand explanations: the creation of money in the banking system; “Nozick’s account of the rise of the so-called ultra-minimal state” (p. 264); the development of early forms of money; Adam Smith’s “equilibrium pricing system that develops within the perfectly competitive market” (p. 270); and segregation in cities, following Schelling (1969).

Below we first look at Ullmann-Margalit’s discussion of eligibility for (and characteristics of) invisible hand explanations. We then discuss how her two ‘moulds’ correspond to the ABMs set out later in the thesis.

### 4.3.1 Eligibility as an Invisible Hand Explanation

What makes an explanation for a “social pattern or institution” (Ullmann-Margalit, 1978, p. 267) eligible to be classed as ‘invisible hand’ in nature? Her answer to this question has a direct bearing on this thesis. She states:

It seems to me to be quite clear ... that the onus of the explanation lies on the process, or mechanism, that aggregates the dispersed individual actions into the patterned outcome: it is the degree to which this mechanism is explicit, complex, sophisticated - and, indeed, in a sense unexpected - that determines the success and interest of the invisible hand explanation in question. (Ullmann-Margalit, 1978, pp. 267-8)

She continues, stating that “the more structured and complex the pattern, the greater the challenge it poses to whoever proposes to explain it invisible-handedly” (Ullmann-Margalit, 1978, p. 269), arguing also that such complex patterns can be associated, mistakenly, with planned order. Her discussion of this point was referenced in the previous chapter.

Furthermore, by “invisible hand process”, she means a “mechanism which takes as ‘input’ the dispersed actions of the participating individuals and produces as ‘output’ the overall social pattern ... it is this process which bears the explanatory brunt of invisible-hand explanations.” (p. 270).

In light of the weight she places on invisible hand mechanisms, perhaps the most surprising part of Ullmann-Margalit’s (1978) paper is that she thinks it is not worth “looking for generalizations over these [invisible hand processes], or to seek to unearth [their] ‘logic’”, preferring to focus on constraints (or characteristics) which she views as “the backbone of good invisible hand explanations” (Ullmann-Margalit, 1978, p. 271) - these are discussed below. The two moulds she refers to are not invisible hand explanations themselves: they are groupings within a typology.

This devaluing of generalizations is very contentious, least of all because Ullmann-Margalit gives no compelling reason for not seeking them. In light of the broader range of invisible hand explanations (than only those pertaining to organic institutions) discussed in [Ullmann-Margalit \(1978\)](#), it is possible that Ullmann-Margalit thinks the range of explanations would be so wide as to make generalizations worthless. This would be a reasonable argument (if supported by empirical evidence); however, this is conjecture about her opinion - she did not state this.

In the context of organic institutions - the focus of this thesis - the discussion of the rationale for studying their emergence in the Introduction constitutes further criticism of Ullmann-Margalit's preference to avoid generalizations. Organic institutions are important phenomena and understanding the nature and timing of their emergence can help inform what steps we might take to catalyse helpful types in the future, e.g., by providing an enabling environment.

### 4.3.2 Characteristics of Invisible Hand Explanations

This criticism aside, [Ullmann-Margalit \(1978\)](#) offers three constraints on (characteristics of) invisible hand explanations<sup>16</sup>:

1. the initial stage of their origination consists of “nothing but the private intentions, beliefs, goals, and actions of the participating individuals, in a specified setup of circumstances.” (p. 271). Let us refer to this as the “individual-based characteristic”.
2. the various stages of any explanation should “sound like a description of the ordinary and normal course of events. It cannot hinge on the extraordinary and the freaky” (ibid). Put another way, an explanation should seem reasonable in its local context. We refer to this as “the normalcy condition”.
3. invisible hand explanations typically involve surprise “at their very existence” (ibid); but, once properly understood, there are “no further surprises within the explanation itself.” (p. 272).

The first of these locates Ullmann-Margalit's invisible hand explanations in the long tradition of spontaneous order. It fits with the definition provided by Ferguson, cited in the previous chapter, and sits comfortably alongside [Smith's \(1776\)](#) invisible hand and Hayek's subjectivism.

---

<sup>16</sup>Strictly speaking, [Ullmann-Margalit \(1978\)](#) includes two constraints followed by an additional “comment” on the nature of surprises. The third is significant in the context of emergent properties so it is included here as a third characteristic.

The second of these characteristics appears innocuous at first and could be interpreted as an appeal to common sense. However, in the context of economics, it is a strong methodological statement.

This second characteristic appears related to the discussion of instrumentalism and realism (Section 2.3.3.10). Ullmann-Margalit's valuing of a 'grounded' normalcy appears consistent with a realist's point of view. However, we should not take this too far because Ullmann-Margalit did not discuss this matter in any detail. The point here is that her emphasis on normalcy has, let us say, a flavour of realism.

In the context of mental models and rationality, it is worth highlighting Ullmann-Margalit's (1978) Endnote 10 (p. 288). She asks whether rationality would "take care of this normalcy condition?" (ibid). She answers that it would, provided "agents act rationally relative to their beliefs and utilities ... [and] relative to the prevailing (and possibly changing) circumstances." (ibid). Ullmann-Margalit's framing here appears to be equivalent to Simon's (1976) *procedural rationality* and not *substantive rationality*. It also sits comfortably alongside the approach taken in the ABMs developed for this thesis, in which agents reason under conditions of uncertainty and change.

The third point listed above is particularly interesting in the context of the complexity sciences, specifically the distinction between *epistemological* and *ontological emergence* (Ladyman, Lambert and Wiesner, 2013). The former is associated with the idea that an emergent property is predictable in principle (Laplace's (1814) demon would be able to anticipate it) whereas the latter is not. Both are associated with surprises but the crucial point here is that the former can be understood, at least in principle. This is very much consistent with Ullmann-Margalit's idea that invisible hand explanations can be surprising but, ultimately, understandable.

Ullmann-Margalit's three characteristics provide a framework for evaluating the results of the ABMs presented in later chapters. The idea is that a coherent and compelling mechanism which explains the emergence of organic institutions features all three characteristics.

We will see that the first was included in the models at the design stage: we are interested in the collective behaviour of individual agents. The second and third are essentially challenges for these models: can (a) compelling invisible hand explanation(s) be found for the emergence of organic institutions? and can this explanation move us from surprise to comprehension?

Now that we have discussed some of the territory highlighted in Ullmann-Margalit (1978), let us consider her two moulds of invisible hand explanations: the aggregate mould and the functional-evolutionary mould.



### 4.3.3 The Aggregate and Functional-Evolutionary Moulds

In discussing her aggregate and functional-evolutionary moulds, [Ullmann-Margalit \(1978\)](#) helpfully distinguishes between explanations of how a phenomenon *comes in to being*, and how it is *maintained* in a population.

By aggregate mould of invisible hand explanation, [Ullmann-Margalit \(1978\)](#) writes that these are associated with “one particular mode of emergence rather than ... intentional design.” (pp. 283-4). This means that the explanation is bound up with the origination of the “social pattern or institution” (*ibid*).

This description by Ullmann-Margalit is still somewhat vague so we will make use of [Gedeon's \(2015\)](#) helpful interpretation to make it clearer. According to Gedeon, the aggregate mould is about “social order as an unintended, spontaneously emerging consequence of individual choices. Spontaneous aggregation of individual actions leads to social *coordination* of individual actions.” (p. 3, emphasis included).

The emphasis here is on coordination, and we can further tie this notion to it being self-sustaining, i.e., a situation in which third party enforcement is not required because the nature of the interaction means the agents do not want to change for their own reasons. This appears to be consistent with Schultz's coordination situations.

For Ullmann-Margalit, the functional-evolutionary mould includes a:

...functional analysis of the pertinent social institution, conjoined with the concomitant evolutionary apparatus presumed to supply the missing causal link, that constitute - together - the invisible hand explanation of that institution. ([Ullmann-Margalit, 1978](#), p. 282)

Put another way, the phenomenon contains two factors: a social function it performs; and there is an evolutionary explanation for it being maintained.

Regarding the first characteristic here, a phenomenon performs “a useful service to the social unit incorporating it.” ([Ullmann-Margalit, 1978](#), p. 284). The reference to “social unit” is important: a function is performed for a population.

Furthermore, this mould is defined by there being some type of ‘evolutionary’ explanation for how the phenomenon is sustained in a population. The:

process of *selection* is supposed to be a non-man-made one: it is visualized as a large-scale evolutionary mechanism that as it were scans the inventory of social patterns and institutions at any given period of time and screens

through to the next those of them that are best adapted to their (respective) roles. (Ullmann-Margalit, 1978, p. 282, emphasis included)

A curious feature of the functional-evolutionary mould is that the origination of such phenomena can be planned. Its 'invisible hand' nature is associated with 'invisible forces' which ensure the phenomenon is maintained<sup>17</sup>.

#### 4.3.4 Links to Computational Models

The first model below, referred to as the Market Emergence Model, is aligned with Ullmann-Margalit's (1978) aggregate mould of invisible hand explanations and Schultz's coordination situations.

We find that, in the default simulations, a single market emerges as an unintended consequence of agents' actions and in the context of an enabling environment. This market is equivalent to a solution in a coordination situation (it is 'non-pure' because agents benefit by being closer to the emergent market). It is both a type of organic institution and a convention.

The second model is motivated by Schultz's (2001) discussion of the fact that if we assume substantive rationality, an "exchange situation" is a collective action situation and not a coordination situation because "there are no normative constraints precluding force or fraud" (p. 67). The first model is aligned with Schultz's first category through the assumption of property rights: agents respect each other's resource holding as their property. In the second model we drop this assumption: the second model is essentially identical to the first but agents can steal from each other.

This second model explores whether an organic institution can emerge in the form of property rights. Moreover, we are particularly interested in whether a functional-evolutionary explanation à la Ullmann-Margalit (1978) is seen in the simulations. We find this is indeed the case when the agents' mental models change as a result of reinforcement learning and habituation: the agents converge (approximately) on two different strategies and we find that the strategy which includes property rights proves more likely to survive than the alternative.

Furthermore, we find that legal rules can help catalyse property rights when they do not emerge endogenously, which fits with Ullmann-Margalit's (1978) idea that the origination of the institution does not matter but the selection and maintenance aspect does.

---

<sup>17</sup>This is similar to but not quite the same as planning that might enable spontaneous order, which was discussed in the Introduction. Ullmann-Margalit is referring here to planned rules.

Now we have discussed key features of Hayek's ontology, and Ullmann-Margalit's invisible hand explanations, let us turn to what Hayek and others view as the paradigm case of spontaneous order, that produced by so-called 'free markets', which relates to what Hayek refers to as 'catallaxy'.

## 4.4 Catallaxy and the 'Market Order'

This type of spontaneous order is seen as one of several 'paradigm' cases (others include money, language, and - more contentiously - English common law). Moreover, Hayek's work was generally concerned with economics and the market order provides not only an example of spontaneous order, it can also be viewed as an analogy for spontaneous order seen elsewhere, notably in "the emergence of institutional structure itself." (Buchanan, 1986, p. 76).

It is helpful to first look at how Hayek differentiated between an economy and 'catallaxy' (the latter is attributed to von Mises). An *economy*:

...in the strict sense of the word in which a household, a farm, or an enterprise can be called economies, consists of a complex of activities by which a given set of means is allocated in accordance with a unitary plan among the competing ends according to their relative importance. (Hayek, 1976c, p. 107)

By contrast, *catallaxy* is:

the order brought about by the mutual adjustment of many individual economies in a market. A catallaxy is thus the special kind of spontaneous order produced by the market through people acting within the rules of the law of property, tort and contract. (Hayek, 1976c, pp. 108-9)

In Hayek's view, economies are associated with deliberate organisation and the coordination of individuals' activity toward some specified end. By contrast, catallaxy is about spontaneous order without any common purpose. Moreover, it is worth highlighting, in the context of the co-adaptation observed in the simulations reported below, Hayek's reference to "the mutual adjustment of many individual economies" in his definition of catallaxy.

Furthermore, for Hayek catallaxy is related to his emphasis on knowledge: "[t]he important point about the catallaxy is that it reconciles different knowledge and different purposes which, whether the individuals be selfish or not, will greatly differ from one person to another." (Hayek, 1976c, p. 110). This is a crucial part of Hayek's body of

thought: the relationship between the order brought about by markets and dispersed, heterogenous knowledge among people with different purposes.

Hayek uses this distinction between an economy and catallaxy to illustrate his argument against central planning. The reference to a “national economy” creates confusion over the nature of order and encourages people to think that the economic activities of a nation (for example) can be planned and organised in the same way “a household, a farm, or an enterprise” can.

Two further points are worth highlighting. First, order is clearly referred to here by Hayek in a positive light: in Chapter 10 of Hayek (1976c) he discusses how catallaxy is related to peace and prosperity in the “Great Society”. There is a large literature concerned with the normative content of Hayek’s framing, and this is discussed briefly in Section 4.5.5 below.

Second, in the quote above containing his definition of catallaxy, Hayek is clear that order should be viewed in the context of appropriate forms of law. This is the subject of Section 4.6 below.

The above statements by Hayek raise the question of the *mechanisms* through which market-based order is maintained. One of these is a body of law whose primary purpose is the protection of individuals’ property. A number of researchers have argued (e.g., Gray, 1998; and Vanberg, 1994b) that Hayek is writing here as a classical liberal whereby a key role of the state is to protect individuals from others.

Another mechanism, related directly to the operation of markets, is about market “[c]ompetition as a discovery procedure” (Hayek, 1979, p. 70), which relates to the roles played by market prices and entrepreneurs. Part of this framing is the familiar ‘story’ of free market economics in which prices equilibrate supply and demand for goods and services. For Hayek, however, this is only part of the story - he also draws on his discussion of dispersed knowledge:

the price mechanism operates as a medium of communicating knowledge which brings it about that the facts which become known to some, through the effects of their actions on prices, are made to influence the decision of others. (Hayek, 1976c, pp. 108-9)

Prices, therefore, act to communicate knowledge between people who need not know each other. This led Gray (1998) to state that “the market is a paradigm case of a social institution having an epistemological role” (p. 41).

Moreover, the entrepreneur plays an important role through which market order is maintained: prices communicate opportunities to make a profit, which entrepreneurs exploit.

The final point to note here is that market order is not the same as the ergodic world of General Equilibrium theory (Arrow and Debreu, 1954). Hayek's framing is consistent with the definition of *non-ergodic* in Chapter 2: he understands that catallaxies re-pattern themselves in unpredictable ways<sup>18</sup>; however, "large-scale disequilibrium would be impossible in a catallaxy of wholly unhampered markets." (Gray, 1998, p. 92). Hayek writes:

The significance of the abstract character of [spontaneous] orders rests on the fact that they may persist while all the particular elements they comprise, and even the number of such elements, change. (Hayek, 1973, p. 38)

#### 4.4.1 Criticisms of Hayek's 'Market Order'

Having outlined Hayek's framing of the 'market order', how might we evaluate it? There are two broad criticisms relevant to this thesis. The first is posed as a question in Vanberg (1986): "what rules can be considered 'appropriate' in the sense of allowing for a beneficial working of the market mechanism[?]" (p. 97). The view that markets can operate without enabling institutions<sup>19</sup> has been criticised by a number of researchers (e.g., Vanberg, 1986; Hodgson, 1988; and Loasby, 2000): markets are institutions which, historically, have had to be supported by various conventions, norms, and legal rules.

Hayek's answer to this question is wrapped up in his theory of cultural evolution, which is "expected systematically to select for appropriate rules." (Vanberg, 1986, p. 97). We discuss this further in Section 4.5 below.

The second criticism relates to the stability of the order Hayek describes: in Hayek's framing there is no room for what we might refer to as *endogenous volatility*, i.e., unwelcome outcomes brought about by agents making decisions in "free markets".

Moreover, Hayek wrote at length about how he attributed much of the volatility in a catallaxy to attempts at planning, most notably by national governments. For example, in *The Denationalisation of Money*, Hayek writes "monetary policy is much more likely to be a cause than a cure of depressions, because it is much easier [for governments to give in] to the clamour for cheap money" (Hayek, 1976b, p. 79)<sup>20</sup>.

In a sense, therefore, if we view governments as outside of a catallaxy, poor policies represent external shocks to the system.

---

<sup>18</sup>New resources might be discovered and new technologies invented, either of which will cascade across a catallaxy via actions which are, in effect, coordinated by the price mechanism.

<sup>19</sup>Typified by Williamson (1975) who writes that in "the beginning there were markets" (p. 20).

<sup>20</sup>Hayek was more open to the idea that private banking systems could be the cause of monetary expansion in Hayek (1933) but he appeared to step back from this in Hayek (1976b).

If, however, we consider some of the principles of complex systems discussed in Chapter 2, notably positive feedback loops, non-linearity, and systems existing far from equilibrium, it seems unreasonable to simply assume that a catallaxy necessarily gravitates to stability.

This is a general observation from the point of view of the complexity sciences but it is supported by research in the past two decades looking at endogenous volatility in economic and financial behaviour. Examples include [Branch and Evans \(2007\)](#), [Al-Suwailem \(2014\)](#), and [Gomes \(2014\)](#).

Related to this is the work of George Shackle (who was a PhD student of Hayek's)<sup>21</sup>. Shackle was not a particularly clear writer so his work is difficult to understand but in light of Hayek's framework, Shackle emphasised the subjective nature of expectations within a context of uncertainty. More specifically, business confidence is viewed as a "matter of animal spirits or creative imagination rather than of rational assessment." ([Gray, 1998](#), p.92).

For Shackle, large scale dis-equilibria can result from this mass subjective decision making, including a 'national economy' operating far below full employment, i.e., depressions can result from endogenous volatility. In a sense, Shackle's views here are an attempt to bridge Hayek's subjectivism (and micro foundations) and [Keynes' \(1936\)](#) emphasis on aggregate demand volatility (and macro focus).

There are two final points to note here regarding endogenous volatility. First, the notion is a powerful challenge to Hayek's classical liberalism (specifically the argument for a minimalist state). Most importantly, this volatility could be used to legitimise some types of "central planning", which Hayek argued vociferously against throughout his career. A good example is, again, John Maynard Keynes' argument that market systems do not necessarily gravitate to full employment; but the point is a broader one. Any form of endogenous volatility could be used as a rationale for government policy.

Second, and related, it is perhaps surprising that Hayek did not give more credence to the idea of endogenous volatility in his work given his discussions of 'complexity' noted in detail above (and his association with some of the major researchers in General Systems Theory and Cybernetics). Clearly we cannot know why Hayek did not attach more weight to endogenous volatility but it seems reasonable to state that it probably would have required a significant shift away from his Austrian intellectual heritage.

Now that we have considered Hayek's interpretation of spontaneous order brought about by free markets, let us turn to his theory of cultural evolution.

---

<sup>21</sup>[Gray \(1998\)](#) refers to Shackle as "Hayek's brilliant but somewhat neglected pupil." (p. 92).

## 4.5 Hayek's Theory of Cultural Evolution

In this section we focus on Hayek's theory but we first briefly discuss what a 'good' theory of social evolution looks like. This will help us evaluate Hayek's own theory and help us analyse the results of simulations based on the second model below, which appear to have some of these features.

This section comes with a health warning equivalent to those stated previously: the literatures concerned with Hayek's theory, the nature of evolution in biology (including the contentious subject of group selection), and the applicability of evolutionary concepts to social science are each enormous. This section attempts only to distil key issues from these literatures which are important for this thesis.

The subject of "generalized Darwinism" (e.g., [Hodgson and Knudsen, 2006](#)) was discussed in Chapter 2 where the three principles of variation, selection, and durability were generalized in two different ways: the degree of abstraction of these words, and the 'level' of explanation within a stratified ontology<sup>22</sup>. It stands to reason that a good theory of cultural evolution would contain mechanisms that correspond to each of these principles.

Here we add a further point from [Hodgson and Knudsen \(2006\)](#), that "Darwinism alone is not enough" (p. 15) because "explanations additional to natural selection are always required to explain any evolved phenomenon." (ibid). Put another way, the three principles are necessary but not sufficient - we also require a compelling explanation.

How does Hayek's theory of cultural evolution fare in this context? We find that it has been heavily criticised mainly because it lacks coherence vis-à-vis the three principles and, over and above this, because it appears weak in the face of empirical evidence. For example, [Vanberg \(1986\)](#) writes that Hayek's "notion of cultural group selection is theoretically vague, inconsistent with the basic thrust of [his] individualistic approach, and faulty judged on its own grounds." (p. 97). Similarly, [Hodgson \(1991\)](#) writes "Hayek is unclear about the mechanisms of socioeconomic evolution and thus, like [Herbert] Spencer's, his account of evolutionary processes still has to fall back on some strange, detached, and universal selective force, emanating from the 'free' market." (p. 78). [Gray \(1998\)](#) does not mix his words when he states that Hayek's "whiggish interpretation of history has been secularized in a pseudo-Darwinian idiom." (p. 152).

Let us look at Hayek's theory in more detail.

He states that "cultural evolution operates largely through group selection; whether group selection also operates in biological evolution remains an open question - one on which my argument does not depend." ([Hayek, 1988](#), p. 25). However, we will see below

---

<sup>22</sup>See Section 2.3.3.5 (p. 54).

that Hayek's account awkwardly straddles an individual-based theory of variation and selection, and that of group selection.

Moreover, it is worth noting that Hayek's group selection operates "always within a given (market) structure" and "ignores the possibility that selection may also be working at the level of structure and substructure, creating a diversity ... of types of economic system or subsystem." (Hodgson, 1991, p. 79).

We might also add that Ullmann-Margalit (1978) aligns Hayek's theory of cultural evolution with her functional-evolutionary mould of invisible hand explanation:

Notable among the authors to whom this conception [of the functional-evolutionary mould] of invisible-hand explanations may be attributed is F. A. Hayek; he is the most explicit in what he says about this matter. (Ullmann-Margalit, 1978, p. 282)

Below we consider Hayek's theory of cultural evolution, which he expounded in various publications, by looking at the principles of variation, selection, and durability in turn<sup>23</sup>. We will then consider whether it is coherent and compelling as an explanation over and above these three principles, specifically with respect to empirical evidence.

After this, we briefly discuss the normative content of Hayek's theory and how rules might become common to groups.

### 4.5.1 Variation

In Hayek's words:

...the evolution of culture [was] made possible by some individuals breaking some traditional rules and practising new forms of conduct - not because they understood them to be better, but because the groups which acted on them prospered more than others and grew. (Hayek, 1976c, p. 161)

This statement is consistent with Hayek's other writings concerning variation: it is about individuals adopting new rules of behaviour, who by "experimenting with new practices act as innovators and generate 'new variants'" (Vanberg, 1986, p. 82).

Hayek's' individual-based account of variation is compatible with Ullmann-Margalit's (1978) first characteristic of invisible hand explanations. The problem, however, is when we consider a wider set of rules, most notably those formally agreed within (or even imposed upon but followed by) groups.

<sup>23</sup>We see in Hayek (1967) a reference to the three principles of Darwinism as that of "a mechanism of reduplication with transmittable variations and competitive selection." (p. 32).



The core problem here has been argued well by Buchanan and Vanberg. [Buchanan \(1986\)](#), for example, writes that constitutions “*cannot emerge in a process of simultaneous coordination* analogous to that which allows us to classify the market as efficient or which characterizes natural selection.” (p. 321, emphasis included). Equivalently, [Vanberg \(2011\)](#) writes that there are “kinds of rules that only groups can experiment with, be it because by their nature they are inapplicable at the level of individual behaviour - such as, e.g., rules for organizing collective action - or because individuals cannot directly benefit from unilaterally adopting them” (p. 21).

Put another way, some collective rules cannot come about by agents acting in isolation. This criticism is made all the more powerful, [Vanberg \(2011\)](#) argues, when we consider that Hayek's theory of cultural evolution is one of spontaneous order so any form of planning (like individuals consciously coordinating their actions or designing enforcement mechanisms) must be excluded.

## 4.5.2 Selection

The problem with Hayek's selection mechanism is that it is incoherent (it combines individual and collective forms of selection) and it also appears to be missing a mechanism that overcomes the free-rider problem.

Hayek's selection argument at the individual level is oriented around function combined with mimicry. If some new rule benefits a ‘deviant’ more than the previously adopted rule in some social context, that individual will replace the old with the new in their repertoire and we would expect it to be copied (selected) by other individuals. In this framework, rules compete against other rules in the social context, and the best wins out. This idea is included in many evolutionary game theoretic models, which are discussed in more detail in the next chapter.

Curiously, while [Ullmann-Margalit \(1978\)](#) aligns Hayek's theory of cultural selection with her functional-evolutionary mould, there is an argument that Hayek's framing of selection at the individual level resembles Ullmann-Margalit's *amplification* rather than an *invisible hand explanation*. It appears more concerned with individuals copying others.

In addition to his individual-based selection, Hayek also proposes that group selection mechanisms operate at the group level.

It is important to note here that Hayek was fully aware of the free-rider problem. This is crucial in the context of group selection arguments because it raises the question of whether individual agents would sustain rules that are beneficial to the group but where they have an incentive to cheat. [Hayek \(1973\)](#) writes of individuals that some rules “they will follow spontaneously because they will be part of their common cultural tradition.

But there will be still others which they may have to be *made to obey*, since ... it would be in the interest of each to disregard them.” (p. 45, emphasis added).

However, despite this recognition, several researchers have emphasised a bias in Hayek's work toward interactions of a coordination type rather than others (which appear to fall into [Schultz's \(2001\)](#) collective action scenarios and therefore involve the free-rider problem).

[Schultz \(2001\)](#) argues that Hume, Menger, and Hayek all conflate his two categories of interaction, focusing much more on the first, which is related to coordination (pp. 61-62). [Hodgson \(2006a\)](#) makes an equivalent point, noting that the difference is significant because “coordination rules can often emerge spontaneously and be self-reinforcing” (p. 14) whereas other forms of interaction (like the Prisoners' Dilemma) involve defection / the free-rider problem (so are not self-enforcing).

[Vanberg and Buchanan \(1988\)](#) mimic [Schultz \(2001\)](#) when they write that there “is a tendency throughout this tradition - from David Hume over Carl Menger to F. A. Hayek - to argue as if the kind of explanation that applies to coordination-type rules can be generalized to other kinds of rules as well, including those of the prisoner's dilemma type.” (p. 143).

Vanberg notes:

Apparently, and strangely enough, Hayek appeals to ... a collectivist, functionalist notion ([Hayek, 1967](#), pp. 70f., 74) when he stresses that, in cultural evolution, a process of “group selection” is of “greatest importance” ([Hayek, 1979](#), p. 202), a selection process which Hayek obviously considers to be different from the process of variation and selection by individual choices... ([Vanberg, 1986](#), p. 84)

Moreover, [Vanberg \(1986\)](#) writes that “with those behavioural patterns that apparently are advantageous to the group in which they are practised, but appear to be disadvantageous on the part of the individual exhibiting them ... Hayek considers recourse to the notion of group selection to be necessary” (p. 87). In Vanberg's interpretation of Hayek, therefore, group selection is necessary vis-à-vis free-rider problems.

In terms of the mechanisms of group selection, Hayek refers to population growth (expected to be greater in more prosperous societies), migration from group to group, and conflict between groups. Over time, the group with ‘better’ customs / traditions / habits will survive and other groups will (relatively) wither.

The problem, however, is that Hayek did not propose a mechanism by which the free-rider problem is overcome within groups other than vague references to “tradition and

group retaliation" (Gedeon, 2015, p. 16)<sup>24</sup> and to similarly loose references to government.

It is well known in economics (e.g., Bergstrom, Blume and Varian, 1986; Andreoni, 1988; and Okada, 2008) that the enforcement of social norms by individuals can be interpreted as a (second-order) public goods problem because retaliation for not following a social norm has an individual cost but a public benefit. This means it is not sufficient for Hayek to resort to selection mechanisms like "retaliation" without explaining how the subsequent collective action situations are then resolved.

In referring to Hayek's comments about the role of government in group-wide rules, Vanberg proposes an important criticism of Hayek's work:

...these aspects are merely added to, rather than systematically integrated with Hayek's theory of cultural evolution. Their systematic incorporation would require Hayek's theory to be much more specific about the nature, scope, and limits of *evolutionary* principles and their relation to or interaction with the forces of organized, political choice in cultural change. (Vanberg, 1986, p. 96, emphasis included)

Overall, we can say that Hayek's selection mechanism is incoherent because it confuses two levels of explanation (individuals and groups) and he fails to provide a compelling mechanism for overcoming the free-rider problem.

### 4.5.3 Durability

When it comes to Hayek's functional arguments of cultural evolution, rules are selected for and endure for reasons of functionality. Therefore, there is little more to add to Hayek's theory in the context of durability.

We will add, however, that in terms of what Hayek (1967) refers to as "reduplication", his view is that "cultural evolution is brought about through transmission of habits and information not merely from the individual's physical parents, but from an indefinite number of 'ancestors'." (Hayek, 1988, p. 25).

There are two points worth considering further from this: the question of Lamarckism, and the role of habits and cultural traditions.

The quote from Hayek above chimes with the Lamarckian principle of 'acquired characteristics', which raises the question of whether this can be compatible with a Darwinian

---

<sup>24</sup>Hayek (1973) discusses "rules which do not simply follow from [individuals'] desires and their insight into relations of cause and effect, but which are normative and tell them what they ought to or ought not to do." (p. 45).

account. When it comes to cultural evolution, Darwin and Lamarck are not incompatible: [Hodgson \(2002b\)](#) writes that in social evolution, “Lamarckism and Darwinism are neither opposites, nor even mutually exclusive.” (p. 270, Footnote 14). The reason is because in the social realm, agents can learn from others via mimicry while, at the same time, the three principles of generalized Darwinism can remain relevant. [Hodgson \(2001\)](#) and [Knudsen \(2001\)](#) also discuss the compatibility of Lamarckism and Darwinism in social contexts.

Moreover, as [Marciano \(2009\)](#) points out, Darwin discussed cultural inheritance in *The Descent of Man* ([Darwin, \[1871\] 1988](#)). For example, the “greater intellectual vigour and power of invention in man is probably due to natural selection, combined with the inherited effects of habit” ([Darwin, \[1871\] 1988](#), p. 372, as quoted in [Marciano, 2009](#), p. 58).

In the quote above from Hayek regarding cultural evolution, he referred to the role of habits. This is worth discussing in more detail because it highlights a particular form of ‘cognitive embeddedness’ which differentiates Hayek’s theory from the idea of “rational habits” (e.g., [Stigler and Becker, 1977](#), and [Becker, 1992](#)).

[Gray \(1998\)](#) argues that while there is some commonality, Hayek’s reference to habits (and cultural traditions) is different to that Becker’s. The latter’s view is that habits, or rules, form which are due to previously rational behaviour: habits minimize information search costs. By contrast, ‘rationality’ plays a limited role in Hayek’s conception of social norms and conventions, and “Hayek has always maintained that a measure of ‘uncritical’ submission to social convention is an indispensable condition of stability as much as liberty.” ([Gray, 1998](#), p. 50).

We will see in the simulation results reported below that organic institutions emerge via reinforcement learning and that these rules become habits over time.

#### 4.5.4 Hayek’s Explanation of Cultural Evolution

At the beginning of this section we noted [Hodgson and Knudsen’s \(2006\)](#) argument that a theory of social evolution ought to contain convincing references to the three principles of generalised Darwinism and a compelling explanation over and above these three. Here we briefly add another criticism of Hayek’s theory of cultural evolution, linked to “good explanations”, related to empirical evidence.

Is Hayek’s theory consistent with observed cultural phenomena? His theory can be summarised by stating that experiments of rules by individuals (with unplanned outcomes) are adopted by them (and others) if they are beneficial; and that groups that adopt such rules will endure better than those that do not.

A criticism is that Hayek's theory does not explain observed cultural differences and changes over time, at least not in a particularly compelling way. [Gray \(1988\)](#) summarises these perspectives well when, while also criticising Hayek's functionalism, he writes that:

...human history is too riddled with sheer contingencies for any monocausal model of institutional development to be at all plausible, and for this reason, a Darwinian explanation of the rise and fall of institutions comes up against many strong counterexamples. For example, Hayek's suggestion that there is a sort of natural selection of religions, in which religions favouring private property and family life prevail over others by virtue of the enhanced survival chances they afford the offspring of their practitioners, neglects the role that the capture of state power has often played in accounting for the triumph of religions over their rivals. In fact, the evolutionary turn in Hayek's thought seems open to all the criticisms and objections that disable the evolutionary-functional sociologies of Herbert Spencer, W. G. Sumner, and (perhaps) Marx. ([Gray, 1988](#), p. 57)

On the historical contingencies point, [Gray \(1998\)](#) writes that this "is a point Hayek implicitly recognises when he acknowledges that barbarous militarist states may win out over more pacific free societies, but it has large implications which may demand a revision of his system." (p. 138). This is reminiscent of [Vanberg's \(1986\)](#) criticism noted above that Hayek's discussion of government in [Hayek \(1973\)](#) appeared to be *in addition* to his body of thought rather than integrated within it.

The argument here, however, is not that group selection mechanisms linked to conventions and social norms have no explanatory power for the existence and evolution of cultural phenomena at all. It is more that these mechanisms appear weak when contrasted with, for example, the exercise of power by monarchs.

### 4.5.5 Order and Efficiency

At this stage it is helpful to briefly consider how Hayek perceives efficiency vis-à-vis spontaneous order, i.e., the question of whether this type of order is necessarily beneficial for a population. Hayek's definition of order is restated here for convenience. Order is:

a state of affairs in which a multiplicity of elements of various kinds are so related to each other that we may learn from our acquaintance with some spatial or temporal part of the whole to form correct expectations concerning the rest, or at least expectations which have a good chance of proving correct. ([Hayek, 1973](#), p. 36, emphasis removed)

As stated this definition “is entirely value-neutral” (Gray, 1998, p. 119). Consistent with this, Hayek was clear in certain parts of his work that the order which arises in a population is not necessarily beneficial for its constituents (Hayek, 1973, Chapter 2).

The reason why Hayek promoted the notion of spontaneous order in his work despite the potential for ‘bad order’, is due to the selection process within his framing of social evolution: more efficient social structures, including institutions, will be selected for over a sufficiently long time period.

Buchanan has been an ardent critic of Hayek concerning this point. The problem, argues Buchanan, is that:

emphasis is placed on the unintended consequences of limited-vision actions, with an implicit faith that these consequences will be benign. It is as if the many entrepreneurial choices, in the small, act always to push the institutional frontier towards efficiency, in the small *and* in the large.e (Buchanan, 1986, p. 79, emphasis included)

These arguments can be augmented by the concept of lock-in, which was discussed in Chapter 2. Detrimental rules can become locked in as can beneficial rules that become obsolete.

These issues are addressed to some extent by the computational models below. Most notably, one of the ‘null’ experiments run using the second model adopts substantive rationality as the agents’ mental models. In that scenario we observe the population of agents (in effect) descend into a “war of all against all”. Similarly, in certain parts of the parameter space when the agents use different mental models to make decisions, the opposite of property rights emerges: agents learn it is preferable to steal from each other and we observe another Hobbesian-like state (equivalent to the Hawk-Hawk outcome in an evolutionary Hawk-Dove game).

### 4.5.6 Mapping from Individuals to Groups

There is a question of how organic institutions become common to groups of agents. Here we briefly emphasise three mechanisms mentioned at various points in this chapter for how rules common to groups might arise from - or among - their constituent individuals. Conventions and social norms are typically common to groups of agents<sup>25</sup> and a helpful question to ask is: how did they become common?

The first mechanism, directly related to Hayek’s theory of cultural evolution, is the simple process of mimicry. It does not require much imagination to see how a rule which

---

<sup>25</sup>We might also add that ‘groups’ (e.g., a nation) can be defined by phenomena like conventions and social norms.

is beneficial to one agent might then be copied by others, and for it to then become common to a group.

The second process was attributed to Becker above. If a particular rule is optimal for all the agents in a group then we can imagine that it might become adopted by all individually, thereby becoming common.

Indeed, we might think of Hayek's and Becker's mechanisms as 'fractal explanations' of group phenomena: what we see at the social level is a reflection, or an aggregation, of individuals.

Third, we can appreciate that deliberately designed legal rules with credible enforcement mechanisms could also ensure that some rule is adopted by a group of agents. This is not to argue that all such rules are necessary useful, nor that they achieve their objectives without detrimental consequences. The point is that in principle such rules might become common among agents.

None of these three mechanisms meet all three characteristics of [Ullmann-Margalit's \(1978\)](#) invisible hand explanations set out above (individual-based, normalcy, and surprise). Both Hayek's mimicry and Becker's rational rules are individual-based and they both seem reasonably realistic, i.e., they do not rely on the "extraordinary" or "freaky" (which is not to say they are sufficient explanations of institutional emergence); however, neither Hayek's nor Becker's mechanisms could be viewed as surprising. Planned rules seem to fall foul of the first and third characteristics.

In [Chapter 5](#) we describe a fourth mechanism, observed in [Hodgson and Knudsen's \(2004\)](#) results, that explains how rules adopted by individuals can become common within a group: via the co-adaptation of agents' mental models. This is also discussed further in [Chapter 6](#) because it is seen in the simulations results of the models developed for this thesis.

## 4.6 Hayek on Common Law and Legislation

For Hayek, his theory of cultural evolution is relevant not only to organic institutions but to English common law too. Hayek refers to the latter as "a law existing independently of anyone's will and at the same time binding upon and developed by the independent courts; a law with which parliament only rarely interfered" ([Hayek, 1973](#), p. 85).

We can see from this quote how Hayek's view both chimes with Hale's perspective on common law (discussed in the previous chapter) and is consistent with his theory of cultural evolution (the reference to independence from will sits particularly well with the idea of order that is not designed). For Hayek, common law is a type of spontaneous

order, which stands in contrast to any legislation that is designed via the mindset of 'constructivist rationalism'.

The aim of this and the next section, which looks at Vanberg's attempt to reconcile some apparent inconsistencies in Hayek's work regarding legislation, is related to the 'liberal legislation' experiments, conducted with the second computational model, in Chapter 12.

When it comes to common law, Hayek follows in the tradition of Hale but he also follows Menger - to some extent - in the latter's distinction between 'organic' (unplanned) and 'pragmatic' (designed) institutions (Menger, [1883] 1985)<sup>26</sup>.

In Hayek's view, common law is formed bottom-up and 'discovered' by the judiciary rather than being deliberately designed in a top-down manner. Gray (1998) explains this well when he writes that "[j]ust as central allocation of economic resources produces chaotic waste and a degree of coordination of activities far less exact than that yielded by the market process, so centralized legislation cannot match the subtlety of common law in responding to complex and changing circumstances." (p. 69).

We can see in this quote from Gray the correspondence between markets and common law, and between central planning and legislation. This is consistent with one of Buchanan's criticisms of Hayek's body of work, that of an (inappropriate, in Buchanan's view) extension of the market analogy to institutions (e.g., Buchanan, 1986, pp. 75-86).

There are two broad criticisms we can level at Hayek's framing of common law and legislation. First, that it is ahistoric and, second, that Hayek appeared to contradict himself. Let us look at these in turn. The second criticism is used as background to Vanberg's (1994b) distinction between unconditional and conditional evolution in the next section.

### 4.6.1 Hayek's Ahistoric View of Common Law

A number of researchers have argued that Hayek's account of common law does not accord with empirical evidence. Related to (but preceding Hayek), Menger writes that the:

theory of the 'higher wisdom' of common law thus not only contradicts experience but is at the same time rooted in a vague feeling, in a misunderstanding. It is an exaggeration carried to the point of distortion, of the true statement that positive legislation has upon occasion not comprehended the unintended wisdom in common law... (Menger, 1985, p. 233)

---

<sup>26</sup>We will see below that Menger did not fully share Hayek's point of view.



This is consistent with [Hamowy \(2003\)](#) who argues that while Hayek's account has some merit,

it fails as an accurate description of the genesis and development of common law. More importantly, it does not address the common law's weaknesses and inadequacies, which were so extensive that it was only by supplementing it with other systems of substantive and procedural rules, particularly the law of equity, that it was able to survive its early history. ([Hamowy, 2003](#), p. 245)

Put another way, common law had to be enabled by legislation and other planned mechanisms in order to survive, especially the law of equity.

Luban summarises the problem well when he writes:

Hayek's sketch of the English common-law judge essentially restates the stylized depictions offered by jurists like Edward Coke and Matthew Hale in the course of seventeenth-century polemics against absolutism. They aimed to show that the common law had developed organically, independent of the central sovereign, but historians from Maitland onward have shown how misleading such a picture is (cf. [Hamowy 2003](#)). ([Luban, 2020](#), p. 76)

Finally, as mentioned previously, [Gray \(1998\)](#) refers to Hayek's "whiggish fallacy" (p. 152), which seems apt in this context.

In summarising, it is important to highlight a comment from [Hamowy \(2003\)](#), that Hayek's "characterization [of common law] has some merit" (p. 245), i.e., we should not fully reject Hayek's understanding of the "wisdom" inherent in common law. It is more that his account is incomplete, notably with respect to the historical role of legislation in enabling it. Given that "Hayek's distrust of social institutions that are clearly the product of deliberate design [runs] deep" ([Hamowy, 2003](#), p. 262), this is an important criticism.

## 4.6.2 Inconsistency in Hayek's Writings

When we look closer at Hayek's work vis-à-vis common law versus deliberately designed legislation, we find what a number of researchers have referred to as an inconsistency.

On the one hand, Hayek lauds the virtues of common law and the idea that it contains some higher wisdom (and he binds it to the notion of spontaneous order). On the other hand, Hayek discusses the need for legislation to do particular things, including the curtailing of common law. [Gray \(1998\)](#) writes that this is because Hayek believed

the “spontaneous development of law may sometimes result in dead ends or practical deadlocks from which it has to be extricated.” (p. 70). For Hayek’s discussion of this, see [Hayek \(1973\)](#), pp. 88-89. Moreover, Hayek linked this risk to his proposed role of an upper parliamentary chamber (the Legislative Assembly)<sup>27</sup>, which is “to correct the evolution of common law” ([Gray, 1998](#), p. 71).

This acceptance by Hayek of a role for legislation vis-à-vis common law sits somewhat awkwardly with the criticisms of Hayek’s framing being ahistoric in the subsection immediately above. As mentioned there, several researchers criticise Hayek for downplaying the historical role of legislation in enabling common law; but at the same time, Hayek writes that there ought to be a role. The issue appears to be that Hayek’s critics cited above are reacting to certain parts of his work and not others. [Vanberg \(1994b\)](#) argues this is the case: many of Hayek’s critics have placed greater emphasis on his later work, including [Hayek \(1988\)](#), than his earlier work.

Nonetheless, the idea of an upper chamber directing common law seems to represent a significant contradiction in Hayek’s work. This would require the people in the upper chamber understanding the intricacies of the whole corpus of common law, who would then consciously plan changes to the benefit of the population. This closely resembles the centralized planning which Hayek argued vociferously against in much of his work: no human being could hold the breadth of knowledge (much of it tacit) in mind that is required to plan on behalf of society.

There is an argument that Hayek understood that an upper chamber would have to do its best when common law found itself in a dead end: the benefits might be thought to outweigh the costs, unintended consequences aside. Perfect planning is impossible but imperfect planning might be net beneficial.

The problem with this argument, however, is that it admits to planning having some potential value. This then raises the question of the circumstances under which spontaneous order should be allowed to unfold unchallenged, and those in which planning can add value. Hayek did not offer a clear answer to this. However, Vanberg attempted to reconstruct Hayek’s arguments in a way that makes them more coherent. Let us look at Vanberg’s arguments now.

## 4.7 Vanberg’s Reconciliation of Hayek Regarding Legislation

In his article *Hayek’s legacy and the future of liberal thought* ([Vanberg, 1994b](#)), Vanberg considered these apparent inconsistencies in Hayek’s work in some detail, stating

---

<sup>27</sup>Hayek’s policy proposals are elaborated in [Hayek \(1979\)](#), Chapter 17.

that Hayek was far from clear about the relationship between planned legislation and spontaneous order. This problem is related to the fact his “work appears to contain a fundamental tension between what I call ‘rational liberalism’ and ‘evolutionary agnosticism.’” (Vanberg, 1994b, p. 179). Barry (1994) and Kukathas (1989) make essentially the same point.

Vanberg (1994b)<sup>28</sup> proposes a way of reconciling these two seemingly conflicting parts of Hayek's work by distinguishing between *unconditional* and *conditional evolution* (developed in more detail in Vanberg, 1994c), arguing that Hayek's views can be reconciled under the banner of conditional evolution.

It is worth looking in more detail at Vanberg's arguments because they have a direct bearing on the ‘liberal legislation’ experiments conducted using the second model, which is described in Chapter 12 below.

Vanberg (1994b) notes many references in Hayek's earlier work to the constructive use of legislation for the “task of improving our institutions” (Hayek, 1960, p. 5, as quoted in Vanberg, 1994b, p. 179, Footnote 1). This is consistent with researchers such as Gray and Kukathas who argued that the aim of much of Hayek's work was to defend classical liberalism.

Hayek's later work tended to emphasise his views on cultural evolution. This is most obvious in *The Fatal Conceit* (Hayek, 1988), which is described by Vanberg as “the book in which Hayek's rational liberalism is least visible” (Vanberg, 1994b, p. 180). However, Vanberg believes that it “has gained disproportional attention and is widely regarded as the definitive and authoritative summary of his ideas” (ibid) because it is Hayek's last book. We should also bear in mind here the question of the authenticity of Hayek (1988) noted earlier in this chapter, which came to light after Vanberg (1994b) was first published. In any case, Vanberg emphasises that we should not lose sight of Hayek's “rational liberalism” as this was an important part of his work for decades.

Nonetheless, the challenge remains of reconciling what appears to be two seemingly contradictory strands of Hayek's work.

Vanberg developed his argument by first distinguishing between two forms of evolution.

Unconditional claims are statements about evolution per se, statements that leave totally unspecified the kinds of constraints under which the evolutionary process occurs. Such claims provide no substantive information about what it is that can be expected to survive. Conditional claims, by contrast,

---

<sup>28</sup>Vanberg (1994b) has also been published as Vanberg (2001): the latter is an extended version of the former (both are cited below).

are statements about the working properties of evolutionary processes under specified constraints. (Vanberg, 1994b, p. 185)

Hayek's cultural evolution, writes Vanberg, is often viewed through the lens of unconditional evolution, which makes it appear incompatible with his rational liberalism. The latter includes the designing of what Hayek himself referred to as "liberal legislation".

It is worthwhile quoting the following from Vanberg (1994b) in full as a helpful summary of Hayek's liberal perspective and because it contains a quote from Hayek - a gardening analogy - that suggests he is sympathetic to the notion of an "enabling environment":

In a handbook article on Liberalism, written in 1973, Hayek included a section entitled "Positive Tasks of Liberal Legislation" in which he refers approvingly to certain "neoliberal" approaches that explicitly address the issue of what the positive content of the legal framework must be in order "to make the market mechanism operate satisfactorily" (Hayek, 1978, p. 146). Though he did not specify which neoliberal approaches he had in mind, his description certainly fits German Ordo-liberals of the so-called Freiburg School, like Walter Eucken and Franz Boehm. It corresponds to their understanding of the role of liberal legislation when Hayek (1976d [1944], p. 18) notes that the "attitude of the liberal toward society is like that of the gardener" who seeks to create favorable conditions for natural growth. (Vanberg, 1994b, p. 182)

Vanberg (1994b) argues that Hayek's rational liberalism and evolutionary agnosticism can be reconciled if we think of his cultural evolution as conditional, i.e., as constrained by liberal legislation. However, we should note Vanberg's own health warning here: Hayek "made very little effort to explicitly state what the relevant characteristics of a beneficial process of cultural evolution are", so we must "reconstruct his notion of a properly constrained process from his writings." (Vanberg, 1994b, p. 195). Vanberg's argument, therefore, contains some conjecture over parts of Hayek's work.

Vanberg's 1994b reference to conditional evolution fits with the distinction he made in Vanberg (2001) to *action interests* and *constitutional interests*, which he discussed in more detail in Chapter 4 of Vanberg (1994a). *Action interests* are "our preferences over alternative courses of action that are open to us under given constraints", and *constitutional interests* are "our preferences over alternative rule regimes under which we may come to live." (Vanberg, 2001, p. 65).

There is a large literature discussing action and constitutional interests (and the distinction is important in the field of Constitutional Economics). The key idea is that individuals can exercise their 'rationality' in different ways in the two 'levels': people

might collectively agree to formal rules which apply to a group of individuals - including themselves - that they would prefer not to follow as individuals. Clearly this is directly related to Schultz's (2001) collective action scenarios, where agents would prefer to change strategies at the collectively preferred outcome (as in the Prisoners' Dilemma).

Vanberg's argument, therefore, is that groups can agree 'liberal legislation' at the constitutional level as constraints on individuals' actions. Economies based on free market competition can then evolve but in a conditional way, consistent with Hayek's description of the market order. In summary, "[i]f read as an unconditional claim about cultural evolution per se, Hayek's evolutionary argument makes little sense, and it would be inconsistent with the liberal thrust of his work." (Vanberg, 1994b, p. 195).

The reconciliation of Hayek's work proposed by Vanberg is seductive in that it seems to neatly tie together various strands of Hayek's work which appear incompatible. It is possible, of course, that Hayek's views had changed significantly by the time *The Fatal Conceit* was published in 1988, and that this book did in fact reflect his 'final say' on these matters. As mentioned previously, Hayek did a poor job of telegraphing and explaining changes in his opinions so another way of resolving what Vanberg referred to as Hayek's 'rational liberalism' and 'evolutionary agnosticism' is simply that Hayek had changed his mind, mostly dropping the former for the latter.

Another reconciliation of sorts is that *The Fatal Conceit*, where Hayek's cultural evolution was most prominent, reflected more the views of the book's editor, William Bartley, than Hayek's. Unfortunately, while there is documentary evidence to suggest some of the text came from Bartley and his colleague, we cannot know the extent to which Hayek agreed to the final draft of the book (Bartley died in 1990).

The final point to discuss here concerns what appears to be a second tension - or lack of clarity - in Hayek's work. On the one hand, Hayek seems open to liberal legislation but on the other he is opposed to deliberate, central planning for reasons related to knowledge, as discussed above. On this subject, Vanberg (1994b) helpfully distinguishes between "two types of 'design approaches'" (p. 196): liberal legislation (or what he calls *Ordnungspolitik*) and *intervention*. The former is concerned with "a beneficial framework of rules within which market processes can unfold" whereas the latter "seeks to bring about particular outcomes by intervening in the market process".

Moreover, Vanberg notes that while Hayek's comfort with liberal legislation appears more related to market processes than cultural evolution, "[l]ogic requires that the same distinction between types of policies be made with regard to the process of institutional competition that Hayek's theory of cultural evolution is concerned with." (Vanberg, 1994b, p. 196).

Vanberg's distinction between liberal legislation and intervention is very helpful and appears to resolve the second tension in Hayek's work that he discussed in [Vanberg \(1994b\)](#). Moreover, Vanberg's description of liberal legislation cited above is very compatible with the idea of an enabling environment, which is discussed in more detail in [Chapter 6](#).

### 4.7.1 Links to Computational Models

As mentioned at the beginning of [Section 4.6](#), the main aim of this and the previous section was to provide context for the 'liberal legislation' experiments set out in [Chapter 12](#).

In simulations based on the second model, we see that property rights emerge endogenously when we apply the default parameter set (this also occurs in other parts of the parameter space). In some other parts of the parameter space, however, property rights do not emerge. This can be interpreted in the following way: property rights as an organic institution emerge when the environmental conditions are sufficiently enabling of such emergence, and vice versa. Furthermore, if property rights do emerge, this enables the emergence of markets.

In the context of Hayek's liberal legislation, this result presents a question: can legal rules be designed that catalyse the emergence of property rights when the environment is not sufficiently enabling? This is the question addressed in the liberal legislation experiments: a set of parameters is chosen which is known not to give rise to property rights, and we then apply various legal rules to the agents' interactions (in fact, we look at four different scenarios when property rights do not emerge endogenously).

These legal rules take the form of fines for agents who attempt to steal from others (compensation is paid to 'victims' in some experiments). These rules look more like liberal legislation than intervention because they relate more to the environment in which markets operate rather than being a direct intervention in the market process.

We find that in all four cases, legal rules enable the emergence of property rights. In a sense, this liberal legislation changes a non-enabling environment into an enabling one. This is discussed in more detail in [chapters 6](#) and [12](#).

## 4.8 Summary

The main aim of this chapter, stated at the outset, was to describe and evaluate Hayek's approach to spontaneous order, especially his theory of cultural evolution.

Hayek's work has taken centre stage because, as stated earlier, he developed the most sophisticated and systematic body of thought concerning spontaneous order. We looked

in some detail at his understanding of 'the market order' (including the notion of catalaxy) and then examined more closely his theory of cultural evolution as this is directly related to organic institutions.

We have seen that while Hayek's ontological foundations bear a great deal of resemblance to CE (and his work preceded that subject by decades), his theory of cultural evolution has not stood up to critical evaluation. However, the weaknesses in Hayek's arguments do not necessarily mean we should give up on the idea that organic institutions can emerge in a spontaneous way, nor that 'liberal legislation' could be used to develop an enabling environment along the lines of Vanberg's constrained evolution. These are open questions toward which the computational models below are orientated.

Moreover, while focused on Hayek's work, this chapter deliberately examined [Ullmann-Margalit \(1978\)](#) and [Vanberg \(1994b\)](#) in sections 4.3 and 4.7 respectively. These papers provide motivation and helpful ways of framing the research questions addressed in this thesis and also the computational models developed below.





# Chapter 5

## Models of Organic Institutional Emergence

Q: Why did the chicken cross the road? A: To maximize its utility.

– Geoffrey Hodgson

In this chapter we switch our attention away from theorizing about spontaneous order to focus on models. More specifically, we are interested in those concerned with the emergence of organic institutions.

Four models / areas of research were chosen as the most relevant to this thesis:

1. [Hodgson and Knudsen's \(2004\)](#) model of a simple traffic convention (Section [5.1](#));
2. parts of the social simulation literature, especially the models completed within the “EMergence In the Loop” (EMIL) Project (Section [5.2](#));
3. attempts in game theory to model institutional emergence (Section [5.3](#)); and
4. a strand of literature concerned with money emergence, which follows [Jones \(1976\)](#) and [Kiyotaki and Wright \(1989\)](#) (Section [5.4](#)).

Before moving forward, we should briefly note two strands of literature which are related and interesting but, ultimately, not discussed in detail here.

### Research Not Covered

One strand of related research makes use of systems dynamics, which has been the focus of Michael Radzicki in particular (e.g., [Radzicki, 1988, 2003, 2009, and 2010](#)). Radzicki

has done a considerable amount of work combining institutional economics and the idea of a ‘system dynamics pattern model’, which is a “system of highly non-linear differential equations that has no exact, analytical solution.” (Radzicki, 1988, p. 653).

Radzicki’s work sits comfortably alongside the ontology of complex systems discussed in Chapter 2 but this material is not evaluated in this chapter for two reasons. First, it is not concerned with the *emergence* of new organic institutions, which is our focus of research.

Second, one of the 11 principles of Complexity Economics (CE) mentioned in Section 2.3.3 is that agents are computational in nature. By contrast, dynamical systems of the type explored by Radzicki are mathematical and based on quantified stocks and flows. This difference is important for this thesis because we are interested in how agents use (computation-based) mental models, including whether institutions might emerge from within these mental models.

Another strand of research worth mentioning is that which follows Vernon L. Smith’s paper *An experimental study of competitive market behavior* (Smith, 1962). This work has become associated with the ‘Hayek Hypothesis’, that strict “privacy together with the trading rules of a market institution are sufficient to produce competitive market outcomes at or near 100% efficiency” (Smith, 1982, p. 223). Crockett (2013) surveys the dynamics of various exchange experiments.

This literature is closely related to Hayek’s conceptualisation of the spontaneous order of markets but in this thesis we are interested in a different type of spontaneous order: the emergence of organic institutions. We will not, therefore, consider this research any further.

Let us now look closer at Hodgson and Knudsen’s (2004) model of a simple traffic convention.

## 5.1 Hodgson and Knudsen’s Traffic Convention Model

Hodgson and Knudsen (2004) is concerned with a number of issues of interest to us, notably habituation. As mentioned in the Introduction, the models developed for this thesis can be viewed both as a continuation of and a response to certain issues raised by Hodgson and Knudsen’s (2004) results.

An additional motivation for examining this paper is that Hodgson and Knudsen (2004) appears to be the only research published which, simultaneously: (i) is concerned with organic institutional emergence; (ii) makes use of an Agent-Based Model (ABM); and

(iii) is written specifically for an Institutional Economics (IE) audience. As was also mentioned in the Introduction, it is noteworthy that in three recently published papers petitioning for the use of ABMs in IE (Gräbner and Kapeller, 2015; and Gräbner (2016, 2018), Claudius Gräbner mentions only Hodgson and Knudsen (2004) as an example of such work<sup>1</sup>.

This sub-section is divided into: (i) a description of the model; (ii) a discussion of the results of the simulations; (iii) one criticism of the paper; and (iv) a look at the paper's explanation for how conventions become common to a group of agents.

### 5.1.1 The Traffic Conventions Model

Hodgson and Knudsen (2004) explore the emergence of a traffic convention: whether to drive on the left or right side of the road. In the context of Schultz's (2001) typology, the model fits into his coordination situations: there are no free-rider problems seen in the interactions.

The environment includes agents who are placed on a  $100 \times 2$  ring that represents a road on which they drive around, either on the right or left hand side<sup>2</sup>. There are 40 agents, 20 of whom drive clockwise around the ring and 20 anti-clockwise. The model is turn-based and the agents make their decisions sequentially: whether to drive on the right or left.

In the standard model the agents can see 10 'zones' ahead: they count the number of cars in both lanes and in which direction they are moving. This information is processed by each agent's mental model when it is their turn to move. The agents also count the number of cars in the zone immediately ahead of them (whether travelling in the same or opposite direction): this gives the drivers the opportunity to avoid these cars.

If the cars collide (when two cars occupy the same zone on the ring, irrespective of the direction of travel) then both agents die. Both are replaced by cars driving in the same direction as the replaced cars.

Hodgson and Knudsen's agents have mental models based on four innate dispositions, three of which are linked to the data corresponding to the position of cars ahead of them. The fourth ( $Habitgene_n$ ) corresponds to a variable ( $Habituation_{n,t}$ ) which reflects an agent's historical decisions to drive on the right or left (linked to the concept of habit). Here,  $n$  refers to a single vehicle; and  $t$  refers to the timestamp.

The five pieces of data used in the agents' mental models to make a decision are:

---

<sup>1</sup>Gräbner and Kapeller (2015) discuss other ABMs in IE (p. 439) but none of these are concerned with the emergence of organic institutions.

<sup>2</sup>Here we use 'cars', 'drivers', and 'agents' interchangeably.

- $S_{L,n}$  is the proportion of cars who are driving in the same direction up to 10 zones ahead who drive on the left, where subscript  $L$  refers to driving on the left ( $S_{L,n} = 0.5$  if there are no such cars);
- $O_{L,n}$  is the proportion of cars who are driving in the opposite direction up to 10 zones ahead who drive on their left ( $n$ 's right) ( $O_{L,n} = 0.5$  if there were no such cars);
- $C_{L,n}$  is the number of cars one zone ahead (moving in either direction) driving on  $n$ 's left;
- $C_{R,n}$  is the number of cars one zone ahead (moving in either direction) driving on  $n$ 's right; and
- Habituation $_{n,t}$ , a variable which measures the historical tendency of agents to drive on the left or right.

The last piece of information is updated in the following way, for each player after each move:

$$\text{Habituation}_{n,t} = \text{Habituation}_{n,t-1} + LR_{n,t}/(K + \text{Moves}_{n,t})$$

$LR_{n,t}$  is a value corresponding to  $n$ 's decision to drive on the left or right at time  $t$ .  $LR_{n,t} = 1$  if  $n$  chooses to drive on the left at time  $t$  and  $LR_{n,t} = -1$  if it chooses to drive on the right.  $K$  is an arbitrary parameter which controls the weight of each decision to drive on the right or left hand side. The variable 'Moves' is the number of moves agent  $n$  has made up until time  $t$ , i.e., its 'age' in the simulation. Habituation $_n = 0$  when the agent is instantiated.

Hodgson and Knudsen state that Habituation $_{n,t}$  is bounded between -1 and +1. Presumably this is 'forced' by the code as the above algorithm that updates Habituation $_{n,t}$  does not on its own constrain this variable within these bounds.

The dispositions agents are born with are as follows:

- SSensitivity $_n$ , corresponding to  $S_{L,n}$ ;
- OSensitivity $_n$ , corresponding to  $O_{L,n}$ ;
- Avoidance $_n$ , corresponding to  $C_{L,n} - C_{R,n}$ ; and
- Habitgene $_n$ , corresponding to Habituation $_{n,t}$ .

All of these dispositions are drawn from a normal distribution with a mean of 1 and standard deviation of  $\delta$ .

In addition to these, four (non-negative) ‘weight’ parameters (common to all drivers) are used to control the weight of inputs in to the agents’ decisions to drive on the right or left hand side (corresponding to the four dispositions above, respectively):

- $w_{Sdirection}$ ;
- $w_{Odirection}$ ;
- $w_{Avoidance}$ ; and
- $w_{Habit}$ .

At time  $t$ , agent  $n$  first uses the following expression to make a (pre-error) evaluation of whether to drive on the right or left:

$$\begin{aligned} LREvaluation_n = & w_{Sdirection} \times SSensitivity_n \times (2S_{L,n} - 1) \\ & + w_{Odirection} \times OSensitivity_n \times (2O_{L,n} - 1) \\ & + w_{Avoidance} \times Avoidance_n \times (C_{R,n} - C_{L,n}) \\ & + w_{Habit} \times Habitgene_n \times Habituation_{n,t} \end{aligned}$$

The model is designed so that positive values of  $S_{L,n}$ ,  $O_{L,n}$ ,  $C_{R,n}$ , and  $Habituation_{n,t}$  are associated with  $n$  choosing to drive on the left (and vice versa) and a positive value of  $C_{L,n}$  tends to make  $n$  drive on the right (and vice versa). If unadjusted by an error, an overall value of  $LREvaluation > 0$  will lead  $n$  to drive on the left, and vice versa.

After each subjective evaluation ( $LREvaluation$ ), the final step is to use an error parameter ( $\epsilon$ ) to (potentially) invert the agent’s subjective decision. We can assume that a decimal number is drawn from a uniform distribution with a minimum of zero and a maximum of 1<sup>3</sup> and this random number is compared with  $\epsilon$  (adjusted between 0 and 0.02 in the standard model). If this random number is lower then the right-left evaluation is reversed.

The agent then moves one zone ahead and implements its decision to drive on the left or right. If there is no collision, the next driver takes its turn until all 40 drivers have moved, after which the clock ticks forward (there are 20,000 rounds in each simulation). As mentioned above, if two cars collide, they are both removed and replaced with two new cars (which have new, randomly determined ‘genes’).

<sup>3</sup>The authors did not make it clear how the random number generator works but this is probably a safe assumption.

### 5.1.2 Simulation Results

The aim of the model is to explore “different points of parameter space and [to assess] the impact of different levels of habit and error.” (Hodgson and Knudsen, 2004, p. 27).

In terms of the convergence of the agents on a particular convention, there was considerable variation within the parameter space. In some parts of the space there was hardly any convergence, and in others it was close to total.

Fig. 5.1 below, reproduced from the paper (with permission), is helpful. It shows a measure of convergence in an area of the parameter space where  $w_{Sdirection} = 1.4$ ,  $w_{Odirection} = 0.9$ , and  $w_{Avoidance} = 0.7$  (these values were chosen because they maximized convergence when  $w_{Habit} = 0$ ). The data reported show the area in the parameter space where  $\epsilon$  varies between 0 and 0.02, and  $w_{Habit}$  varies between 0 and 2.




Figure 5.1: Taken from Hodgson and Knudsen (2004), p. 29. The caption in the original paper reads “Fig. 2. Degrees of convergence with 200 runs for each level of habit and error  $w_X = \{1.4, 0.9, 0.7, w_{Habit}\}$ .” The vertical axis shows Convergence (0.5 means the cars drive equally on the left and right side and 1 means all the cars drive on the same side); the y-axis (error) corresponds to  $\epsilon$  above; and the x-axis corresponds to Habituation $_{n,t}$  above. Reproduced with permission from Elsevier B. V.

We can see from this graph that increasing the habit weight ( $w_{Habit}$ ) from 0 to 1 increases the measure of convergence by approximately 0.3 (averaged across all values of error shown). The authors refer to this measure as the *habit effect*.

Fig. 5.2 below shows this habit effect across others parts of the parameter space. The authors normalize  $w_{Sdirection}$ ,  $w_{Odirection}$  and  $w_{Avoidance}$  so these three weights average 1, allowing this 2-D representation.




Figure 5.2: Taken from [Hodgson and Knudsen \(2004\)](#), p. 31. The caption in the original paper reads “Fig. 3. The habit effect in parameter space.”  $C_0$  is convergence when  $w_{Habit} = 0$  and  $C_1$  is convergence when  $w_{Habit} = 1$ . The difference between these two ( $C_1 - C_0$ ) is termed the *habit effect*. Reproduced with permission from Elsevier B. V.

In light of these results, the main conclusion the authors came to was that “convergence is never achieved by the force of habit alone. Furthermore, convergence can sometimes occur with low or zero level of habit. Crucially, habit helps convergence only when it is combined with selection pressure on the fixed ‘instincts’ in the population of cars.” ([Hodgson and Knudsen, 2004](#), p. 30).

The statement here, that “convergence can sometimes occur with low or zero level of habit” looks curious in light of Fig. 5.1 above. In that diagram, when ‘Habit’ is zero, the degree of conversion is approximately 0.57 - 0.64 (depending on  $\epsilon$ ). This hardly looks like convergence at a zero level of habit. What is probably happening here is that the conversion data shown in Fig. 5.1 are mean values over 200 runs: there was probably some variation between runs such that convergence was seen in some of them.

There are two final points worth noting before we discuss and critique the model and results. First, in some simulations, a group of agents learned to be agile in their driving rather than remaining on the same side of the road. [Hodgson and Knudsen \(2004\)](#) write

that a “‘cycling’ pattern can occur, when cohorts of agile drivers repeatedly move safely and laterally to avoid other oncoming groups.” (p. 28). This result of agility emerging instead of a convention is fascinating from a complexity science point of view.

Second, the authors look at whether “inertia” is any different to habituation. This means, essentially, replacing  $Habituation_{n,t}$  with serially correlated behaviour in the form of adding its current (pre- $LR_{n,t}$ ) position and its previous two positions ( $LR_{n,t-1}$  and  $LR_{n,t-2}$ ). Clearly, these terms on average will be correlated with  $Habituation_{n,t}$  but the difference lies in the nature of memory and ‘stickiness’: “Habits are like a crude summarized memory [which] are built up steadily once a repeated behaviour emerges.” (Hodgson and Knudsen, 2004, p. 35).

The results show that inertia, as defined, has a much weaker effect than habituation.

Now that we have understood the results, let us discuss and critique the model.

### 5.1.3 Discussion and Critique

We can see from the model described above that each agent’s decision making uses four agent-specific parameters<sup>4</sup>. These “are akin to instincts: they are fixed for the lifetime of each car.” (Hodgson and Knudsen, 2004, p. 35).

In this context, Hodgson and Knudsen make an important distinction between two broad approaches to habit in the literature, stating that the question “is whether rational choice is the foundation of habit, or whether the reverse is true.” (ibid, p. 22). This is reminiscent of Gray’s (1998) reference to Hayek’s theory of mind in which rational processing supervenes on rules within the sub-conscious.

As mentioned in the previous chapter, some researchers (e.g., Stigler and Becker, 1977; and Becker, 1992) have argued that current habits reflect previously rational decisions, i.e., that rational choice is the foundation of habit. Stigler and Becker (1977) write that “habit is often a more efficient way to deal with moderate or temporary changes in the environment than would be a full, apparently utility-maximizing decision.” (p. 82). Here, “economizing” refers to the costs of acquiring and analysing information.

Hodgson and Knudsen (2004) argue in favour of “defining habit as a disposition or propensity” (p. 36). In doing so, they place themselves “in the pragmatist tradition of Charles Sanders Peirce, William James, George Herbert Mead and John Dewey [for whom] any rational deliberation is always seen as grounded on habit.” (p. 22).

<sup>4</sup>Note that Hodgson and Knudsen refer to  $Habituation_{n,t}$  as a parameter (p. 35) but it changed throughout each simulation for each agent so convention (!) dictates it should be classified as a variable.



Let us step back and consider this in a bit more detail because it is at the heart of human mental models and how these should be represented in computational terms.

Let us assume that people's mental models include: (i) a perceived ontology of the environment in which they exist (including other individuals); (ii) the aim(s) of the agents; (iii) forms of information processing; and (iv) learning processes. This is not meant to be an exhaustive list of the constituents of human cognition, rather it is meant to help us understand Hodgson and Knudsen's model.

We can think of [Stigler and Becker's \(1977\)](#) rationality approach in this framework too. Their ontology is a Neoclassical one (with a specific focus on addiction) in which agents choose between consuming goods / services, and the aim of all the agents is to maximize utility. These agents process information via substantive rationality, i.e., they deduce from the information they have (including given preferences) their preferred choice. Learning does not seem to happen in [Stigler and Becker's \(1977\)](#) model in the way it does in [Hodgson and Knudsen's](#), where we see a change in a variable. However, we might interpret the formation of a habit, post-rational decision, as learning of a type.

For [Stigler and Becker \(1977\)](#), therefore, habits in a sense substitute for substantive rationality as a form of information processing.

[Hodgson and Knudsen's \(2004\)](#) framework includes agents interacting on a road (and all of the cognitive features that entails<sup>5</sup>). The aim of all the agents is to remain alive by not colliding with other cars. The nature of the information processing is very different to [Stigler and Becker's](#): agents use instincts which are fixed at instantiation to make decisions and, as stated above, habit forms part of this process. Importantly, habit work 'alongside' rather than instead of other processed information ( $Habituation_{n,t}$  is a component of  $LREvaluation_n$  above).

The question of learning is taken up below.

### Downward Causation / Effects

In their discussion of results, [Hodgson and Knudsen \(2004\)](#) refer to two types of downward causation<sup>6</sup>. These play an important role in this thesis so let us consider them further.

Following [Campbell \(1974\)](#), Hodgson and Knudsen discuss how a weak form of this phenomenon means that as a convention begins to emerge, non-conforming agents tend to be removed, i.e., "those that survive tend to be those that conform" ([Hodgson and Knudsen, 2004](#), pp. 38-39).

---

<sup>5</sup>For example, knowing what a car is, how to drive, the nature of a road, the difference between left and right, etc.

<sup>6</sup>As mentioned in the Introduction, Hodgson later preferred downward and upward *effects*. We keep to the terminology used in [Hodgson and Knudsen \(2004\)](#) here.

This is of course a crude process because a non-conforming agent might also take out a conforming agent; however, we can appreciate that conforming agents “tend” to survive.

A form of downward causation occurs at the agent level also, via  $\text{Habituation}_{n,t}$ . Following Sperry (1969), Hodgson and Knudsen (2004) refer to this as a strong form of downward causation that ‘reconstitutes’ the agents’ mental models. They write that “as the left/right convention begins to emerge, more and more surviving cars develop the habit to drive on the left or the right” (p. 39).

### Habits Versus Reinforcement Learning

We now look at one criticism of Hodgson and Knudsen’s (2004) model and their interpretation of the results. This is about whether we should interpret the changes to their  $\text{Habituation}_{n,t}$  variable as habituation or reinforcement learning. The authors are aware of and refer to this question in their paper (see Footnote 19, p. 38, in particular).

Hodgson and Knudsen mention Erev and Roth’s (1998) model of reinforcement learning, stating that “our concept of habituation is close to that of reinforcement in their generalized model.” (Hodgson and Knudsen, 2004, p. 38, Footnote 19). For our purposes, what is important in this quote is “close to”, which means that while Hodgson and Knudsen view changes to  $\text{Habituation}_{n,t}$  as similar to reinforcement learning, it is not the same thing.

We can use Erev and Roth’s (1998) three principles of reinforcement learning (see Section 1.4.6) to consider this more closely. First, the *Law of Effect* seems to hold when a driver makes a decision and survives: the change in  $\text{Habituation}_{n,t}$  means it will be more likely to make this decision again in the future, *ceteris paribus*. Note, however, that this is asymmetric: drivers cannot learn from ‘bad’ decisions because they are removed from the simulation.

The *power law of practice* seems to hold in that changes to  $\text{Habituation}_{n,t}$  decelerate over time (and are kept within bounds of  $-1$  and  $+1$ ).

Finally, the third principle, that *choice behaviour is probabilistic* appears not to be relevant here with the very weak exception of errors being determined probabilistically.

Given these three principles, there is an argument that  $\text{Habituation}_{n,t}$  in Hodgson and Knudsen (2004) can be interpreted as a proxy for reinforcement learning.

The difference between reinforcement learning and habituation is clearly important: they are significantly different cognitive processes. The former is sensitive to success and failure relative to some goal(s) whereas the latter is not.

We should be clear, however, that this second criticism does not mean Hodgson and Knudsen are blatantly ‘wrong’: their results are consistent with habituation playing a

role in the emergence of conventions. The problem is that their results are also consistent with an alternative hypothesis that is based on reinforcement learning.

This criticism is one of the motivators for the models developed for this thesis: do organic institutions emerge as a result of renouncement learning or habituation, or both? To answer this, we explore organic institutional emergence with and without reinforcement learning, and with and without habituation.

#### 5.1.4 Commonality of Rules within the Group

In the previous chapter we identified three mechanisms by which rules might become common within a group of individuals (mimicry, identical rational decisions, and a legal rule). None of these appear relevant to [Hodgson and Knudsen's \(2004\)](#) results but conventions common to all the agents were observed in certain parts of the parameter space: how so?

A fourth process appears to be at work here: when observed in the simulations, conventions seem to emerge via a specific form of co-adaptation of the agents' mental models, i.e., the agents iteratively respond to each others' actions in such a way that rule commonality is achieved (the convention). Symmetry breaking appears to be the specific mechanism at play here.

This is described well by Hodgson and Knudsen although they use different language. As a convention begins to emerge, more of the surviving cars start to follow the convention ( $\text{Habituation}_{n,t}$  tends to + or -1 for these cars); and on average more non-conforming cars are removed than those that conform. The phenomena Hodgson and Knudsen point to are reconstitutive downward causation and upward causation.

Now that we have discussed [Hodgson and Knudsen \(2004\)](#) in some detail, let us turn to the literature concerned with computational models of social norms.

## 5.2 Social Norms Within the Social Simulation Literature

There is a large literature concerned with the computational modelling of social norms in what we are referring to here as the social simulation literature, and it spans across different fields of study (including sociology and social psychology).

To help present this work, in this section we focus on research that is not based on game theory, leaving a discussion of game theoretic work to the next section.

It is important to note two points before we proceed: first, these literatures vary considerably in their assumptions, framings, terminology, and aims; and, second, our focus here is on economics.

Perhaps the most notable difference between a lot of the social simulation literature (which appears most heavily influenced by sociology) and economics is the emphasis placed on the free-rider problem. This thesis uses [Schultz's \(2001\)](#) framing of interactions in which his second category (collective action scenarios) includes the free-rider problem (where at least one agent is incentivised to change its behaviour at some socially preferred outcome). In discussing organic institutions in economics, this problem is fundamentally important. By contrast, and by way of example, the EMIL Project Report (discussed in more detail below) was a 245-page paper that summarised a 3-year multi-million Euro project focused on the computational modelling of social norms: the free-rider problem was not mentioned once.

This is not meant as a blanket criticism of sociological approaches. In fact, writing in broad brush strokes, researchers publishing in that literature seem to have been more open-minded about the nature of mental models and socially emergent properties than orthodox economics for decades. Also, and related, these researchers appear to have been much more open to computational modelling than IE has been. The issue being emphasised here is the conceptual distance between computational models based on sociology and those located in economics.

The implication of this point is that we have to be careful translating between fields. Indeed, the first issue we face is terminological. Most notably, the definitions of social norms used in much of this literature appears to be different to that used in economics and also in this thesis.

The interpretation of conventions seems broadly aligned, however. For example, the EMIL Project Report states that a convention is based on “the agent’s goal of conforming to that behavior in order to act like the others, the mutual expectation that the others will conform to that behavior as well.” ([Andrighetto and Conte, 2010](#), p. 11).

By contrast, in the EMIL Project Report a norm is “a prescribed guide for conduct which is generally complied with by the members of society.” ([Ullmann-Margalit, 1977](#), as quoted in [Andrighetto and Conte, 2010](#), p. 10). Moreover, given the definitions of conventions and social norms in the EMIL Report, the distinction “is not clear-cut.” ([Andrighetto and Conte, 2010](#), p. 11).

In this thesis we align social norms specifically with [Schultz's \(2001\)](#) ‘collective action’ category of interactions, which include the free-rider problem. This is consistent with Schultz’s use of the term.

Let us now look at this literature in more detail. As [Elsenbrioch and Gilbert \(2014\)](#) note, this literature can be roughly divided into three. These are models of:

1. the *diffusion* of pre-existing norms, some of which has made use of complexity science framings;
2. the *use/application* of norms, notably within the field of distributed artificial intelligence (AI); and
3. the *emergence* of social norms (which have mostly been game theoretic in nature).

In addition to these three broad categories, we should note that this literature has also considered issues of *internalisation* (as self-policing), e.g., [Villatoro et al \(2015\)](#); the relationship of social norms to *obligations*, e.g., [Savarimuthu et al \(2010\)](#); *reciprocity* and *sanctions*, e.g., [Younger \(2005\)](#); *social responsibility* and *meta-agency*, e.g., [Conte and Paolucci \(2004\)](#); *altruism*, e.g., [Jaffe \(2002\)](#); and *inequality*, e.g., [Saam and Harrer \(1999\)](#). This research is noted here to emphasise the amount of work that has focused on social norms through computational models. The work referenced above is, however, less directly relevant to this thesis so it is mentioned in passing.

Returning to the above three categories, the third is clearly more relevant for this thesis; however, the first two are worth noting briefly, partly as background to the third.

### 5.2.1 Norm Diffusion

Concerning norm diffusion, one such model is based on the Sugarscape model, developed in [Epstein and Axtell \(1996\)](#). In this experiment, agents might adopt a pre-existing norms when making a decision, and the adoption of a norm depends on those used by an agent's neighbours. This model is useful in that it identifies how norms might spread within a population and it included "habitual behaviour when the agents are in a stable norm-environment." ([Elsenbrioch and Gilbert, 2014](#), p. 77).

There is also a large amount of work regarding norm diffusion that has used the Belief-Desire-Intention (BDI) framework and its variant, the Belief-Obligations-Intentions-Desires (BOID) architecture. A detailed discussion of these frameworks can be seen in [Neumann \(2010\)](#): they provide a cognitive architecture for agents that, as their names suggest, incorporate beliefs, intentions, and desires; and obligations in the case of BOID.

[Neumann \(2008\)](#) provides a helpful review of some of the literature focused on norm diffusion, citing [Axelrod \(1986\)](#) as seminal and representative of the game theory approach. He also notes that "[t]he classical model in the tradition of models employing cognitive rich agents is the model described in [[Conte and Castelfranchi \(1995b\)](#)]." ([Neumann,](#)

2008, para. 5.2<sup>7</sup>). The latter model has been replicated several times “and is still the reference point for authors in this tradition.” (ibid, para 5.2). However, its orientation is around the diffusion of pre-existing norms so we will not discuss it any further.

The papers reviewed in Neumann (2008) are considered in the context of whether they add value in any of three ways: (i) the transmission of norms between agents; (ii) the transformation of the agents (i.e., their mental models); and (iii) “the function of norms for a society” (para 2.5). It is noteworthy that Neumann did not consider whether the papers and models he reviewed add value via the origination of new norms.

## 5.2.2 Multi-Agent Systems & Artificial Intelligence Models

Models in this literature tend to take one of two forms: (i) offline design, in which “systems specify what norms a system [of agents] will follow and encode them directly into the agents” (Hollander and Wu, 2011, para. 3.37); or (ii) autonomous innovation, which “requires the agents of a system to create new norms without external influence” (ibid).

Clearly, the second type is of more interest for this thesis. However, Hollander and Wu (2011) note that this tended to be tackled in two ways. The first was through the use of game theory; and the second was through machine learning, which takes “the same selection approach as game theory.” (para. 3.38). It seems, therefore, that when it comes to the origination of norms, this literature loops us back to game theoretic approaches, which is the subject of the next section.

Let us now turn to non-game-theoretic models that are focused on the emergence of social norms.

## 5.2.3 The Emergence of Social Norms

Interestingly, Elsenbrioch and Gilbert (2014) associate models of norm emergence with game theoretic models. They write “[o]ne of the most famous simulations of norm emergence is presented in Axelrod (1986) ... [which features] agents playing iterated prisoners’ dilemmas” (p. 76). We defer our discussion of game theoretic models to the next section.

Here we focus mainly on the computational models reported in the EMIL Project: this was a sophisticated project (both in theoretical and computational modelling terms) which made some important advances in the framing and modelling of social norms.

---

<sup>7</sup>Note that papers in the Journal of Artificial Societies and Social Simulations (JASSS) are published online and therefore do not have page numbers. References therefore use paragraph numbers.

Before we do this, however, let us briefly look at two simple but elegant models of norm emergence contained in [Caldas and Coelho \(1999\)](#).

### 5.2.3.1 Two Simple Coordination Models

[Caldas and Coelho \(1999\)](#) first develop a simple model where 100 agents have to choose from one of 16 colours. The ‘return’ of the chosen colour simply reflects the number of other agents that had chosen the same colour, i.e., the more agents with your chosen colour, the better. The total number of agents choosing each colour is announced at the end of each round and agents use this information in their decisions in the next round. The paper does not explain this in detail but presumably the weight attached to each rule is associated with its frequency in the previous round, i.e., there is some sort of memory.

In this model agent rules consist of the 16 colour types. Each simulation shows that the agents very quickly converge on a single rule-colour and each agent’s return is maximized as a result. This is clearly a pure coordination-type interaction.

We might add here, briefly, that this convergence on a single colour is reminiscent of behaviour seen in ants when a single ant, who has discovered a food source, ‘recruits’ other ants. This can occur through a variety of mechanisms, including laying down a pheromone trail to the food or physically stimulating others at the nest as a signal to follow. See [Kirman \(1993\)](#)<sup>8</sup> for a discussion of this. Such behaviour can lead to symmetry breaking whereby ants who had been wandering randomly are all recruited to visit the food source.

In their second model, [Caldas and Coelho \(1999\)](#) attach different values to each colour. When the authors ran the simulations they noticed that agents converge on the highest value colour only half of the time. In the other half agents become locked in to a lower-value colour: this happens if the colour has a relatively high frequency in the population at instantiation such that the value of replicating this colour exceeds the colour’s intrinsic value. This is related to the concept of lock-in, which was discussed in Chapter 2.

[Caldas and Coelho’s](#) second model is equivalent to the Market Emergence Model developed for this thesis (Chapter 7): both can be interpreted as non-pure coordination games.

Let us now consider the EMIL Project and the four computational models developed within it.

---

<sup>8</sup>This paper is a rare example of a CE approach (with computational modelling) used to examine institution-like phenomena.

### 5.2.3.2 EMergence In the Loop (EMIL)

Elsenbrioch and Gilbert (2014) write that the EMIL Project (which ran between 2006 and 2010) helped move the study of norms beyond the relatively restrictive BDI and BOID architectures. Furthermore, “EMIL achieved its goal of designing a normative agent that does not rely on a static set of norms but can learn normative behaviour from its environment.” (p. 142). Given this outcome, it is worth considering the project in more detail.

In the EMIL models, changes to an agent’s mental model were based on reinforcement learning (from the agent’s own experiences and observations of others) and also ‘norm invocation messages’ received from other agents. The latter can be admonishments for transgressions or assertions about normative behaviour.

The use of reinforcement learning means that these EMIL models form a neat complementarity with Hodgson and Knudsen’s (2004) traffic convention model, which is focused on habituation.

This project had two broad parts: a theoretical component and four computational models, which “were initially carried out for exploratory purposes, but many of them were later replicated for validation.” (Conte and Edmunds, 2010, p. 5). Put another way, the models produced in the project helped develop the theory and to subsequently demonstrate its validity.

An important theoretical feature of this project was the inclusion of Conte and Castelfranchi’s (1995a) *cognitive emergence* and *immergence*. Both of these concepts were defined in the Introduction.

Moreover, immergence and some sort of emergence outside of the agents’ boundaries are thought to work together:

...our simulations show that under given social conditions, namely in multi-scenario worlds, norms operate in society while operating in the mind, neither after nor before. In those circumstances, the two components of norm dynamics – emergence and immergence – need to be strictly intertwined, such that one cannot occur without the other. (Conte and Edmunds, 2010, p. 6-7)

This simultaneous change in immanent and external phenomena will be discussed in more detail in the next chapter and in the Conclusion. For now, it is worth noting how the above quote chimes with Hodgson’s reference in Hodgson (2006a) to institutions, which are “like Klein bottles: the subjective ‘inside’ is simultaneously the objective ‘outside’.” (p. 8).



In critiquing the EMIL computational models, we look at: (i) a summary of the project's outcomes; (ii) an evaluation of the four chapters which each cover one of the models; (iii) the use of reinforcement learning and norm-invocation messages; (iv) habituation; and (v) a conclusion.

### **EMIL: Summary of Outcomes**

The project helped to develop the idea of immergence in a theoretical sense but it also helpfully straddled theory and computational modelling: as mentioned above, the models allowed for both theory development and 'proof of concept'. The project participants did a particularly good job of demonstrating the technical aspects of developing models which combine 'inner' immergence and 'outer' emergence.

However, as we discuss in more detail below, the chapters that describe the models were so weighted toward technical matters that detailed discussions and analyses of the simulation results were missing. This made it impossible to understand precisely what happened in these simulations.

### **Four Computational Models**

The models are as follows:

- Chapter 14 ([Lotzmann, 2010](#)) describes a traffic model in which pedestrians have to cross a road, which - at the same time - car drivers want to travel along as quickly as possible;
- Chapter 15 ([Lotzmann, Emde and Troitzsch, 2010](#)) sets out a model based on the creation and criticism of Wikipedia entries;
- Chapter 16 ([dos Anjos, Lotzmann and Pauli, 2010](#)) is based on a micro-finance model (linked to empirical work in Mexico); and
- Chapter 17 ([Campenni et al, 2010](#)) describes a model in which agents have to navigate three events at an airport (baggage reclaim, customs, and queuing for a taxi).

One of the interesting features of all these models is that they are all grounded in real-world situations. This means the agents' mental models are context-specific with respect to goals and how information is processed. This contrasts with the highly abstract nature of many game theoretic models. Moreover, this grounded-ness is consistent with both the pragmatist orientation of this thesis and sits well alongside [Ullmann-Margalit's \(1978\)](#) 'normalcy' condition for invisible hand explanations.

The agents' mental models are all based on collections of rules, each of which is associated with a probability for its enaction. These probabilities are adjusted via reinforcement

learning given specific goals (consistent with the discussion in Section 1.4.6, p. 25) and norm-invocation messages.

In all four chapters, the authors report that the movement of certain probabilities toward 1 equates to the emergence of social norms. This is the main result of these models and is discussed further below.

Taking a step back, in evaluating computational research (including that developed for this thesis) we typically look for the following four things (at least):

1. theoretical context;
2. a clear description of the model (ideally minimizing coding jargon to make it accessible to a wider audience);
3. a description and analysis of the results - typically this contains the results of some 'default' parameter set, and a discussion of the results when the parameter space is explored; and
4. a discussion of how the results relate to theory.

Taking the EMIL Report as a whole, the theoretical context for the four models is excellent and goes far beyond what we normally see in such research (chapters 2-4 of the report are focused on "theoretical foundations").

On the whole, however, the chapters focused on the computational models are relatively poor vis-à-vis the other three aims stated above.

There was some variation between these chapters, however. Most include a clear description of the real world situation the models were attempting to mimic but fail to include a description of key features of the model (with the possible exception of the airport model). Without clarity of the models' mechanics, results interpretation is made more difficult. Replication of the models would be impossible in all four cases.

The results reported fell significantly short of what we would expect in such research. For example, in Chapter 14 (which is 9 pages in length), the results are reported in the last two paragraphs. Moreover, any analysis was focused on the aggregated outcomes and not on how agents interact and how their mental models change over time. Any discussion of results across the parameter space was meagre and confined mainly to statements that this work had been done (but not reported).

Regarding the fourth aim above, there is some discussion in chapters 18 and 19 about how the results of these models are related to theory. The main point made is that these models and their simulations appear to corroborate the theory. This may indeed

be the case given the authors' exposure to the computational work undertaken but it is difficult to confidently draw this conclusion from chapters 14-17 as they are written.

As mentioned above, an important qualification to these criticisms is that one of the aims of the project was technical in nature: to describe how such models can be built in practice. A substantial proportion of chapters 14-17 was devoted to modelling techniques, e.g., Chapter 15 includes a block of code which takes up 2 of its 20 pages. If we also consider that the report has 22 chapters (in 245 pages) and probably had a word count limit, we can understand why chapters 14-17 were light on results and analysis.

It is, however, surprising that none of these chapters were followed up with published research that developed the results in more detail<sup>9</sup>.

### **Reinforcement Learning and Norm-Invocation Messages**

Reinforcement learning appears to play a central role in the emergence of norms in the EMIL models, and this seems to be an important result.

The role of norm-invocation messages is worth considering in more detail. There are two broad issues. First, when these messages took the form of admonishments in the EMIL models, there appeared to be no cost incurred by the admonishing agent, and the recipients of the messages were assumed to be influenced by them, e.g., the equivalent of feeling embarrassed. From an economics point of view, this seems unsatisfactory because it implies that punishment can be applied at no cost.

Second, norm invocation messages can only be applied if a norm already exists. This mechanism is, therefore, more one of norm diffusion than origination. That is not to argue, however, that it plays no role in the immergence of a norm for individual agents, nor in the 'global cascade' (Watts, 2002) of norms across a population. However, it appears that in these models, the only underlying mechanism of immergence is reinforcement learning.

Let us consider how we should interpret norm emergence via reinforcement learning. This has a direct bearing on the models developed for this thesis because they also use this mechanism as a basis for organic institutional emergence.

As mentioned in the Introduction, reinforcement learning is based on feedback. Let us, approximately speaking, divide such feedback into three (not mutually exclusive) sources:

1. direct from the environment (not other agents), e.g., a fee for importing declarable goods (used in the EMIL Report's airport model);

---

<sup>9</sup>Andrighetto et al (2010) reported some of the results but, again, the details were relatively scant.

2. direct from other agents, e.g., admonishments for contravening what others perceive as a norm (as mentioned above, this can play a role in the emergence of norms for individual agents); and
3. resulting from the nature of the interaction, e.g., taking longer to travel a section of road because you have collided with a pedestrian, as in the first model above; or payoffs in a prisoners' dilemma game because you cooperated but the other agent defected.

An important problem with the models reported in the EMIL Report is that it is difficult to be certain what type of feedback the agents react to. Also, the precise nature of how the agents interacted was not clear in the models (with the possible exception of the airport model). This lack of clarity means it is difficult to understand how reinforcement learning (and norm-invocation messages) worked in the EMIL simulations.

Consider *in extremis* a scenario in which, say, 100 agents receive feedback from the environment only and in which there is no interaction between the agents. If all the agents learn to adopt the same rule then the resulting group-wide 'social norm' is in fact an aggregation of personal norms. This would fit into Ullmann-Margalit's (1978) use of "aggregation" in the context of invisible hand explanations. Furthermore, this process would be of no surprise to us, i.e., it does not meet Ullmann-Margalit's (1978) third characteristic of invisible hand explanations (Section 4.3.2).

Given the lack of detailed analysis it is difficult to be absolutely certain but this appears to be what happens in the airport model at the first two stations: baggage reclaim and customs. The agents did not interact at either of these stations but the authors claim that social norms emerge over time. One of the interesting features of this model is the combination of normative agents who use reinforcement learning and "social conformers" who simply copied those around them: the results show that the latter group helped to amplify the norms the former group converged on.

At the third station, there is (punishment) feedback from other agents if an agent jumps the taxi queue but the interaction between the agents is meaningful in that there is a zero sum game: one agent jumping the queue saves time equal to the time lost by the other agents. However, without data or analysis, it is impossible to understand from Chapter 17 precisely how reinforcement learning gave rise to the social norm of queuing, and we do not know which parts of the parameter space this occurred in and where it did not.

Looking at the Wikipedia model, the feedback agents received appears to have been from all the sources noted above but it is impossible to understand how these interact and under what conditions specific social norms emerge. The same appears also true of the traffic and micro-finance models.

It is worth noting, however, that while the mechanics of emergence are unclear, the results of the traffic scenario appear particularly interesting: in the reported results, pedestrians learned to cross the road where a striped line was marked on the road, and car users learned to slow down when pedestrians crossed at this location. How this came about remains, nonetheless, unknowable from the data and commentary presented.

### **Habits**

The role of habits is discussed in the EMIL Report. Most notably, the authors state that:

...while giving a motivated and reason-based foundation to autonomous norm-obedience, one cannot ignore that even intentional actions become automatic, when habitual. The related behaviour is no longer really decided nor deliberated, but it is just executed as a response to the recognition of a given stimulus in a given context; and the corresponding action is performed under reduced controls and a higher attentive threshold. ([Andrighetto, Campennì and Conte, 2010](#), p. 80, Footnote 43)

This and similar comments elsewhere in the report suggest that in principle, habits should be viewed as playing a role within norm internalisation. Unfortunately, however, this is not explicitly incorporated into the models. We might guess that when a probability associated with an internal rule converges on 1, the authors would interpret this as habit-like, but this is conjecture.

### **EMIL: Conclusions**

As [Elsenbrioch and Gilbert \(2014\)](#) state, the EMIL Project was a good step forward in our understanding and modelling of social norm emergence. The development of the notion of immergence, its combination with some outer-emergence, and related computational modelling were all strong contributions.

However, as discussed in detail above, the simulations based on the four EMIL-like models were poorly described, analysed, and reported. Moreover, there were important questions left open about how reinforcement learning might give rise to social norms, and also how this type of learning corresponds with the process of habituation.

In the same way that the models developed for this thesis can be viewed as extending [Hodgson and Knudsen's \(2004\)](#) line of research, we can say the same about the EMIL Project. Most importantly, reinforcement learning and habituation are employed in the models developed for this thesis but we analyse the results of the simulations in fine detail to understand how the agents' mental models and actions change over time.

## 5.3 Game Theoretic Approaches

In this section we discuss approaches to institutions from a game theoretic perspective. We do this by looking at the subject through the lens of CE, as defined in Chapter 2, and related literature.

It is worth highlighting up front that this literature is vast and a comprehensive survey is well beyond the scope of this thesis.

Despite the scale of this literature it is possible to criticise: (i) the ability of game theoretic models to explain institutions that are assumed in their analyses; and (ii) other assumptions made in much of the game theory literature, especially those related to uncertainty, deductive reasoning, and utility maximization.

Consistent with Arthur (2013), we can think of these models as *first-order approximations* of many interaction types rather than providing realistic explanations based on bottom-up, practical situations (including conditions of uncertainty).

In the next section we discuss models of monetary emergence based on Jones (1976), including Kiyotaki and Wright's (1989) game theoretic analysis. This paper demonstrates two Nash equilibria<sup>10</sup> (which one exists depends on the parameters used) but empirical studies based on the same model show that these equilibria are not observed in full.

Moreover, while this chapter is mostly focused on models of organic institutional emergence, here we focus more on critiquing game theoretic approaches in general rather than specific models. The only exception is Calvert (1995), which we discuss in detail as an exemplar.

Section 5.3.1 below describes the two main conclusions (regarding institutions) drawn from the game theoretic literature: that institutions are either equilibria or correlating devices that achieve equilibria. We look at evolutionary models and those based on repeated games as the main types of equilibrium-based research.

Section 5.3.2 looks at Aoki (2001) as an attempt to bridge the institutions-as-equilibria approach with that of institutions-as-rules. It also discusses Hindriks and Guala (2015), which also attempts to bridge these two.

The third sub-section (5.3.3) considers Field's criticisms of game theoretic attempts to 'explain' the emergence of institutions. We look at Field (1979, 1981, and 1984) in particular.

Section 5.3.4 develops a specific criticism, which is the over-generalization of utility maximization and the resulting de-valuing of mechanisms that explain institutional emergence

---

<sup>10</sup>Specifically, in their 'Model A'.

in the real world. This criticism follows Hodgson (2012) in particular and corresponds to the discussion of instrumentalism and realism in Section 2.3.3.10.

The final sub-section (5.3.5) considers the implications of the various criticisms above for the models developed for this thesis.

### 5.3.1 Institutions via Games

In looking at how game theory has been used to ‘explain’ institutions, the literature appears to be divided, approximately, in to two (Gräbner and Ghorbani, 2019). The first defines institutions as equilibria in games; and the second emphasises institutions as correlating devices<sup>11</sup> that achieve equilibria<sup>12</sup>.

The second strand of this literature is of interest to us but here we focus on the first.<sup>13</sup>

As Aoki (2001) states, the research that represents the ‘equilibrium’ view of institutions can be further divided into two: evolutionary models and repeated games. Let us briefly look at the first category before looking at the second, which includes a discussion of the model developed in Calvert (1995).

#### Evolutionary Models

Foundational texts in this area include Smith and Price (1973) and Smith (1982). These gave rise to an enormous amount of evolutionary game theoretic work that spans many different fields of study, including in the natural and social sciences.

The three characteristics of generalized Darwinism, discussed in Section 2.3.3.5, are helpful for understanding these evolutionary models (see, e.g., Bowles, 2004, pp. 69-84, for a more detailed discussion). First, agents adopt strategies that are subject to *variation* over time: this can take a wide range of forms, including mutation in biological models, and copying. Second, *selection* occurs via what is referred to in the literature as *differential replication*, e.g., how outcomes in bilateral interactions are determined

---

<sup>11</sup>This type refers to devices “external to the game that selects a particular equilibrium by ensuring the coordination of agents’ expectations” (Gräbner and Ghorbani, 2019, p. 14).

<sup>12</sup>Hédoin (2012) provided an example of a correlated equilibrium: imagine two cars approaching an intersection with payoffs for the drivers that mean there are no dominant strategies and no Nash equilibria. An external rule such as “the driver to the right has priority” would represent a way to coordinate the drivers’ actions. This rule is, however, imposed from “outside” of the game and would not constitute an organic institution that has emerged.

<sup>13</sup>There are three reasons for this: (i) this literature overlaps considerably with the research concerned with institutions-as-equilibria; (ii) the thesis is primarily orientated around the emergence of organic institutions; and (iii) while the literature concerned with correlating devices is relevant to the second research question, we are interested in legal rules specifically in the context of mechanisms that correspond to processes observed in the real world (like reinforcement learning and habituation). Below we see that game theoretic models do not appear to explain these mechanisms, which makes the literature concerned with correlating devices less relevant here.

(such as in the Hawk-Dove game or the Prisoners' Dilemma). *Durability* is typically determined by *replicator equations* - a common approach is for the growth of some strategy to be determined by its payoff relative to the mean payoff of the population.

Evolutionary models are typically mathematical, with the population normalized to 1: different strategies, therefore, compete against each other in a proportional sense. These models are typically open-ended and their dynamics are analysed to identify features that correspond to phenomena observed in various types of system.

Institutions in social systems are examples of the phenomena studied. A simulation based on such models can arrive at an Evolutionary Stable State (ESS): this is defined as an equilibrium to which a system will return after some disturbance. Note that these ESSs share some of the characteristics of Nash equilibria but the two are not identical, e.g., [Apaloo et al. \(2015\)](#) argue that Nash equilibria is a concept best applied to stationary (classical) and not evolutionary games, in which ESS is a more relevant concept.

Notable research in this field includes [Sugden \(1989, 2005\)](#), [Young \(1998\)](#), and [Okazaki, Okuno-Fujiwara, and Greif \(1998\)](#).

Equilibrium states are viewed as institutions within this literature. Indeed, it is worth highlighting an important difference between the definition of institutions stated in the Introduction and institutions as equilibria: the former is a rule whereas the latter is a state. [Gräbner and Ghorbani \(2019\)](#) provide a helpful discussion of the resulting definitional tensions, arguing that the stipulative definition of equilibria is difficult to reconcile with taxonomic definitions like [Hodgson's \(2006a\)](#), which is stated in the Introduction.

Finally, [Hodgson's \(2015\)](#) reference to such equilibria as secondary "to the relational framework that generates their possibility." (p. 500), mentioned in the Introduction, is worth repeating here. An equilibrium state might emerge but what are the (primary) mechanisms by which this occurs?

### Repeated Games

This literature is also enormous. Noteworthy research includes [Axelrod \(1984, 1986, 1997\)](#), [Greif \(1989, 1994, 1998\)](#), [Schotter \(1981\)](#), [Milgrom, North, and Weingast \(1990\)](#), [Greif, Milgrom, and Weingast \(1994\)](#), [Calvert \(1995\)](#), and [Bicchieri \(1993, 2006\)](#).

Models have tended to be of three broad types: two agents only; multiple agents playing bilateral 2-player games; and N-player games (where  $N > 2$  players play single games over successive rounds). N-player games overlap with Common Pool Resource problems because these have been modelled as multiple agents playing a simultaneous prisoners' dilemma game over time, e.g., [Schindler \(2012\)](#).



In earlier work in multi-agent systems, researchers typically modelled a world of agents that randomly interact with other agents. Subsequent work developed structural features like networks with agents who interact with neighbour agents only. Classic network structures have been investigated e.g. scale-free and small world, to consider their impact on cooperation. For example, [Seltzer and Smirnov \(2015\)](#) combine an iterated prisoners' dilemma model with social networks, showing greater cooperate with fewer degrees of separation.

Agents in these models typically have two strategies to choose from in each game, e.g., defect or cooperate in iterated prisoner dilemma games; however, they can also develop dynamic strategies where an agent might choose from a range of such strategies depending on which has historically gained the most benefit. A classic example of such a dynamic strategy is 'tit-for-tat' where (in 2-player games) an agent mimics the behaviour of its opponent from the previous game: the opponent will be 'punished' if it defected previously, and vice versa. An example of this work is [Isaac \(2008\)](#), and a classic text is [Nowak and May \(1992\)](#).

The orientating question in this literature seems to be about the conditions under which cooperation is a stable strategy for all or most players. In particular, is a broadly cooperative system robust given an 'invading defector'?

Some of the research in this literature has investigated various human characteristics which might facilitate sustained cooperation. An example is [Schindler \(2012\)](#), which discusses 'socio-psychological factors' including parameters linked to cooperativeness, positive reciprocity, fairness towards others, and risk aversion.

This literature is both useful and revealing in that many authors have demonstrated how certain traits, under specific conditions, are consistent with sustained cooperation in multi-agent systems. Whether this equates to the emergence of an institution or not depends on one's definition of that term, e.g., [Calvert \(1995\)](#) argues there is no institution because of the bilateral nature of the dynamic strategies observed.

Let us now focus on [Calvert \(1995\)](#), which developed [Schotter's \(1981\)](#) view of "an institution as an equilibrium of behavior in an underlying game." ([Calvert, 1995](#), p. 58).

*Calvert (1995)*

Aligning his model with substantive rationality, [Calvert \(1995\)](#) states that institutions as "patterns of behavior and expectations must be consistent with utility maximization of each individual." (p. 59).

More specifically, Calvert argues that we define an institution as:

... an equilibrium in which individuals' actions are dependent upon the past actions of many others, or upon expectations about the future reactions of many other players, to one's present actions. (Calvert, 1995, p. 59-60)

The model Calvert develops is based on an iterated prisoners' dilemma game involving  $N$  agents who are randomly paired in each iteration of the model. Importantly, Calvert states the conditions under which this type of game achieves an equilibrium over time under the assumption that the agents employ a type of tit-for-tat strategy. Furthermore, he develops variations of the model in which: (i) multilateral communication is possible between the agents; and (ii) a formal rule is applied.

Calvert states clearly that his model is not intended to show the emergence of an institution. It only provides information about the parts of the parameter space in which we would expect an institution to be maintained within Calvert's games.

In concluding, Calvert states:

... there is, strictly speaking, no separate animal that we can identify as an institution. There is only rational behavior, conditioned on expectations about the behavior and reactions of others ... all in an attempt to maximize utility ... . *Institution* is just a name we give to certain kinds of equilibria. (Calvert, 1995, p. 74, emphasis included)

For the discussion below, we should ask what Calvert (1995) says about the situations in which his model and results are viewed as relevant. He focuses on scenarios in which the 'players' in the game could be thought to be behaving in a broadly "rational" way, e.g., in his conclusion he discusses "congressional committees and the exPost Veto" (pp. 75-79).

However, he also states that his approach should be helpful "for all kinds of institutions, from the most particular matters of legislative procedure to *the broadest aspects of social order*." (Calvert, 1995, p. 83, emphasis added).

Whether or not Calvert's type of model is applicable to "the broadest aspects of social order" (he gives no examples of what he means by this) is a fundamentally important question. To what extent can models that assume deductive reasoning and utility maximization generalise to interactions in which neither is true? We discuss this further below.

### 5.3.2 The Rules-as-Equilibria Approach

Here we discuss a strand of literature that attempts to combine the equilibrium approach to institutions with the view that institutions are rules. The main text we look at is Aoki (2001) but we also discuss Hindriks and Guala (2015). Both of these refer to Greif's

work as a strong influence, e.g., Greif, Milgrom, and Weingast (1994), and Greif (1994, 1998); and both mention Lewis [1969] (2008) as significant.

Looking in more detail at Aoki (2001), it is perhaps better to think of this work as a variation of the equilibrium view rather than an approach that treats rules-as-equilibria with equal weight. Aoki's analysis is still heavily reliant on the framing and terminology of game theory. Moreover, Aoki developed his thinking in subsequent work, notably in Aoki (2011) but, for our purposes, focusing on Aoki (2001) is sufficient for understanding the thrust of his work.

By way of additional motivation, it is worth noting that Aoki (2001) contains a number of features that sit comfortably with the CE ontology promoted in Chapter 2.

The main aim of Aoki's research is to marry the equilibrium definition of institutions (discussed above) with an agent-level explanation of institutions, i.e., to "subsume [Greif's (1994)] player-of-the-game view" (p. 9). To understand this, it is helpful to state Aoki's (2001) tentative characterization of an institution at the beginning of his book, as "a *self-sustaining system of shared beliefs* about a salient way in which the game is repeatedly played." (p. 10, emphasis included).

This definition can be interpreted as a taxonomic definition, which is different to the stipulative definitions we often see in Game Theory. Gräbner and Ghorbani (2019) contains a helpful good discussion of this point: the goal of a "stipulative definitions in mathematics ... is to allow for very precise analysis and logical derivation [whereas] the primary goal of taxonomic definitions is to establish shared meanings about phenomena within a scientific community" (Gräbner and Ghorbani, 2019, p. 3-4). This distinction is important for the discussion below.

Aoki refers to the above characterization of institutions as his *equilibrium-summary-representation* approach. In Part 2 of his book, Aoki develops this in order to articulate his more advanced *shared-beliefs cum equilibrium-summary-representation* framing, which is also discussed below.

Aoki's (2001) reference to shared beliefs in the above definition correspond to agents having a "*summary representation (compressed information) of an equilibrium* of the game." (p. 10, emphasis included). The reference to "summary representation" and "compressed information" are concerned with the agents having limited information and cognitive constraints.

Furthermore, by "salient", an agent might only be aware, tacitly, of some feature of the equilibrium; or there might be a symbolic representation outside of its mind.

Aoki's framing also contains the view that institutions help individuals to form expectations about each other, which is both intuitively appealing and fits other researchers'

reference to this role, e.g., [Hayek \(1973\)](#), [North \(1990\)](#), [Ostrom \(1991\)](#), [Hodgson \(2006a\)](#), and [Gräbner and Ghorbani \(2019\)](#).

In terms of detailed aims, [Aoki \(2001\)](#) refers to his *synchronic* approach that corresponds to understanding institutions across the economy in the present; and his *diachronic* approach, which is to “understand the mechanisms of institutional evolution/change in a framework consistent with an equilibrium view of institutions, but allowing for the possibility of the emergence of novelty.” (pp. 2-3). This second, dynamic, aim is related to our first research question but it is focused on the *evolution* of existing institutions.

In addition to all of the above theoretical work, in Part 1, [Aoki \(2001\)](#) develops a number of empirical examples of institutions, in the context of the definition stated above. Overall, this part of Aoki’s book fits fairly comfortably with CE as it was defined in Chapter 2 above.

In Part 2, Aoki attempts to develop a more advanced characterization of institutions from his tentative definition. Surprisingly, having set out a taxonomic definition of institutions earlier in the book, in Chapter 7 Aoki falls back on the framing and language of game theory to develop this more advanced definition. [Gräbner and Ghorbani \(2019\)](#) refer to this type of definition as stipulative, as discussed above. The ‘core’ mathematical work, including what Aoki refers to as an institution, is contained in Section 7.2 (pp. 197-202).

[Gräbner and Ghorbani \(2019\)](#) characterise Aoki’s definition in the following way<sup>14</sup>:

Institutions in their deep structure are commonly cognized, salient patterns of the ways in which societal games are recursively played and expected to be played. Institutions in their substantive forms are social artefacts that ensure the societal games are in equilibrium. ([Gräbner and Ghorbani, 2019](#), p. 6)

The distinction between stipulative and taxonomic definitions noted above is important because it points to a considerable gap between the niche language and framing of game theory and the general discussion of institutions in the literature. We might think of a stipulative definition as a jigsaw piece that stands alone, not quite fitting into the overall picture. [Gräbner and Ghorbani \(2019\)](#) is helpful again by arguing that a stipulative “definition of institution in a particular model should refer to and complement, but not replace a taxonomic definition” (p. 5). [Aoki \(2001\)](#) does not provide such a ‘bridging’ taxonomic definition.

In the context of our first research question, another problem with [Aoki \(2001\)](#) is that the focus of his diachronic work is on the *evolution* of existing institutions (see chapters

---

<sup>14</sup>Note that this definition is not from Aoki and that it incorporates two later articles: [Aoki \(2007\)](#) and [Aoki \(2011\)](#).

7, 9, and 10 in particular) rather than the *emergence* of new institutions. Takizawa (2017) argues that Aoki “has not yet achieved the second goal of his book (Aoki, 2001): to understand how a novel institution is created.” (p. 536). However, this is an unfair criticism. Aoki (2001) is clear (see pp. 2-4) that his aim is to understand how institutions might evolve as a result of novel change: his focus is not on the novel emergence of new institutions.

In summarising Aoki (2001), we can say that this book was a commendable and novel attempt to develop an approach to institutions that marries both the institutions-as-equilibria and institutions-as-rules views<sup>15</sup>; and it shares a number of features with CE, including relaxing some of the restrictive assumptions of game theory (notably in Chapter 9). However, given the criticisms of game theoretic approaches to institutions below, it is disappointing that Aoki fell back on game theory to develop his more advanced definition of institutions. Furthermore, his book contained no attempt to explain the origination of new institutions, which is the main focus of this thesis.

Hindriks and Guala (2015) is another attempt to marry institutions as equilibria and rules. We discuss this briefly as its part-reliance on game theory means it shares some problems with Aoki (2001).

At the heart of Hindriks and Guala’s (2015) article is an attempt to use Searle’s (1995) framing of *constitutive rules* to synthesise the equilibrium and rules approaches to institutions. Constitutive rules differ from ‘regulative rules’: the former are structured as “X counts as Y in C”<sup>16</sup> whereas the latter generalise as “in circumstance X do Y” (as stated in the Introduction).

A detailed discussion of Searle’s framing (which is summarised well in Searle (2005), Takizawa (2017), and Gräbner and Ghorbani (2019)) and how it is used by Hindriks and Guala (2015) to create a “unified theory”, is beyond the scope of the thesis. Here we constrain ourselves to specific criticisms of their work.

The overarching criticism is that the argument Hindriks and Guala (2015) developed is not compelling: creating a unified theory that brings together two major strands of literature is a major undertaking and their article was 22 pages long. This means that parts of their argument were left without requisite substantiation or support.

For example, Hindriks and Guala (2015) include a short discussion of the equilibrium approach to institutions in game theory, using the Hawk-Dove game and the institution of property as an example. Immediately after this section they state that “So institutions must be *correlated* equilibria of *coordination* games with multiple equilibria.” (p. 466,

<sup>15</sup>in addition to developing *Comparative Institutional Analysis*.

<sup>16</sup>Here, “the X term is an object, a person, or state of affairs, the Y term a status function assigned to X, and C a context.” (Takizawa, 2017, pp. 532-533).

emphasis included). This is a strong statement in light of their preceding section, which did not attempt to deal with the many criticisms of game theoretic approaches to institutions in the literature.

Hindriks and Guala's (2015) discussion of Searle's constitutive rules was also too light given the weight of their claim. For example, these types of rule rest on the idea of collective intentionality, which "covers not only collective intentions but also such other forms of intentionality as collective beliefs and collective desires." (Searle, 2005, p. 6).

Importantly, Gräbner and Ghorbani (2019) argue that this "concept of collective intentionality [is] specialized, yet in the philosophical literature still [a] contested concept." (p. 24). If Hindriks and Guala (2015) want to base their unified theory on Searle's framing then they have to discuss this (vague) concept in more detail and support it with robust arguments.

We should note that Hindriks and Guala (2015) also did not support their arguments with any substantial empirical evidence, as Aoki (2001) did. This is another important omission in the development of a unified theory.

The final point to note here is that Hindriks and Guala (2015) was "more concerned with understanding the general nature of institutions" (Hodgson, 2015, p. 498) rather than on the emergence of institutions.

Now that we have discussed various approaches to institutions in the game theory literature, let us now turn to the criticisms of such models.

### 5.3.3 Field's Criticisms

Here we summarise Field's criticisms of game theoretic models that claim to 'explain' institutional emergence. The first two papers considered below (Field, 1979 and 1981) are targetted at "Neoclassical" and "rational choice theory" models but they provide useful background to Field (1984), which is focused on game theoretic models.

Field (1979) helpfully draws attention to the distinction between *explaining* and *analysing* economic phenomena, defining explanation as "to make clear the cause or reason of, or to account for." (p. 50).

He argues that "economists generally have been reluctant to describe their activities as explanation" (Field, 1979, p. 51), preferring instead to focus on the weaker notion of analysis (cf partial, general, and microeconomic analysis)<sup>17</sup>.

---

<sup>17</sup>Although he notes that some economists "have often, in an effort to attract attention to the power of their analytical framework, claimed to have explained, whereas they have only analyzed." (Field, 1979, p. 51).

In their use of rational choice theory, Field (1979) writes that “whereas economists have in some sense been able to analyze institutional structures using economic theory, they have not come any closer to explaining them” (p. 52). We will return to this criticism of explaining institutional emergence in the next section when we look at utility maximization.

Field (1981) adds an additional argument in criticising “Neoclassical” models of institutions, targeting North and Thomas (1973) and Posner (1973) in particular. These models “attempt to make endogenous structural or institutional characteristics which have previously been treated as noneconomic givens” (Field, 1981, p.174). Here, Field is querying the extent to which society’s institutional structures can be explained by purely economic phenomena.

He develops this argument further in Field (1984), which questions the ability of game theoretic models to explain the very institutions they assume in their analyses. Field believes that such models can be useful in “the case of a limited number of regulative rather than constitutive rules ... in the sense that the posited choice among rules presupposes shared language as well as the prevailing more fundamental set of rules.” (Field, 1984, p. 691). However, these models cannot explain “the origins of language or of the more fundamental constitutive rules of a group or society are concerned” (ibid).

Regarding language, Field (1984) points out that even simple forms of cooperation require language to bring about a cooperative solution, which in turn challenges the idea that language itself can come about via cooperative games. This is the infinite regress problem noted in the Introduction.

Field also refers to “fundamental constitutive rules” as not being explainable by a game theoretic analysis. Here, Field seems<sup>18</sup> to be referring to:

norms established through the process of socialization, perhaps ‘voluntarily’ accepted or affirmed, perhaps building on certain genetic predispositions, provide part of the framework within which individuals pursue their self-interest. (Field, 1984, p. 705)

We should note that Field’s (1984) criticisms do not apply to all of the research based on game theoretic models. Most notably, Aoki (2001) was very clear that he understood how many pre-existing institutions like language, and related ‘cultural phenomena’, had to be assumed in analysis.

---

<sup>18</sup>Field is not unambiguous about this.

Overall, Field provides a number of significant criticisms of game theoretic efforts to ‘explain’ institutions. Most importantly, he raises questions about the nature of ‘explanation’ in these types of analyses and also the types of institutions game theory models can reasonably claim to ‘explain’. Both of these are discussed further below.

### 5.3.4 The Value of Game Theory Models of Institutions

When we consider the question of what value game theoretic models like that developed in Calvert (1995) have, we find two broad groups of opinion.

One group (referred to here as the ‘narrow-value’ group) believes that such models are helpful<sup>19</sup> only in situations that are consistent with the assumptions and framing of these models. For example, game theoretical analyses were extremely helpful during the cold war for making sense of nuclear arms escalation, and in understanding the requirements for significant de-escalation (such as in the START Treaty of 1991).

This point of view means that these models, however, would not be useful for situations when the assumptions do not hold, e.g., under conditions of extreme uncertainty when agents cannot reason deductively (discussed further below).

A good example of this narrow-value point of view is expressed in Tuomela (2002). A model in which:

a social institution is an equilibrium point in a repeated game ... is restricted in that it applies only to activities that can be modelled as a game in the sense of game theory [and] it is idealized in requiring equilibrium behavior ... Thus, the requirement of equilibrium behavior, if informatively specified, often is false in actual life, or else it tends to be a tautological requirement if left fuzzy and vague. (Tuomela, 2002, p. 157)

Tuomela (2002) expands on his criticisms of Schotter’s (1981) model in Endnote 2, pp. 254-256. This 10-point criticism can be summarised as arguing: (i) that these game theoretic models only have value in very specific, narrow circumstances that match the assumptions used; and (ii) the space of other types of interactions observed in the world is enormous<sup>20</sup>.

The second group (referred to here as the ‘broad-value’ group) takes the opposite view: models like that seen in Calvert (1995) help us make sense of interactions even when the assumptions do not hold. We can see this perspective across a lot of work in game theory, e.g., Bicchieri (2006) describes how such game theoretic models help explain

---

<sup>19</sup>Let us say helpfulness here is made up of: (i) making sense of reality; and (ii) choosing between options given some aim(s).

<sup>20</sup>Note that this narrow-value perspective can also be seen in Field (1984), discussed above.



social norms as “The Grammar of Society” (the title of her book) as well as “The Nature and Dynamics of Social Norms” (the sub-title).

These two groups are clearly very different: how do we make sense of them?

In the context of the instrumentalism / realism discussion in Chapter 2, the narrow-value group appears aligned with realism (i.e., game theoretic models are helpful when their narrow assumptions hold but not otherwise) and the broad-value group seems more aligned with instrumentalism. The latter would mean the results of the models ought to aid prediction but the internal mechanisms need not correspond to reality.

Let us briefly consider the implications of realism in the context of the game theoretic models discussed above.

The narrow-value group’s affinity with realism generally leads to a greater appreciation of uncertainty in real-world situations than game theoretic models. Here, uncertainty includes that resulting from a lack of information (Schofield, 1985; and North, 1990) and knowledge (North, 2005); cognitive limitations (Heiner, 1983; and North, 2005); mutual contingency<sup>21</sup> (Binmore, 1987; Arthur, 1994; and Rosser, 1999); and ontological emergence<sup>22</sup> (Ladyman, Lambert and Wiesner, 2013).

In questioning the prevalence of uncertainty in economies, we might recall from Chapter 2 North’s (2005) statement that “uncertainty is not an unusual condition; it has been the underlying condition responsible for the evolving structure of human organization throughout history and pre-history.” (p. 14).

Furthermore, deductive reasoning breaks down in the face of uncertainty (Arthur, 1994), which raises the question of how agents make sense of their situations and make decisions. Put another way, what forms of mental models (Holland et al, 1986; and Denzau and North, 1994) are appropriate in these situations?

Game theoretic models typically assume certainty in the agents’ context, and in terms of mental models they tend to assume utility maximization and deductive reasoning. The latter is not true of all game theoretic models but utility maximization appears to be a consistent assumption. Let us look more closely at this important assumption now.

### **Hodgson’s Criticism of Rational Choice Theory**

Hodgson (2012) is a particularly helpful paper for this subject. The author has elsewhere advocated a pragmatist stance and encouraged the use of the complexity sciences in

---

<sup>21</sup>This arises when two or more agents are trying to anticipate each other’s actions, simultaneously.

<sup>22</sup>This is a ‘strong’ form of emergence, the outcome of which is impossible to predict, even in principle.

institutional research so his paper on the limits of rational choice theory seems especially relevant here.

In this context, it is tempting to interpret the narrow-value group above as promoting an empirically valid account of human nature and interaction. However, [Hodgson \(2012\)](#) argues that an appeal to empirical evidence is something of a red herring because any theory that assumes utility maximization can be made to fit a concept that is unobservable and therefore hypothetical. This follows [Samuelson \(1937\)](#).

The more fundamental problem with a utility maximization approach, Hodgson argues, is that of excessive generalisation<sup>23</sup>, to the point of being close to useless in specific situations of interest. [Hodgson \(2012\)](#) writes that there are at least two important issues missing in a utility maximization account: first, “it neglects the problem of *explaining the causes* of behaviour. Second it fudges the question of the individual *development* of capacities and dispositions.” (p. 99, emphasis included). We might think of the first of these as about explaining behaviour in the moment and the second about learning over time.

It is helpful here to differentiate between two different (but easily conflated) arguments for empirical accuracy. One is a more general belief that game theoretic models should be accurate vis-à-vis human cognition, motivations, and environmental conditions in principle. The second treats empirical evidence as important because the aim is to *understand real mechanisms*, e.g., concerning organic institutional emergence.

[Hodgson’s \(2012\)](#) arguments suggest that the narrow-value group, like Tuomela, are on much firmer grounds if their appeal to empirical evidence is based on the second of these.

Hodgson’s argument of the excessive generality of utility maximization is also helpful because it applies to a wide variety of mental models, i.e., the cognitive means to achieve these ends. Notably, game theoretic models that have sought to include some of the interesting results in behavioural economics still adopt utility maximization as an aim.

Regarding the avenues of research that economists have followed, [Hodgson \(2012\)](#) writes that a problem is that some “still cling tenaciously to the principles of rationality, in a manner that is reminiscent of Ptolemaic astronomers, fitting the evidence of the apparent circular movements of the stars into complicated models.” (p. 103). This is reminiscent of [Kuhn’s \(1962\)](#) reference to researchers tinkering at the edges of some paradigm before it is superseded by another.

The ontology this thesis is based on, which includes the bottom-up orientation of the complexity sciences and pragmatism (notably Dewey’s), leads us to agree with [Hodgson’s](#)

---

<sup>23</sup>Rational choice theory has been applied to, for example, “politics, marriage, religion, suicide, and much else” ([Hodgson, 2012](#), p. 102); but also to “a large portion of the animal kingdom as well.” (ibid, p. 99).

(2012) arguments that utility maximization is too generalised. What is needed is “some real institutional and cultural flesh and blood.” (ibid, p. 103) if we are to explain the mechanisms of organic institutional emergence. As noted above, Field (1984) came to the same conclusion regarding the limited ability of game theory to ‘explain’ such phenomena.

Hodgson’s conclusion chimes particularly well with the following quote from Arthur (2013):

Equilibrium of course will remain a useful first-order approximation, useful for situations in economics that are well-defined, rationalizable, and reasonably static, but it can no longer claim to be the center of economics. (Arthur, 2013, p. 19)

This quote is targeted at Neoclassical theory in general but it seems appropriate to the equilibrium-orientation of game theoretic models more specifically.

Finally, looking at game theory from a Cognitive Institutional Economics (CIE) perspective, Ambrosino, Fontana, and Gigante (2018) argue that the “agents’ interactions leading to the emergence of institutions is far more complex than how [New Institutional Economics] describes them in its game theory models.” (p. 777). As a result, “CIE takes a critical stance on game theory.” (ibid, p. 784).

The conclusion we draw is, as stated previously, that the game theoretic models of institutions are helpful as first-order approximations but, as Field (1984) and Hodgson (2012) emphasise, they tell us little about mechanisms of behaviour in the real world. This is important to us in this thesis: recall from the Introduction (and Chapter 2) the discussion of practical institutional problems that arise as an economy changes, including the topical example of cryptocurrencies. Without a realistic understanding of the mechanisms of institutional emergence, how do we move forward?

These points both act as the underlying criticism of game theoretic models of institutional emergence and they provides a motivation for the models developed for this thesis. Let us consider this, briefly, now.

### 5.3.5 Implications for Modelling Organic Institutional Emergence

The models developed for this thesis are based on the ontology of economic systems described in Chapter 2.

In the context of game theoretic models of institutional emergence, an emphasis is placed on the bottom-up orientation of the complexity sciences (cf Section 2.1.1) and

on pragmatism (Section 2.3.3.11). This is consistent with Hodgson's (2012) conclusion that to find explanations of behaviour we must draw upon "psychology, anthropology, sociology, and other disciplines" (p. 101). Moreover, this aim sits comfortably alongside Ullmann-Margalit's (1978) normalcy condition for invisible hand explanations.

More concretely, in the models below, the agents make decisions under conditions of uncertainty, including situations when information is scarce but also when interactions involve mutual contingency, i.e., when the agents cannot make decisions deductively because they cannot deduce what other agents will do.

As a result, the agents' mental models are based not on deductive reasoning but on pragmatic forms of reasoning, learning, and habit formation. As stated above, Arthur (1994) argues that deductive reasoning, which is implicit in substantive rationality, breaks down in the face of uncertainty.

In addition, the agents must consume two resources to survive. There is no utility maximization although the aim is still quantified.

We should note, however, that the development of models that precisely replicate cutting edge "psychology, anthropology, sociology, and other disciplines" in a detailed way is beyond the scope of a single thesis. More modestly, the aim of the models is to take a step *towards* such an approach. As mentioned previously, these models can be viewed within the family of those developed in Hodgson and Knudsen (2004) and in the EMIL Project.

## 5.4 Models of Monetary Emergence

This section discusses a specific strand of literature concerned with the emergence of money. This research is often traced back to Kiyotaki and Wright (1989) but their work is in turn based on Jones (1976). Looking even further back, this literature can be viewed as exploring Menger's ([1890] 1981, [1883] 1985) account of monetary emergence, which was discussed in Section 3.2.3 above.

It is worth noting briefly that money is thought of as a convention in this literature, i.e., it provides a coordination role in which (self-sustaining) Nash equilibria are thought to exist. The 'quality of money' argument discussed in Section 3.2.3 is not taken into account. This literature also assumes that the need for money arises because of a division of labour: there is no consideration of why agents would specialise in the absence of media of exchange.

One of the most interesting features of this literature is the combination of: (i) game theoretic analysis (Kiyotaki and Wright, 1989); (ii) agent-based models (notably Marimon et al., 1990); and (iii) empirical studies (Brown, 1996; Duffy and Ochs, 1999; and Duffy, 2001 in particular). In this section we summarise and critique these three.

We first look at Kiyotaki and Wright's (1989) Model A to orientate us around how the problem is typically framed in most of the literature. This is the subject of Section 5.4.1, which includes a statement of the conditions under which two Nash equilibria exist. Sethi (1999) supports Kiyotaki and Wright's (1989) results by demonstrating that their Nash equilibria would be reached in a dynamic setting.

The second sub-section below (5.4.2) looks at Marimon et al.'s (1990) computational model, which was an attempt to consider whether Kiyotaki and Wright's (1989) Nash equilibria would be arrived at by interacting agents using classifier systems.

Section 5.4.3 considers the empirical evidence: studies have sought to replicate Kiyotaki and Wright's (1989) model in laboratory settings with real human subjects. We find that this empirical evidence is inconsistent (though not fully) with both Kiyotaki and Wright's (1989) and Marimon et al.'s (1990) analyses. Duffy (2001) represents an attempt to reconcile the empirical evidence with an agent-based model.

### 5.4.1 Kiyotaki & Wright's Model of Monetary Emergence

Kiyotaki and Wright's (1989) Model A is summarised in Fig. 5.3 below<sup>24</sup>. There are three agent types: each produces a good which it does not want to consume and the circularity of production and consumption means there is never a 'double coincidence of wants' (Jevons, 1875) corresponding to any two agents' production and consumption goods. This sets up a simple interdependence problem and the question is whether any of the commodities would emerge as a form of money<sup>25</sup>.

In the model, an infinite number of agents (denoted  $i$ ) interact bilaterally in each round. They gain  $u_i > 0$  of utility (identical for all agents) if they consume their preferred good, and zero otherwise. To compare utilities between rounds, a discount factor is used ( $\beta < 1$ , also identical for all agents). There are an infinite number of rounds, agents have infinite lives, and the aim of each agent is to maximize the net present value of its present and future utility.

A transaction occurs only if both counterparties agree to sell the resource they hold. If either refuses to sell, there is no transaction.

---

<sup>24</sup>Note that Kiyotaki and Wright's (1989) Model B is the same as Model A except the production / consumption arrows in Fig. 5.3 are reversed. This is not symmetrically the same as Model A because of the asymmetry of storage costs of the three resources.

<sup>25</sup>Type  $i$  agents consume Resource  $i$  and produce Resource  $(i + 1) \bmod 3$ . In Model B they produce  $(i + 2) \bmod 3$ .




Figure 5.3: A summary of [Kiyotaki and Wright's \(1989\)](#) production and consumption environment (Model A). A group of agents is divided into 3 types. Each arrow relates to a specific resource: the square at the base of each arrow correspond to production and the pointy ends correspond to its consumption by each agent type, i.e., Type 1 agents consume Resource 1 and produce Resource 2; Type 2 agents consume Resource 2 and produce Resource 3; and Type 3 agents consume Resource 3 and produce Resource 1. The design of the problem space ensures there is no double coincidence of wants vis-à-vis any two agents' production and consumption resources.

Agents can only store 1 resource unit between rounds. If an agent does hold a resource between rounds, it incurs a storage cost, denoted by  $c_i$ . These storage costs differ such that  $0 < c_1 < c_2 < c_3$ . Furthermore, if an agent trades for its preferred resource, this is consumed immediately and the agent produces its corresponding production good (at no cost) before the round is finished.

[Kiyotaki and Wright \(1989\)](#) identify two Nash equilibria (depending on  $c_2$ ,  $c_3$ ,  $\beta$ , and  $u_1$ ). They refer to these as a 'Fundamental Equilibrium' and a 'Speculative Equilibrium' (illustrated in figures 5.4 and 5.5 below, respectively).

The Fundamental Equilibrium shown in Fig. 5.4 is quite simple: Type 2 agents intermediate between Type 1 and Type 3 agents, accepting Resource 1 from Type 3 agents and then trading it for Resource 2 from Type 1 agents. Type 1 agents will only store Resource 2 and Type 3 agents will only store Resource 1. Type 2 agents store Resource 1 half the time and Resource 3 half the time.

This steady state equilibrium exists when the following condition holds:

$$c_3 - c_2 > \frac{1}{2} \frac{\beta}{3} u_1 \quad (5.1)$$




Figure 5.4: Kiyotaki and Wright's (1989) Fundamental Equilibrium. Types 1 and 3 only transact with Type 2 who essentially intermediates between them. Resource 1 is viewed as a form of money in this pattern.

We can understand this condition intuitively. The LHS of the inequality represents the (opportunity) cost of Type 1 agents holding Resource 3 instead of Resource 2 (they will only consume Resource 1 and never hold it). The RHS is the probability-weighted expected utility (discounted) of holding Resource 3 instead of Resource 2 in the current period, and in equilibrium. If the LHS exceeds the RHS then Type 1 agents will refuse to hold Resource 3 at any time.

A Speculative Equilibrium (depicted in Fig. 5.5 below) exists if:

$$c_3 - c_2 < (\sqrt{2} - 1) \frac{\beta}{3} u_1 \quad (5.2)$$

In this type of equilibrium, Type 1 agents believe it is worthwhile holding Resource 3 when offered (the expected utility of this strategy exceeds its opportunity cost), hence Type 1 agents will 'speculate' that they might sell Resource 3 for 1 in the next round. In this literature, resources 1 and 3 are seen as having money-like properties.

It is worth emphasising here that the difference between the fundamental and speculative equilibria is in whether Type 1 agents store and trade Resource 3.

Note that when:

$$(\sqrt{2} - 1) \frac{\beta}{3} u_1 < c_3 - c_2 < \frac{1}{2} \frac{\beta}{3} u_1$$




Figure 5.5: Kiyotaki and Wright's (1989) Speculative Equilibrium. This is identical to the Fundamental Equilibrium outcome except this time Type 1 agents will accept Resource 3 from either Type 2 or 3 agents. Both resource 1 and 3 are viewed as money-like: Resource 1 is used as 'money' by Type 2 agents and Resource 3 is used by Type 1 agents.

there is no determinable equilibrium.

The simple, interdependent framework developed and analysed in Kiyotaki and Wright (1989) seems a neat way of thinking about money as an organic institution (here, a convention). Also, the Nash equilibria are helpful in helping us understand which outcomes would be self-sustaining.

In the context of this thesis, there remains a crucial question of whether such institutional equilibria can *emerge* in a population.

Related to this, Sethi (1999) demonstrates that not only are Kiyotaki and Wright's (1989) Nash equilibria both stable, we should also expect convergence on these equilibria. This is true of the equilibria pertaining to Model A but also Model B: for the latter, convergence to either equilibrium can occur but which emerges depends on the initial conditions.

It is worth noting, also, that Sethi (1999) states that "stringent assumption regarding information and rationality are not necessary in order to explain the use of money as a medium of exchange." (p. 246). This seems a strong statement in light of the fact he does not explore whether a wide variety of mental models give rise to these Nash equilibria. Marimon et al. (1990), to which we now turn, indicates that this is not true.



### 5.4.2 Marimon et al's Computational Model

In the words of Marimon et al. (1990), “while [Kiyotaki and Wright (1989)] studies stationary equilibria in which beliefs about ‘media of exchange’ are *consistent* with trading patterns, we study economies in which particular commodities *emerge* as media of exchange.” (p. 330, emphasis included).

In this subsection we focus on Marimon et al.'s (1990) evolutionary model and then comment on other evolutionary models at the end.

The economic environment in which the agents are located is essentially the same as that of Kiyotaki and Wright's (1989) Model A<sup>26</sup>. However, here, the agents use *classifier systems* as described in Holland (1975), which means the agents' mental models are based on rules rather than substantive rationality.

Each agent has two classifier systems: one related to exchange (concerned with whether an agent should offer to trade their stored resource) and another related to consumption (whether or not they should consume the resource they hold at the end of the round). Within these two systems, classifiers are categorised according to the conditions an agent might face such that for a given set of conditions (e.g., I hold Resource 2, my counterpart Resource 3) there are a number of potential rules which apply (e.g., offer to trade or not).

Each rule in a group of classifiers is given a weight that changes over time depending on experience: at any given time, the rule with the highest weight is chosen<sup>27</sup>. The change in weights occurs via a type of reinforcement learning: if the agent benefits from a rule, its weight increases; and vice versa. Note this reinforcement learning can occur inter-temporally, e.g., if an agent chooses a rule which means it keeps its stored resource at time  $t$  but which it then trades for its consumption resource in  $t + 1$ , the weight of the rule applied in  $t$  increases.

Marimon et al. (1990) conduct two sets of experiments. In one the agents are allowed access to every rule possible; and in the second, the agents have a limited set but they use a genetic algorithm “as a device for periodically eliminating some rules and injecting new rules into the population...” (Marimon et al., 1990, pp. 330-331).

#### Initial Conditions and Parameters

At instantiation each agent ( $i$ ) is endowed with the resource it can produce:  $(i+1) \bmod 3$ .

<sup>26</sup>The authors also explore Model B but most of their simulations are based on Model A.

<sup>27</sup>This is referred to below as a “Winner Takes All” approach. In the default simulations of the first model presented in Chapter 7 below we consider this and also a “Roulette Wheel” approach whereby the *probability* a rule is chosen is determined by its relative weight.

Marimon et al. (1990) are interested in both the fundamental and speculative equilibria so they (attempt to) adjust parameters to correspond with these scenarios.

For simulations related to a Fundamental Equilibria, they select the following parameters:

$$u_i = 100$$

$$c_1 = 0.1$$

$$c_2 = 1.0$$

$$c_3 = 20$$

When focused on the speculative equilibrium, the same parameters are used except  $u_i = 500$ .

An important difference between Kiyotaki and Wright's (1989) framework and Marimon et al.'s (1990) architecture is that  $\beta$  (the discount rate applied to future utility) is not incorporated into the latter. One implication of this is that we cannot specify whether the Speculative Equilibrium inequality (5.2 above) holds or not.

In fact, the exclusion of  $\beta$  from Marimon et al.'s (1990) model raises a larger question about its comparability with Kiyotaki and Wright's (1989) equilibrium-based analysis. This simple difference highlights the distinction between the types of first-order approximations seen in game theoretic analyses and attempts to consider mechanisms closer to those seen in reality. This was the main theme of the last section.

In Kiyotaki and Wright's (1989) analysis,  $\beta$  is necessary to enable agents to calculate their expected utilities (to infinity) of different strategic choices in the present. This corresponds to Hodgson's (2012) reference to treating "individuals as capable of emulating incredible super-calculators with unbounded cognitive capacities, without any consideration of how they would manage to do this." (p. 100).

We can think of Marimon et al.'s (1990) model as a step away from a first-order approximation and closer to a realist approach (not that they stated this). Agents make decisions via rules and the probabilities of selecting those rules change over time. In this model, the agents learn through reinforcement triggered by utility gains and storage costs. Moreover, while we can imagine a version of this model which discounts expected utility in the future to the present, this is not done in Marimon et al.'s (1990) model:  $\beta$  is not relevant in this framework<sup>28</sup>.

<sup>28</sup>Below we will see that Duffy and Ochs (1999) and Duffy (2001) in effect 'force'  $\beta = 0.9$  on their live subjects by applying a 10% probability that the 'game' will finish at the end of the current round.

To the extent we want to compare [Marimon et al. \(1990\)](#) and [Kiyotaki and Wright \(1989\)](#), we can determine the values of  $\beta$  implied by inequalities 5.1 and 5.2 by applying the former's parameters.

If we apply the above Fundamental Equilibrium parameters to Inequality 5.1 above we find that  $\beta < 1.140$ . Given that we would expect  $0 < \beta < 1$ , this condition will hold, i.e., the above parameters appear consistent with a Fundamental Equilibrium.

However, if we apply [Marimon et al.'s \(1990\)](#) Speculative Equilibrium parameters to Inequality 5.2 above, we find  $\beta > 0.275$ . This is more troublesome because we cannot say for sure what  $\beta$  ought to be. This requirement leads [Marimon et al. \(1990\)](#) to fudge the issue somewhat by stating that “for high enough discount factors the unique stationary equilibrium of Kiyotaki and Wright's economy is the so-called speculative equilibrium...” (p. 359).

In addition to the comparability problem that  $\beta$  points to, there are at least two more criticisms of [Marimon et al.'s \(1990\)](#) simulations: (i) the authors did not explain why these two specific sets of parameters were chosen (other than attempting to replicate inequalities 5.1 and 5.2); and, related, (ii) they did not discuss results that used different parts of the parameter space<sup>29</sup>.

Because the authors did not address these issues, they are at risk of being accused of ‘cherry picking’ parts of the parameter space. Unfortunately, we require an analysis of results from different parts of the parameter space to know the extent to which this criticism is important!

## Results and Discussion

This potential criticism of cherry-picking aside, [Marimon et al. \(1990\)](#) report and discuss various simulations of these models (A and B).

When the Fundamental Equilibrium parameters were applied to Model A and the agents had a full set of classifiers available to them, “the exchange and consumption strategies implemented by the system of winning classifiers virtually coincide with the [fundamental] equilibrium strategies.” (p. 356), i.e., there is successful convergence to this Nash equilibrium.

Furthermore, when the agents have a restricted set of classifiers (along with a genetic algorithm as described above) under the same conditions, the results were about the same but “it takes longer to converge...” ([Marimon et al., 1990](#), p. 359).

---

<sup>29</sup>It is possible that the authors covered these issues in the working paper that preceded the published paper (the former is referred to in the latter) but this appears no longer available. An attempt was made to retrieve this from the publishers (the Hoover Institution) but they could not find it.

Turning to the Speculative Equilibrium parameters (now  $u_i = 500$ ), the authors report that the results essentially reflect the Fundamental Equilibrium and not the Speculative Equilibrium they had targetted. This was true when the agents used full and restricted sets of classifiers.

It appears that the agents fail to learn a combination of strategies via their classifiers required for the Speculative Equilibrium to emerge. For this to happen, the agents have to learn to: (i) exchange Resource 2 for 3; (ii) hold Resource 3 between rounds (i.e., not consume it and incur a relatively high storage cost); then (iii) exchange Resource 3 for 1 (potentially); and then (iv) consume Resource 1. After this sequence of events, the utility consumed (500) would feed back to the classifiers, as a type of reinforcement learning, that allowed this consumption to occur.

From [Marimon et al.'s \(1990\)](#) data and comments, it appears the Type 1 agents will do step (i) above but the agents learn to consume their Resource 3. This seems odd but it allows them instead to produce and hold Resource 2 (with a lower storage cost) between rounds, which they might then trade. The net effect is that the pattern of resource holdings closely resembles the [Kiyotaki and Wright's \(1989\)](#) Fundamental Equilibrium<sup>30</sup>.

It is possible that the seemingly large difference in storage costs ( $c_2 = 1.0$  and  $c_3 = 20.0$ ) plays a role here and we will see below that the models of both [Başçi \(1999\)](#) and [Staudinger \(1998\)](#) reduce this differential.

One of the noteworthy results identified by [Marimon et al. \(1990\)](#) is referred to in a subsection entitled 'Patience requires experience' (pp. 361-2). They observe the problem of lock-in when it comes to winning classifiers that dominate the agents' mental models. If an agent happens upon beneficial but sub-optimal classifiers, these are often retained for the life of the simulation. This is like climbing to the top of a local peak in a fitness landscape when higher peaks exist elsewhere.

Put another way, the agents did not experiment sufficiently enough to find an optimal strategy of exchange.

[Marimon et al. \(1990\)](#) also run simulations corresponding to [Kiyotaki and Wright's \(1989\)](#) Model B. This model also has two equilibria but this time they co-exist across the whole parameter space.

When the agents use a full set of classifiers, [Marimon et al. \(1990\)](#) observe trading patterns consistent with one of the equilibria but eventually the economy gravitates to

---

<sup>30</sup>It is not clear why the exchange and consumption classifiers had to be linked here. An alternative design could have simply used the utilities gained from consumption (and costs from storage) as feedback to both sets of classifiers but independently of each other. Alternatively, the consumption classifiers could have been assumed away by simply granting the agents the intelligence of consuming their own preferred resource and neither of the other two. Our interest lies more with exchange so the consumption classifiers seem like a distraction.

the other equilibrium. When a restricted set of classifiers is used, the economy does not converge on an equilibrium after 2,000 rounds but “the economy seems to be moving towards [one of the equilibria].” (p. 366).

We conclude by quoting [Hodgson and Knudsen \(2004\)](#): the results of [Marimon et al. \(1990\)](#) are “qualified and partially inconclusive.” ([Hodgson and Knudsen, 2004](#), p. 21).

### Other Evolutionary Models

Here we briefly note two papers. The first is [Başçi \(1999\)](#) who also uses classifier systems: the question explored is whether reducing  $c_3 - c_2$  and introducing imitation would induce Type 1 agents to adopt a speculative strategy. [Başçi \(1999\)](#) finds that “neither modification by itself results in a significant increase in the speed of convergence to the speculative strategy profile ... but that the combination of these two modifications does enable the classifier system to achieve convergence .. with a high frequency.” ([Duffy, 2001](#), p. 302). Note, however, that the RHS of Inequality 5.2 ( $(\sqrt{2} - 1)\frac{\beta}{3}u_1$ ) must be significantly greater than the RHS ( $c_3 - c_2$ ) for this to be true, in addition to agents imitating others.

The second paper is [Staudinger \(1998\)](#), which also adjusts  $c_3 - c_2$  but here the author applies a genetic algorithm to exchange strategies only. [Staudinger](#) finds that “if the difference between benefit and costs is sufficiently big the economy converges to the speculative equilibrium.” (p. 98). This is essentially the same result as [Başçi \(1999\)](#) but without the use of classifier systems and imitation.

We can think of [Başçi \(1999\)](#) and [Staudinger \(1998\)](#) as exploring specific parts of the parameter space, which [Marimon et al. \(1990\)](#) failed to do.

### 5.4.3 Empirical Studies of Monetary Emergence

[Brown \(1996\)](#) was the first to test the [Kiyotaki and Wright \(1989\)](#) model on live human subjects; however, here we discuss [Duffy and Ochs’s \(1999\)](#) experiments because these were more rigorous and applied more broadly than [Brown’s](#) (and their results included the latter’s main result).

[Duffy and Ochs \(1999\)](#) attempted to create a set of experiments that was as close to [Kiyotaki and Wright’s \(1989\)](#) environment as possible. A detailed description can be found in their Section III ([Duffy and Ochs, 1999](#), pp. 853-857).

In fact, [Duffy and Ochs](#) even went as far as to ‘force’ a discount rate ( $\beta = 0.9$ ) by telling their human subjects there was a 10% probability that the ‘game’ would end before the next round.

The main results from [Duffy and Ochs \(1999\)](#) are presented in Table 5.1 below.

	$u_i = 20$		$u_i = 100$		
	KW Exp	D&O(p)	KW Exp	D&O(p)	D&O(e)
Type 1	0	0.30	1	0.36	0.37
Type 2	1	0.98	1	0.95	0.96
Type 3	0	0.07	0	0.25	0.16

Table 5.1: A Summary of Duffy & Ochs' Experiment Results (Model A). Two sets of results are presented: one where  $u_i = 20$  (Table 4, p. 859) and another where  $u_i = 100$  (Table 7, p. 865). Other parameters are identical for both:  $\beta = 0.9$ ,  $c_1 = 1$ ,  $c_2 = 4$ , and  $c_3 = 9$ . The first experiment corresponds to [Kiyotaki and Wright's \(1989\)](#) Fundamental Equilibrium and the second to their Speculative Equilibrium. Values of 1 in this table mean that Type  $i$  agents would always exchange Resource  $(i+1) \bmod 3$  for  $(i+2) \bmod 3$  (0 means they would not). The results referred to as "KW Exp" are the values expected in [Kiyotaki and Wright's \(1989\)](#) Nash equilibria for each agent type. Those referred to as "D&O" refer to the results observed in [Duffy and Ochs's \(1999\)](#) experiments (second half of results only): "D&O(p)" refers to these results when the subjects start each game holding their production goods, and "D&O(e)" when they start with the good expected in [Kiyotaki and Wright's \(1989\)](#) Nash equilibria. There are two main differences: (i) Type 1 agents do not behave as expected in [Kiyotaki and Wright's \(1989\)](#) Nash equilibria (in both experiments) - notably, their behaviour does not appear to change between the two experiments; and (ii) Type 3 agents do not behave as expected when  $u_i = 100$ .

Like [Brown \(1996\)](#) before them, [Duffy and Ochs \(1999\)](#) observe a "Type 1 problem" in that the behaviour of Type 1 agents does not accord with [Kiyotaki and Wright's \(1989\)](#) Nash equilibria. In the first set of experiments, when  $u_i = 20$ , we expect Type 1 agents never to exchange Resource 2 for Resource 3 and yet they do so 30% of the time<sup>31</sup>.

Likewise there is inconsistency in the Speculative Equilibrium setting ( $u_i = 100$ ). Here [Duffy and Ochs \(1999\)](#) presents two sets of data: one where the subjects start the experiments with the resource they produce ("D&O (p)")<sup>32</sup> and a second where they start off as if in [Kiyotaki and Wright's \(1989\)](#) Speculative Equilibrium ("D&O (p)")<sup>33</sup>.

Type 1 agents exchange resources about 37% of the time under both initial conditions, versus 100% in [Kiyotaki and Wright \(1989\)](#). Duffy and Ochs state that a statistical significance test suggests Type 1 agents did not change their behaviour from when  $u_i = 20$  and  $u_i = 100$ .

[Duffy and Ochs \(1999\)](#) also observe an inconsistency for Type 3 agents, when  $u_i = 100$ . These agents exchange Resource 1 for 2 16-25% of the time (depending on initial conditions) when in [Kiyotaki and Wright's \(1989\)](#) Speculative Equilibrium they are expected never to do so. Statistical significance tests indicate that the 0.25 result is significantly different from zero and that 0.16 is not significantly different from 0.25.

<sup>31</sup>When the data is disaggregated and analysed, Duffy and Ochs find that approximately 48% of agents behave according to the Fundamental Equilibrium and the remaining agents are roughly evenly split in exchanging Resource 2 for 3 between 5% and 100% of the time.

<sup>32</sup>See [Duffy and Ochs \(1999\)](#), Table 4, p. 859.

<sup>33</sup>ibid, Table 7, p. 865.

Duffy (2001) attempts to reconcile the empirical evidence observed in Duffy and Ochs (1999) with an ABM (and to conduct further live experiments).

After making a number of (reasonable) simplifications, Duffy's (2001) computational model boils down to a single probability for each agent type as to whether it would exchange Resource  $(i+1) \bmod 3$  for  $(i+2) \bmod 3$ . To implement this, the model recorded the number of times the agents chose to exchange these resources and then whether they were subsequently able to trade for their consumption resource  $(i)$  or not. The model recorded the same data for when the agents chose not to exchange these resources.

These data were then transformed into a probability that was used to determine the agents' decisions (see pp. 304-305 of Duffy, 2001 for further details). The updating of probabilities was referred to by Duffy as 'hypothetical reinforcement' and was equivalent but not identical to Roth and Erev's (1995) strict reinforcement learning.

Duffy's (2001) main results are shown in Table 5.2 below.

	u = 100			
	KW Exp	D&O(p)	D&O(e)	D Sims
Type 1	1	0.36	0.37	0.32
Type 2	1	0.95	0.96	0.99
Type 3	0	0.25	0.16	0.04

Table 5.2: A Summary of Duffy's (2001) Experiment Results. The parameters are identical to Duffy and Ochs's (1999) Speculative Equilibrium experiments (RHS of Table 5.1 above). The results referred to as "KW Exp" are the values expected in Kiyotaki and Wright's (1989) Nash equilibria for each agent type. Those referred to as "D&O" refer to the results observed in Duffy and Ochs's (1999) experiments (stated in Table 5.1 above). "D Sims" refers to Duffy's (2001) main simulation results.

Duffy (2001) concludes that, overall, "qualitatively, the 'fit' of the artificial agent simulation statistics to those ... from the experimental data appears to be quite good." (p. 309). The largest difference between Duffy and Ochs's (1999) data and Duffy's (2001) simulation results is for Type 3 agents: Duffy seems comfortable that this is not significant.

#### 5.4.4 Comparing the Literature

When we consider the computational models of Marimon et al. (1990), Başı (1999), and Staudinger (1998), and the empirical results of Brown (1996) and Duffy and Ochs (1999), we are in the curious position of being able to compare game theoretic and computational analyses with empirical data (all largely based on the same model).

The main observation we can make is that of inconsistency between the three types of study. Moreover, given the emphasis on realism in Chapter 2, it seems reasonable to use the empirical evidence as a yardstick by which to judge the other two.

On that basis, we can state that, on the whole, [Kiyotaki and Wright's \(1989\)](#) Nash equilibria appear to be a first-order approximation of [Brown's \(1996\)](#) and [Duffy and Ochs's \(1999\)](#) empirical evidence. This is certainly true of type 2 and 3 agents in the Fundamental Equilibrium experiments (LHS of [Table 5.1](#)) but it is less true of Type 1 agents in these experiments. Furthermore, there appears to be a significant difference between both Type 1 and Type 3 agents in the Speculative Equilibrium experiments relative to the equilibrium behaviour expected in [Kiyotaki and Wright \(1989\)](#).

The idea that [Kiyotaki and Wright's \(1989\)](#) Nash equilibria are first order approximations (inaccurate but helpful) is very much consistent with the theme of the last section.

What about the simulation results of [Marimon et al. \(1990\)](#), [Başçi \(1999\)](#), and [Staudinger \(1998\)](#) versus the empirical evidence?

[Marimon et al. \(1990\)](#) predates the empirical tests so these did not influence their ABM. The key result of this paper is that their Type 1 agents appear to play the Fundamental Equilibrium strategies in both the fundamental and speculative environments. [Duffy and Ochs \(1999\)](#) also show consistency between the two but only partially.

[Başçi \(1999\)](#), and [Staudinger \(1998\)](#) were published after [Brown \(1996\)](#) but [Başçi \(1999\)](#) saw [Duffy and Ochs's \(1999\)](#) results before they were published.

Unfortunately, the data published in [Başçi \(1999\)](#), and [Staudinger \(1998\)](#) make it impossible to analyse the differences between their results and the empirical evidence in detail.

[Başçi \(1999\)](#) states that “our algorithm qualitatively mimics the experimental results of [Brown \(1996\)](#) and [Duffy and Ochs \(1996\)](#).” (p. 1582). This is a vague statement and is the only reference to a comparison with empirical evidence in the paper.

[Staudinger \(1998\)](#) states that a “large number of simulations shows that individuals of type 1 prefer to hold the good with the lowest costs. These results are consistent with those of [Brown](#) and [Marimon et. al.](#)” (p. 98). While this appears broadly true, data is not presented to allow us to judge this quantitatively.

It is worth highlighting that [Başçi \(1999\)](#) and [Staudinger \(1998\)](#) share an interesting feature: both appear more interested in replicating [Kiyotaki and Wright's \(1989\)](#) Nash equilibria than reproducing the empirical evidence. Both papers refer to these Nash equilibria as “optimal”. For example, [Staudinger \(1998\)](#) states that the “main questions dealt [sic] with are: Do the economy converge to a steady state equilibrium, and if so, do individuals with bounded rationality learn the optimal strategies of this equilibrium[?]”.

Furthermore, neither of these papers deal with the question of uncertainty for the subjects who took part in [Brown's \(1996\)](#) and [Duffy and Ochs's \(1999\)](#) experiments, and its implications for decision making. Notably, [Kiyotaki and Wright's \(1989\)](#) Nash equilibria



require not only that each agent is substantively rational, they also require that agents assume all other agents are also substantively rational<sup>34</sup>. It would be entirely reasonable for subjects in experiments not to know how other subjects make their decisions, which would have to be based on guesswork about how other agents think about the problem (and how these agents are guessing how others are thinking...).

As discussed above, [Duffy \(2001\)](#) also developed a computational model but his was more focused on the empirical evidence.

## 5.5 Concluding Comments

From the discussions in the preceding four sections of this chapter, three themes in particular are worth highlighting. These provide challenges for the models of organic institutional emergence developed below.

### 5.5.1 Reasoning, Learning and Habits

The models evaluated or mentioned in this chapter all include combinations of forms of reasoning, learning, and habit formation. However, none of these models included all three (at least not explicitly), all of which seem to be important and interesting cognitive phenomena.

One of the challenges this poses is to develop agents' mental models that allow them to reason in some way about their environment, to learn, and to form habits.

The EMIL models evaluated in Section 5.2 are orientated around reasoning and (reinforcement) learning, and are focused on emergence. If the convergence of some cognitive rule to a probability of 1 can be interpreted as a habit then one could argue that these models also included habit formation; however, this was not explicitly mentioned.

Overgeneralizing somewhat, [Hodgson and Knudsen's \(2004\)](#) model can be seen as the mirror image of the EMIL models. Agents' decisions are driven entirely by innate tendencies applied to observed data and a single habituation variable, which simply reflected previous decisions. Changes in the habituation variable were seen by [Hodgson and Knudsen \(2004\)](#) as akin to but not the same as reinforcement learning.

The challenge here is related to the following quote from [Hodgson and Knudsen \(2004\)](#) who state that "more complex learning algorithms would clearly be appropriate in decision environments involving more learning parameters and behavioral choices than are present in our model." (p. 23, Footnote 7).

---

<sup>34</sup>This assumption is necessary when agents decide whether or not to hold a particular resource between rounds: what other agents do between rounds matters for this decision.

## 5.5.2 Beyond First-Order Approximations

The main theme of Section 5.3 was that game theoretic approaches can be viewed as helpful approximations of reality, and this was neatly exemplified in Section 5.4 when we contrasted Kiyotaki and Wright's (1989) Nash equilibria with comparable empirical evidence. However, if our aim is "*explaining the causes of behaviour*" (Hodgson, 2012, p. 99, emphasis included), or at least to move a step closer to this, then we have to look beyond approaches that combine utility maximization and substantive rationality.

This fits with the emphasis on realism and pragmatism in this thesis and also with Ullmann-Margalit's (1978) normalcy condition for invisible hand explanations.

Another feature of a number of models analysed above (notably that of Hodgson and Knudsen, 2004 and the four EMIL ABMs) was the grounding of models in lifelike situations. This stands in contrast to the abstract nature of many 'games' in game theory and it sits comfortably with pragmatism.

Furthermore, we would ideally develop models of both convention and social norm emergence, corresponding to Schultz's (2001) two categories of interaction.

With all these aims in mind, the models developed for this thesis are grounded in the important economic challenges of markets, property rights, and the division of labour. Mental models are designed to help agents make decisions in uncertain environments. More specifically, the emergence of a market is based on a coordination-like challenge (agents finding each other to trade); and the emergence of property rights is based on interactions that include a free-rider problem (the temptation to steal others' resources which runs counter to the socially preferred outcome where they do not).

## 5.5.3 Thoroughness of Analysis

This final theme relates to criticisms of the write-ups of the four EMIL models and of Marimon et al.'s (1990) ABM. The exploration of the parameter space was poor in all of these cases, and there were no explanations for why specified parameters were chosen over others.

As mentioned in Section 5.2.3, it is important in computational research to report on simulation results for areas outside of some chosen point in the parameter space. Ideally data would be provided (like those in figures 5.1 and 5.2 above, which were copied from Hodgson and Knudsen, 2004) in addition to analysis.

The aim in doing this is not only to understand when some phenomenon occurs but also to understand the nature and degree of the impact of parameters (or different methods) on the phenomenon. In the models described in chapters 7 and 9 below, exploring the parameter space is also important for understanding when an environment is enabling

of (conducive to) an institution emerging, and when it is not. Exploring and analysing the parameter space is essential in computational research.



# Chapter 6

## Models: Rationale, Design and Results

All models are wrong, some are useful.

– George Box

This chapter has four aims:

1. the identification of two research questions;
2. the reasons why specific architectural decisions were made when designing the two computational models presented in later chapters;
3. an overview of these models; and
4. a summary of the main results of the simulations.

It is perhaps surprising that a summary of the main results is presented before the models and simulations are discussed in detail in later chapters. There are two reasons for this.

The first stems from the fact the models were not developed in order to investigate a pre-conceived hypothesis; rather, they were used as *tools of investigation*. This means that it would be disingenuous to state some hypothesis in this chapter and then to ‘test’ it subsequently with some Agent-Based Models (ABMs). The two research questions below provide us with orientation but the main results are presented in lieu of some theory or hypothesis to be tested.

In practical terms, the research proceeded iteratively between conceptualising and modelling. Some of the open questions in the spontaneous order and Institutional Economics

literatures acted as ‘compass bearings’ (including earlier versions of the research questions) but the research phase should be thought of more as an exploration of terrain.

In fact, the need for this approach was augmented by the focus on ‘surprising’ emergent phenomena. There is a tension in trying to understand and model such processes: the ‘surprise’ element means they should be beyond our (unassisted) cognitive horizons; but we need to understand them sufficiently enough to develop an initial model. The solution to this paradox was to employ a feedback loop: to develop a model, investigate and analyse the resulting simulation data, and then to make adjustments to the model. The process can perhaps best be described as ‘trial and error’. Moreover, we can think of the computational models as *cognitive annexes* that help overcome the cognitive horizon problem.

The authors of the EMIL Project Report stated something similar to this: their models were used both to develop the theory and to validate it (Conte and Edmunds, 2010, p. 5).

The second reason for discussing the main results in this chapter is that the models described in chapters 7 and 9 are relatively large (the reason for which are discussed below) and summarising the results helps with orientation and understanding.

The rest of this chapter is divided in to four sections which are aligned with the four aims listed above.

## 6.1 Research Questions

The two research questions stated below arise from the ontology of Complexity Economics (CE) set out in Chapter 2; the spontaneous order literature discussed in chapters 3 and 4; and the models / model types evaluated in Chapter 5.

Let us first state the research questions and then discuss them:

1. Can organic institutions emerge spontaneously across a population while also im-merging within individuals’ mental models via reasoning, learning, and habitua-tion? and
2. Can ‘liberal legislation’ catalyse institutional emergence when it does not occur endogenously?

To better understand these questions, it is helpful to pull them apart and address various topics.

## Organic Institutional Emergence

In this thesis we are interested in a specific type of spontaneous order: organic (unplanned) institutions. For clarity, it is helpful to repeat Ferguson's statement, mentioned in the Introduction, of order that is "the result of human action, but not the execution of any human design" (Ferguson, 1767, p. 205).

The main challenge here stems from the lack of understanding of mechanisms by which organic institutions emerge. As was mentioned previously, Hayek's framing of spontaneous order is viewed by many (e.g., Gray, 1998; and Luban, 2020) as the most sophisticated but we saw in Chapter 4 that his theory of cultural evolution has been heavily criticised<sup>1</sup>.

This lack of a coherent theory of organic institutional emergence chimes with the following quote from Hodgson (2002a).

At both the theoretical and methodological level, there is no clear consensus among modern researchers as to what would constitute an adequate or acceptable explanation of the process of emergence of an institution. This question is at present under-researched. (Hodgson, 2002a, p. 112)

Field wrote something similar:

While we are making methodological and substantive progress on macro-level explorations of the consequences of institutional variation, attempts to provide a satisfactory microanalytics, one that explains systematically how institutions are created, how they are sustained, and why they vary, have been far less successful. (Field, 2007, p. 1)

There has also been a lot of work in game theory which claims to be concerned with organic institutional emergence. However, we argued in Chapter 5 that while these models contain helpful 'first-order approximations' (Arthur, 2013), they have mostly failed to identify the *mechanisms* by which organic institutions emerge in the real world. Field's (1979, 1981, 1984) criticisms and Hodgson's (2012) discussion of the limits of rational choice theory were used to support this point<sup>2</sup>.

---

<sup>1</sup>In previous chapters, we saw that Hayek is not the only theorist who has been concerned with spontaneous order. In Chapter 3 we discussed a wide range of research pre-, intra-, and post-Enlightenment. Nonetheless, it is not unreasonable to state that Hayek's framing has been at the heart of the spontaneous order literature for decades.

<sup>2</sup>As mentioned in Chapter 5, this is not to argue that game theory is completely useless. The discussion of Kiyotaki and Wright's (1989) money emergence model, for example, showed that their Nash equilibria are interesting but, ultimately, not fully consistent with empirical evidence.

### **Can These Institutions Emerge?**

A question worth repeating, discussed in Chapter 4, is whether or not we should expect beneficial institutions to necessarily emerge.

There are two parts to this. From a complexity science point of view, whether or not some property emerges depends on whether the environment is sufficiently enabling.

Second, might good and bad institutions emerge under different conditions?

Hence, the first research question asks ‘can’ institutions emerge? And, similar to Hayek’s definition of order, the question is agnostic as to whether any emergent institution is beneficial or detrimental. Note, however, that in the first model we start from a state of non-coordination and explore the agents unintentionally coordinating: no ‘bad’ institutions can emerge in these simulations; but this is not true of the second model.

### **Emergence and Immergence**

As mentioned previously, [Hodgson \(2006a\)](#) uses the analogy of a Klein bottle to illustrate the idea that, for institutions, “the subjective ‘inside’ is simultaneously the objective ‘outside’” (p. 8). From a complexity science perspective this image seems appealing because it highlights both some externally observed structure co-existing with institutions within agents’ mental models. An example can be seen in [Hodgson and Knudsen’s \(2004\)](#) simulations when the agents choose to drive either on the left or right of the road: the objective emergence is the observed phenomenon of all the agents driving on the same side whereas the inner, cognitive phenomenon is habituation.

Moreover, the idea of immergence appears to have been catalysed by [Conte and Castelfranchi’s \(1995a\)](#) reference to cognitive emergence; and further developed in [Castelfranchi \(1998\)](#), the EMIL Project, and subsequent work, e.g., [Conte et al \(2013\)](#).

An important part of the first research question, then, is to explore further this combined emergence and immergence as the process counterparts to Hodgson’s Klein bottle analogy.

### **Reasoning, Learning and Habituation**

While emergence and immergence are of general interest, we are more interested in specific mechanisms by which these occur.

As stated in the conclusion to the previous chapter, the models critiqued in that chapter seem to contain three important features (deployed in different ways): (i) types of reasoning under conditions of uncertainty; (ii) learning, notably that of reinforcement learning ([Roth and Erev, 1995](#); and [Erev and Roth, 1998](#)); and (iii) habituation.



The EMIL Project models probably come closest to including all three of these features but were weak on habituation (and simulation evaluation); and [Hodgson and Knudsen \(2004\)](#), which focused on habits, commented that forms of learning could be added to their model.

The first research question includes the challenge of developing computational models that accommodate all three of these phenomena in a coherent and compelling way. Furthermore, the design of the agents' mental models also accounts for the idea of institutions as rules<sup>3</sup> and also the idea of different rules competing within the mind. This latter point is consistent with Hayek's theory of mind, [Holland's \(1975\)](#) classifier systems, and [Holland et al's \(1986\)](#) discussion of how agents identify patterns and reason under conditions of uncertainty.

### Liberal Legislation

Turning now to the second research question, there was a discussion in Chapter 4 of what Hayek refers to as 'liberal legislation' in some of his work ([Vanberg, 1994b](#)). This is the idea of legislation that enables efficient markets (but which is not designed to bring about specific outcomes) and is contrasted with 'interventions' in economic activity that have concrete objectives in mind.

Related to this, let us recall a point made in [Ullmann-Margalit \(1978\)](#) about her evolutionary-functional mould of invisible hand explanations. This is when some process or feature can *maintain* something which has been deliberately designed, rather than it having emerged spontaneously. This means the mechanism that sustains 'order' can be categorized as 'invisible hand' even though the outcome has been planned.

With these points in mind, a set of experiments is designed (based on the second model) whereby legal rules are imposed on a population of agents in conditions when we know property rights do not emerge endogenously. The question we explore is whether such rules enable the emergence of property rights in the population.

### Computational Models

In addition to the conceptual material considered above, an implicit aim of the research questions is to add to the body of computational models concerned with institutional matters.

Most notably, and as mentioned previously, [Gräbner](#) has argued in three recent papers ([Gräbner and Kapeller, 2015](#); [Gräbner, 2016](#); and [Gräbner, 2018](#)) that agent-based models appear well suited to institutional analysis. However, the only model that focuses specifically on organic institutional analysis cited in those papers is that of [Hodgson and](#)

---

<sup>3</sup>that in "circumstances X do Y." ([Hodgson, 2006a](#), p. 3).

Knudsen (2004), which was discussed in detail in Chapter 5. The models presented below add to this strand of literature.

Now that we have stated and discussed the research questions, let us consider some of the factors that influenced the design of the models over and above these questions.

## 6.2 Model Design

Here we consider the main influences on the design of the models.

Recall from the Introduction various challenges in the literature identified in Hodgson's (2002a) *Agenda*. These all influenced the design of the models and will not be repeated here.

Below we look at five themes before discussing three risks of agent-based modelling, as identified in Gräbner (2016), and how these risks were managed in the models.

### Spontaneous Order

Consistent with the ontology of CE described in Chapter 2, we design our agents with mental models that allow them to make decisions, given their aims, under conditions of uncertainty. The agents in these models are deliberately myopic with no means to directly coordinate or cooperate.

Furthermore, setting aside the 'liberal legislation' experiments for now, the models were developed in a way that meant there was *space* for some organic institution to emerge but care was taken to ensure nothing was designed into the models that created a bias towards such emergence.

In the first model, the agents' mental models were designed to allow them to find other agents to trade with: no attempt was made to force or encourage the creation of a market. In the second model, agents were endowed with two *propensities* (to steal and to defend their resource holdings) that evolved via reinforcement learning. Again, there was no attempt to create an outcome one way or the other, including property rights.

### Bottom-Up Pragmatism and Realism

In Chapter 5 we discussed a number of models that are based on representations of real-world situations, e.g., Hodgson and Knudsen's (2004) traffic convention model and the four models discussed in the EMIL Project Report. By contrast, game theoretic models tend to abstract away from practical situations.

The models developed for this thesis follow the pragmatic, ground-up approach, which fits with the CE principles discussed in Chapter 2. The main rationale for this links to

Hodgson's (2012) emphasis on seeking to identify mechanisms that explain phenomena seen in the real world.

More specifically, the models are focused on economic matters. To that end, three of the core features of market-based economies were chosen as 'target domains': property rights, markets, and the division of labour.

The first model below assumes the agents respect each other's property, allowing us to focus on market emergence and agent specialisation (leading to a division of labour). The second model is based on the first but it relaxes the property rights assumption (agents are now free to steal from each other): the question is whether property rights emerge endogenously which might then enable market emergence and the division of labour.

We should be clear that a 'market' here is defined in a narrow sense: it is a solution to a coordination problem in which agents want to find others to trade with. Markets can be defined in many different ways, e.g., Hodgson (2019) discusses various meanings of this word in the context of taxonomic definitions.

### Invisible Hand Explanations

The design of the models is also linked to Ullmann-Margalit (1978). This paper is helpful in two ways. First, as discussed in Chapter 4, Ullmann-Margalit proposed two 'moulds' of invisible hand explanations: the aggregate and functional-evolutionary moulds.

These two moulds correspond neatly with the two models developed below: the first model with the aggregate mould and the second with the functional-evolutionary mould. In addition, there is an approximate correspondence between the models and Schultz's (2001) coordination situations (the first model) and collective action situations (the second model): free-rider problems were not observed in the simulations based on the first model<sup>4</sup> but they were seen in the second.

The second way Ullmann-Margalit's (1978) paper is helpful is that it proposes three characteristics (or 'constraints') for invisible hand explanations. These should: (i) be individual-based; (ii) meet a 'normalcy' condition; and (iii) be surprising in some way.

The first of these characteristics was designed into the models, which is typical of ABMs anyway. The second and third conditions both helpfully guided the research during the iteration between models and theory. Also, we use these characteristics to evaluate any institutions that might emerge (see Section 6.4.6 below).

---

<sup>4</sup>We found the problem was equivalent to a non-pure coordination game rather than a pure one.

### Principles of Complexity Economics

Given the emphasis on the eleven CE principles in this thesis (Section 2.3.3), it would be appropriate to mention how the models were consistent with these features. The models:

1. are *computational* in nature;
2. have no predisposition toward *equilibrium* or dis-equilibrium conditions;
3. create a space in which institutions might or might not emerge / *form*;
4. allow for a *stratified ontology* whereby social patterns (including respect for property, a market, and division of labour) might or might not emerge from the agents' mental models;
5. allow for the *evolution* and co-adaptation of the agents' mental models and skills;
6. are open to *non-ergodicity* in that certain structures within the system might emerge or change;
7. include *uncertainty* in that agents lack information, have limited cognition, and must handle mutual contingency;
8. include agents that reason, learn, and form habits within their *mental models*;
9. are inspired by various *disciplines*, including economics, the complexity sciences, spontaneous order, psychology, and sociology;
10. are designed with realism in mind, and they focus on key features observed in free market systems, namely respect for property, markets, and the division of labour; and
11. are based on, from an individual agent's point of view, grounded problems.

### Other Factors

The division of labour was mentioned in the previous section. This is of secondary importance to us because it is not related to the research questions but it is not unimportant: we find that it occurs only after (is enabled by) the emergence of a market. The concomitant specialisation of the agents raises productivity levels in the pseudo-economy.

Moreover, there is a link here between Hayek's emphasis on dispersed knowledge and markets as institutions, which is discussed in more detail in Section 6.4 below.

In terms of data and parameters, the models were designed to ensure that a considerable amount of information was produced for each run (and each set of runs), and so that

it was relatively easy to explore the parameter space. Some of the models discussed in Chapter 5 contained little data and their analysis of results was weak (notably the EMIL Project models), and/or they failed to explore the parameter space sufficiently well, e.g., [Marimon et al. \(1990\)](#). Steps were taken to ensure the models generated helpful data, which is analysed in later chapters.

### 6.2.1 Risks with ABMs

[Gräbner \(2016\)](#) describes three risks that are associated with agent-based modelling (these were mentioned briefly in the Introduction). We discuss these, and how our models mitigated them, here.

[Gräbner's \(2016\) first risk](#) is that “ABMs tempt researchers to take a constructionist-instrumental standpoint” (p. 255), which would lead them to “not try to describe ... reality accurately but consider their theories to be mere instruments replicating observed data.” (ibid). This risk was dealt with by adopting a pragmatist stance in the design of the models and having an aim of seeking to understand the mechanisms by which organic institutions emerge in the real world.

The **second risk** is an “[i]mplicit focus on predictive power” (p. 256), in line “with Friedman’s methodological instrumentalism.” (ibid). Prediction was never an aim of the models developed for the thesis. In the simulations reported below we take an interest in the conditions under which organic institutions do and do not emerge but this is not the same as prediction as described by Gräbner.

The **third risk** is that of “[o]verparameterisation and decreased transparency.” (p. 256). We should discuss this in more detail because it is relevant to the models developed for the thesis, which are relatively large. What Gräbner is referring to here is the risk of a researcher “adding variables, processes, and methods until one gets ... the patterns one wishes to explain.” (p. 256).

In framing this issue, we identify a potential tension in modelling, when the models attempt to describe real-world phenomena, between parsimony and description.

This is addressed particularly well (and succinctly!) in [Elsenbrioch and Gilbert \(2014\)](#)<sup>5</sup>. The authors argue that while parsimony is undoubtedly valuable, this does not mean that all valuable models are parsimonious (see [Edmonds, 2004](#) for a discussion of this point<sup>6</sup>). To confuse the two risks falling foul of the association fallacy and it gives rise to another risk: that detailed but useful models are rejected because they are not parsimonious.

---

<sup>5</sup>See Section 10.1, p. 144, entitled “KISS vs KIDS”.

<sup>6</sup>Recall also that Hayek addresses this in [Hayek \[1964\] \(2014\)](#), which was discussed in Chapter 4.

In modelling (and elsewhere), parsimony is associated with the acronym *KISS* (“Keep It Simple Stupid”). By contrast, [Edmonds and Moss \(2005\)](#) argue for the *KIDS* principle (“Keep It Descriptive Stupid”): start with a model that is “as descriptive as given data and evidence as possible. Once the model is understood, it can be simplified if parts are found to be superfluous...” ([Elsenbrioch and Gilbert, 2014](#), p. 144).

The models developed for this thesis were more aligned with *KIDS* than *KISS* for three (related) reasons. First, the emphasis on real world (economic) phenomena requires that we represent them in a model: this is very much aligned with *KIDS*.

Second, there is an argument that, given the CE principles discussed in Chapter 2, we should develop a model of a *whole system* rather than a partial representation of it. This would allow us to explore certain feedback effects which would be otherwise missing. Such models are, however, likely to contain more details (and parameters) than partial models.

Third, we want to address the question of organic institutions emerging within an enabling environment. To explore this thoroughly in a model requires that we develop multiple parameters that can be adjusted to observe the conditions under which an organic institution emerges and when it does not.

Overall, therefore, there were good reasons for designing relatively detailed models, which were different to [Gräbner’s \(2016\)](#) (entirely reasonable) warning of overparameterisation for the sake of obtaining desirable results.

There are, of course, risks in designing a relatively large model. Here we consider two. The first risk concerns the identification of the causal mechanisms we seek: this is more difficult in larger models. The worst outcome is that we misinterpret the results and get the mechanisms wrong.

The second risk is what [Gräbner \(2016\)](#) referred to as transparency: this would be when the causal mechanisms are accurately identified but this is unclear to everyone except the researcher who was immersed in the models and simulations.

We can think of research time as the cost incurred to minimize these risks. This meant designing models that generated data which enabled analysis of the results, which were then re-designed to generate more data when it was not clear what happened in the simulations. This is part of the model iteration process mentioned above. Put another way, a considerable amount of time (about 4 full-time equivalent years) was spent on modelling, which included making sure that the mechanisms described below are an accurate reflection of what was observed in the models.

This discussion of model size is continued in the Conclusion, in light of the simulations results and, most importantly, the identification of a generalised framework of organic institutional emergence, which points to how this simplification might be achieved.

Now that we have discussed some of the factors that influenced the design of the models, let us now summarise them.

### 6.3 Model Description

Detailed descriptions of the two models are included in chapters 7 and 9 below. Here we provide a summary that is sufficient to make sense of the results discussed in the next section.

We begin by describing the first model and assume its ‘default parameter set’<sup>7</sup>: the second model is a variation of this.

A rudimentary economic system was developed that includes 2 resources within an environment and 25 agents (initially) who have the goal of surviving. Each agent has to consume both resources to survive and is subject to a metabolism cost of 1 unit of each resource at the end of each round. Survival requires the agents to maintain positive stocks of both resources in their ‘personal resource arrays’: if either of an agent’s resource stocks fall below zero, it dies and is removed from the simulation (without replacement). The agents can undertake work in the form of foraging for the 2 resources, which they can subsequently trade with other agents.

Each round of the simulations is divided into two phases: (1) foraging; and (2) interaction. Fig. 6.1 below illustrates the main features of each round.

During the foraging phase, agents can visit either of two resource ‘fountains’ and each has 5 time slots in which to attempt to find one resource unit (their ‘foraging strategy’ is a list of which fountains they visit in the 5 slots). Whether they are successful in finding a resource unit in each time slot depends on their skill in foraging and the amount of each resource remaining at the fountain (the fountains are replenished at the beginning of each round and depleted by foraging).

After the foraging phase is complete, the second phase starts in which agents have an opportunity to trade the resources they hold. At the beginning of this phase each agent is placed at different ‘home’ locations in a geographic space (a  $50 \times 50$  torus). The aim of the agents in this phase is to find other agents and to trade when both agents deem it

---

<sup>7</sup>For simplicity, a ‘default’ set of parameters is specified for each of the two models. When a simulation using these parameters is run, this is referred to as the ‘default scenario’. This set is not meant to represent a ‘correct’ collection of parameters - it merely helps us present the model and the simulation results in the simplest way possible. The main parameters are varied when we explore the parameter spaces.

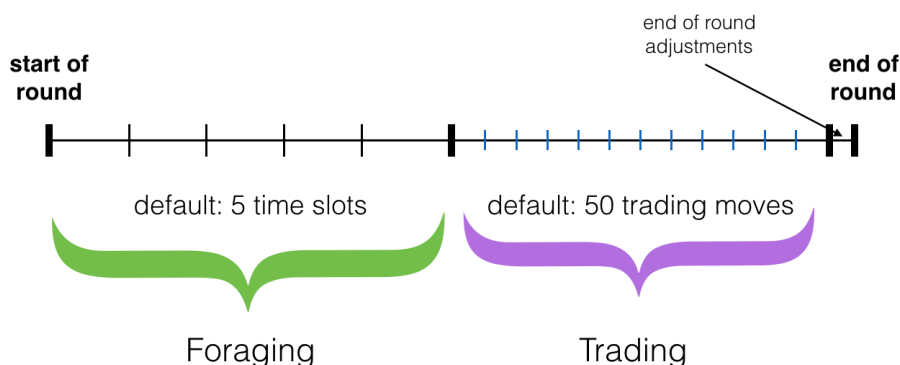


Figure 6.1: A breakdown of each round in the first model. In the first stage of each round, agents forage from resource fountains. In the second stage they try to find each other within a geographic region, in order to trade any resources they might have. At the end of the round, various adjustments are made and some of the agents communicate.

beneficial. Each interaction phase is divided into 50 trading ‘moves’: in each move, an agent can either travel to an adjacent grid square or offer to trade with another agent if one is co-located on its current square.

At the beginning of the simulations, the agents start by moving around in a random walk. After they have met another agent and transacted, they store up memories of transaction locations (but they can also be told of locations by other agents at the end of the round - see below). Agents build up a database of transactions locations (i.e., memories) and these are used by the agents in subsequent rounds to decide where to head to on the torus to find other agents.

Put another way, if an agent has no memories of transactions, they walk around the grid randomly; but if they have memories of transactions, these are used to decide on a target location.

We can think of each agent as having 2,500 potential rules, each corresponding to a grid square: the weight of each rule is determined by the number of transactions at each location known to the agent. These weights decay between rounds as a form of memory loss.

The change in memory weights (due to known transactions and decay) is equivalent to the reinforcement learning seen in the four EMIL models.

If an agent has memories, at the beginning of an interaction phase a ‘roulette wheel’ approach is used to decide which grid square an agent heads towards, i.e., selection is proportional to relative weight. The agents stay at their ‘target’ grid square for the duration of the interaction phase.

At the end of the round (after the 50 trading moves and before a new round begins), a number of processes are completed:



- The agents have an opportunity to communicate with others. In each round there is a 1% probability that any pair of agents will do so: if this happens, the two exchange information about the locations of transactions they were involved in during that round.
- The agents consume whatever resources they have left after foraging and trading, and a metabolism cost of 1 unit of each resource is deducted from each agents' resource stocks. All resources perish in between rounds so the agents cannot carry any forward.
- The agents' foraging skills are updated by the model (simple probabilities that start at 0.5 for each resource, for each agent). These skill changes depend on the number of time slots an agent spent foraging for a particular resource: if an agent spent more time foraging for one than the other, this skill increases, and vice versa.
- The agents then update their 'foraging strategies'. Each has the opportunity to change which fountain it visits in one (randomly chosen) of the five time slots: this depends on which resource the agent is most deficient in (its aim is to increase this in the next round), its foraging skills, and the likelihood of being able to trade in the next round.
- If two agents in the population have resource stocks above a specific threshold, they sire a child. Some resources are deducted from both parents and transferred to the child.

After these processes are completed, the model ticks over to the next round and both resource fountains are replenished. The model is run for 1,000 rounds.

The second model is essentially a variant of the first but, here, respect for property is dropped: the agents can steal from each other. This time the challenge for the agents is about learning to trade or steal, and whether to defend their resources or to acquiesce when another agent attempts to steal from them. If theft is attempted by either agent and neither acquiesces, they both incur a 'cost of fighting', which is determined by a parameter in the model.

Agents are endowed with two propensities: a *propensity to steal* which determines whether an agent attempts to steal from another; and a *propensity to defend* which determines whether an agent defends its resources (if challenged) or acquiesces. In the default simulations, these propensities change via reinforcement learning only: if a choice proves beneficial, the agent is more likely to do it again in the future, and vice versa. Note that in the habituation experiments, the process of habituation also changes these propensities (agents are more likely to repeat past decisions).

The agents' propensities can exceed 1 and fall below zero. At the point of interaction, probabilities (of attempting to steal and defending its resource holdings) are generated by truncating an agent's propensities so they are within the bounds of 0 and 1. These probabilities are then applied in decision making. The rationale for this 'probabilities as truncated propensities' approach is discussed in Chapter 9.

In the default simulations of both models, the agents' mental models change only as a result of reinforcement learning. Additional experiments were conducted to test the impact of habituation on the emergence of organic institutions. These are discussed in Section 6.4.3 below.

## 6.4 Summary of Results

We start with an overview of the main results and then proceed to discuss specific features in more detail.

### 6.4.1 Main Results

In the simulations of both models that use the default parameter sets, organic institutions are seen to emerge and immerge spontaneously. When the parameter space is varied we find that this emergence is conditional on there being an enabling environment (discussed further in Section 6.4.2 below).

In simulations based on the first model we observe:

- Several 'local' (or proto) markets emerge within approximately the first 30 rounds (typically we see 4-7 of them). The agents either 'bump' into others and trade or they hear about transaction locations from other agents.
- A process of symmetry breaking occurs which, in net terms, leads agents in the smaller (by volume) local markets switching allegiance to the larger markets. Eventually a single market dominates. See Fig. 6.2 (described further below) for an illustration of this symmetry breaking in a typical simulation.
- The establishment of a consistent, efficient market leads agents to specialise in foraging for a single resource. This does not happen until agents are confident of being able to trade: it is conditional on market emergence.
- Specialisation leads to foraging skills (as probabilities) for the resources specialised in increasing from 0.5 to 1: productivity increases.

- These productivity gains mean the initial group of agents accumulate resources, sire children, and the population grows until it plateaus at approximately 43 agents<sup>8</sup>.
- When habituation is added to the mental models, we observe an increase in the number of markets: agents are more conservative about their target location choice so symmetry breaking is disrupted. Nonetheless, the results of the default simulations are replicated. There are a number of other interesting results, including habituation eventually dominating the agents' mental models, which are discussed in more detail in Section 6.4.3 below.

Fig. 6.2 below is helpful for visualising the breaking of symmetry (note habituation was not included in the simulation depicted). It shows a time series of the total number of agents visiting each of 7 different proto-markets ( $[x, y]$  denote coordinates on the torus). This run is particularly interesting because it shows two locations dominating (at  $[31, 16]$  and  $[22, 30]$ ) until the latter wins out. It is not shown in Fig. 6.2 but all the agents visit location  $[22, 30]$  only from Round 68.

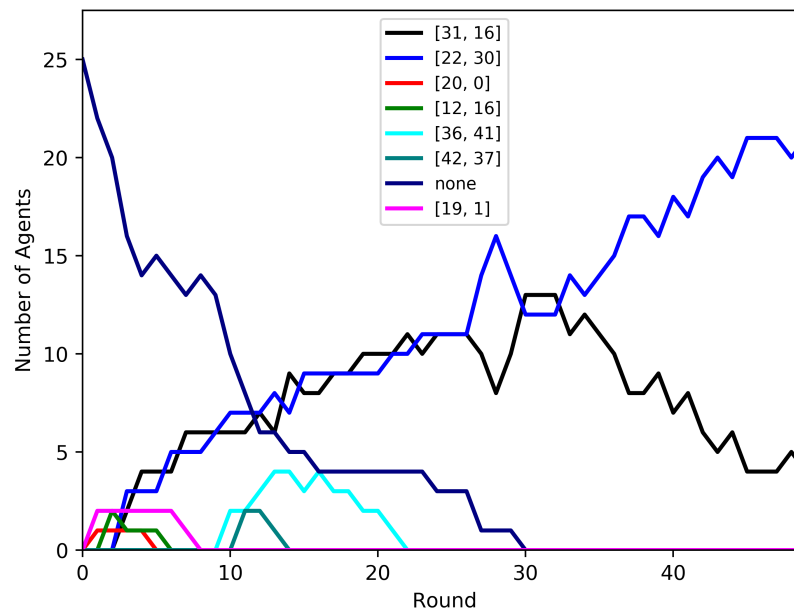


Figure 6.2: A time series of the number of agents targetting different locations on the grid during the first 50 rounds of a typical simulation that uses the default parameter set. This chart demonstrates a scenario when two markets dominate for a period but symmetry breaking means that, eventually, one dominates.

In simulations based on the second model we see:

- The agents quickly learn it is preferable to defend their resources (propensities to defend increase).

<sup>8</sup>This is not to argue that ‘success’ should be measured by total population size: the child-birth algorithm could easily have been changed so agents simply accumulate resources.

- Fig. 6.3 below is helpful for illustrating what happens in the default simulations vis-à-vis agents' propensities to steal (data are taken from a typical simulation). It depicts a 'cloud' of these propensities (measured on the y-axis) and time (rounds) on the x-axis. Each dot represents one agent's propensity to steal at the end of each round. The red dots are those of agents who die before the end of the simulation, the blue dots are of those who survive, and the lines correspond to mean values of each type. The red lines ends when all the agents who die in this simulation have all died.

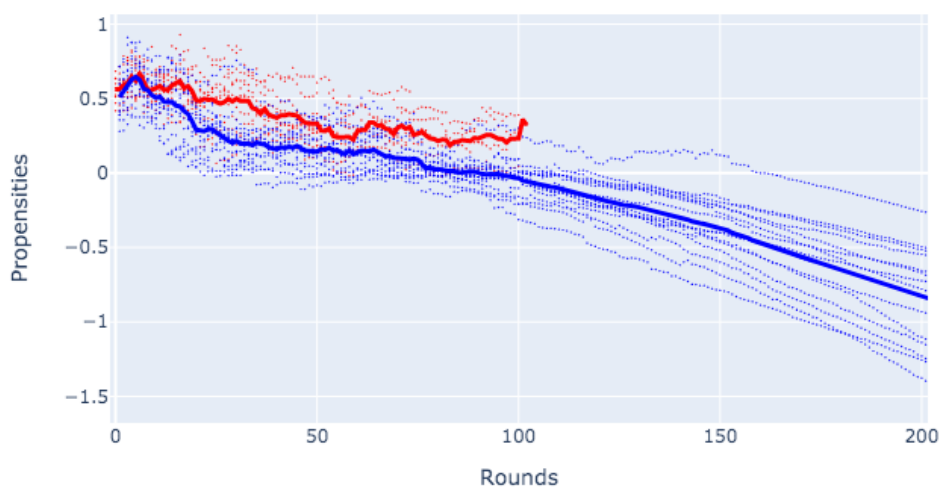


Figure 6.3: The 'cloud' of agents' propensities to steal over 200 rounds in a typical simulation when the default parameter set was used. The chart shows the propensities to steal of the agents who survived until the end of the simulation as blue 'dots'. The blue line is the mean of these values. The red dots represent the propensities to steal of the agents who died before the end of the simulation, and the red line is the mean of these values (it ends in Round 102 when the last of these agents died). The chart shows how, on the whole, the propensities to steal of the agents who died was sustained above those who survived.

- The change in each agent's propensity to steal depends on the 'social environment' it experiences. Six 'patterns' were identified which influence these propensities.
- The first pattern is that when the agents' propensities to defend are below approximately 0.8 on average, the agents learn it is preferable to steal than to trade, i.e., their propensities to steal increase. The benefits of theft on the whole exceed the costs of fighting.
- Second, after the agents' propensities to defend increase above 0.8 on average, the agents learn it is preferable to trade than to steal, i.e., their propensities to steal

decline. This is mainly attributable to the debilitating effect of fight costs, which on the whole outweigh the benefits of theft.

- The third pattern is that when the agents' propensities to steal decline to close to zero, the rate of decline slows down because agents fight less.
- The fourth pattern is the benefit of transactions, which maintains some downward pressure on the agents' propensities to steal.
- Unexpectedly, however, a fifth pattern is observed which encourages propensities to steal higher. The mechanism is described in more detail in Section C.6 of Appendix C. In situations when one agent has a positive propensity to steal and all the other agents have negative propensities, the single agent with a positive propensity benefits disproportionately from a resource concentration effect which is 'centred' around this agent. It is encouraged to steal in this specific social context. Eventually, however, other agents learn from this agent and total fight costs increase again because multiple agents have positive propensities to steal.
- The sixth and final pattern is the locking in of negative propensities to steal when all the agents' propensities are negative. When this occurs, agents only ever trade and they can only benefit from trading, so their propensities can only decline.
- The combination of these six upward and downward patterns and the (eventual) deaths of agents with positive propensities to steal mean that, ultimately, all the surviving agents have propensities to steal below zero. This gets locked in because these agents only ever transact: in this situation we can say that property rights have emerged across the population.
- After property rights have emerged, an efficient market emerges and the agents once again become specialised, i.e., the results of simulations based on the first model are replicated. In a sense, the second model's simulation results show that property rights, which were assumed in the first model, are 'endogenised'.
- The results of habituation and liberal legislation experiments are discussed further below, in sections 6.4.3 and 6.4.5, respectively.

## 6.4.2 Parameter Space Exploration and Enabling Environments

Chapters 8 and 11 below contain summaries of the results of the parameter space explorations for the two models, so this will not be repeated here. In addition, appendices B and C provide much more detailed explanations of these results, respectively.

There are three main points to note.

The first is that the various changes in the parameters have a realistic and understandable impact on the organic institutions. This is because the models are designed to mimic real world situations. For example, when the probability of communication at the end of each round is reduced from 1% to zero, the speed of symmetry breaking slows down in simulations based on the first model. This seems reasonable because communication plays a catalytic role in this process. In extremis, when the agents are prevented from communicating, there is no breaking of symmetry: in a steady state we observe about 9-10 local markets with approximately 4 agents each on average. In these simulations, the ratio of actual to desired transactions is nonetheless high enough for the agents to specialise.

The second main point is that the simulation results demonstrate that organic institutional emergence is conditional on there being an enabling environment. We see this when we start from the default parameters set and relax each of these (or multiple) parameters. For example, when the population is too sparse (the ratio of geographic area / agent exceeds  $100^2$  versus  $10^2$  in the default simulations), the agents never find each other so no markets form.

Note that this idea of an ‘enabling environment’ can give the impression that a specific institution is in a sense ‘waiting’ for the right conditions before emerging. However, it is perhaps better to think of an environment in which a range of social structures (including institutions) have the *potential* to emerge whereby each structure can be catalysed by a different set of circumstances. Different starting points in the environment’s state space might enable different emergent patterns: in a sense we can think of these as manifestations of the starting conditions.

The third main point is that detrimental institutions can emerge. This is observed when we explore the parameter space of the second model, e.g., when the cost of fighting is close to zero, the opposite of property rights emerges: all the agents learn, eventually, to steal in all interactions. No efficient market emerges, none of the agents specialise, and the agent population collapses.

### 6.4.3 Habituation

As noted above, habituation was applied to each model in the form of a set of experiments. The details and results of these are discussed in more detail in sections 7.6 and 10.3 below. For both, we examine the results when the agents’ mental models change as a result of: (i) habituation only, i.e., reinforcement learning is ‘switched off’; and (ii) both habituation and reinforcement learning.

Here we summarise the results of these experiments.

The approach taken in both sets of experiments mimics Hodgson and Knudsen's (2004) method of applying habituation to their model. In their simulations, a habituation variable accrues over time in a way that means agents are more likely, *ceteris paribus*, to repeat a past decision.

In the first model below, a habituation parameter ( $ha$ ) is added to the weight of an agent's target location in their memories after they have selected a target in each round. In the second model, a similar parameter ( $ha_2$ ) is added to or subtracted from the agents' propensities to steal or defend immediately after they make a corresponding decision. The contribution of these habituation parameters (to memory weights or propensities) do not decay over time and they are not influenced by learning in any way<sup>9</sup>.

### Habituation in the First Model

As mentioned above, in simulations based on the first model, habituation leads to agents behaving more conservatively in their choice of target locations at the beginning of the trading phases. This should not be surprising: it reflects agents being more likely to repeat past decisions. Symmetry breaking is either slowed down or prevented, depending on the strength of the habituation parameter.

Nonetheless, in general the agents transact enough to specialise and, eventually, to bear children, i.e., the results of the default simulations are replicated.

There are three additional points worth noting. First, the emergence of multiple markets means some of the markets are illiquid (e.g., with only two agents<sup>10</sup>). In turn this means allocative efficiency is less than when all the agents visit one market only. However, the distribution of resources is more equal than in the case of a single market. Overall, therefore, the markets are less allocatively efficient but they are more equitable.

To understand the equity point, it is helpful to note that we see in the default simulations agents benefiting if their home location is closer to the single emergent market than further away. They profit from playing an intermediary role more than agents who arrive later<sup>11</sup>, which means resources are allocated in a (slightly) unequal way. In the habituation experiments, multiple markets means agents generally live closer to the market they frequent, so this effect is lessened.

The second noteworthy result is that habituation on its own (when reinforcement learning is switched off) can give rise to markets with sufficiently high turnover for agents to

---

<sup>9</sup>This is slightly different to Hodgson and Knudsen (2004) where habituation accrual decelerates over the life of the agent (see Equation 5.1.1, p. 136), and this variable is kept within the bounds of  $-1$  and  $+1$ .

<sup>10</sup>These markets look more like private contracts than institutions.

<sup>11</sup>These agents obtain 'good' prices from newly arriving agents (who sell 'in bulk') and then sell these on to other agents, resulting in small profits.

specialise. About 12 markets emerge on average: on the whole, the process is less efficient (allocative efficiency is lower and more agents die) than in the default simulations.

The third result worth noting is seen when agents' mental models change due to reinforcement learning and habituation, and the parameter of the latter is relatively low ( $ha < 1$ , approximately). The key point is that markets emerge as a result of reinforcement learning but, over time, habituation dominates.

This combination is illustrated well by figures 6.4 and 6.5 below, which show the time series of a location weight (for the market an agent frequents from Round 31), taken from a representative agent's memories. This weight is decomposed into contributions from reinforcement learning and habituation. Fig. 6.4 shows the data for 200 rounds and Fig. 6.5 for 1,000 rounds.

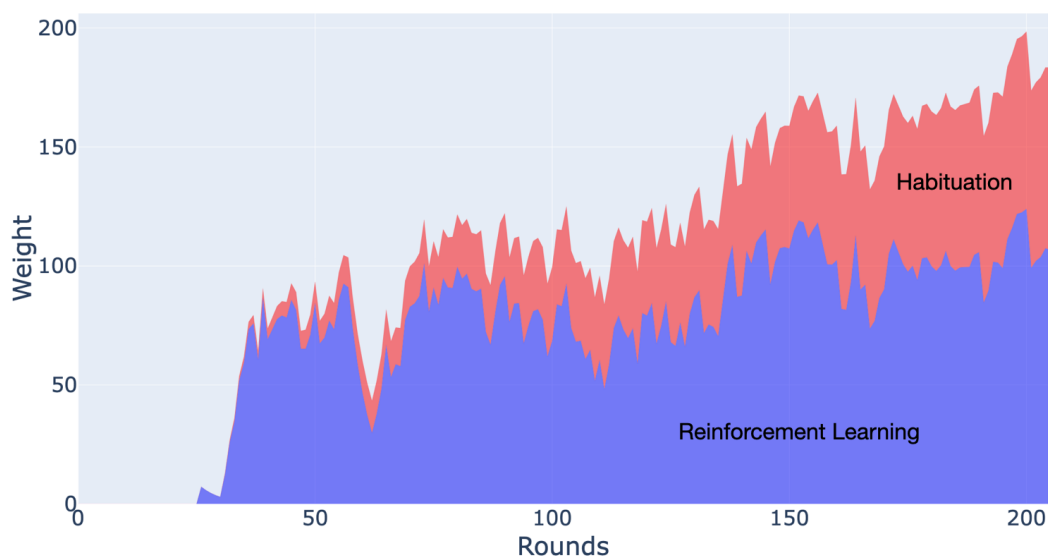


Figure 6.4: A time series of the contributions to a target location's weight in memory over the first 200 rounds of a typical simulation (for one agent). The two 'areas' denote contributions to the total weight over time: the blue area shows the contribution of reinforcement learning and the red area shows that of habituation. This chart shows how the agent 'learned' to visit the location depicted, i.e., reinforcement learning dominated during the emergence of the market.

Fig. 6.4 tells us that the contribution from reinforcement learning dominates the total weight over the first 100 rounds or so, which is when the market is emerging.

Fig. 6.5, however, shows how the contribution from reinforcement learning plateaued over time while the contribution from habituation increased persistently. So, while the agent learned to visit this location early on, the emergent 'rule' in the agent's mental model looked much more like a simple habit by the end of the simulation.






Figure 6.5: The same time series as shown in Fig. 6.4 above but over 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the weight of this location over time.

### Habituation in the Second Model

There are three main findings worth noting.

First, reinforcement learning is necessary to ensure property rights emerge in all simulations. This is observed irrespective of the strength of habituation ( $ha_2$ ), and whether or not mental models change via reinforcement learning. When habituation is ‘strong’<sup>12</sup> we observe scenarios in which 1-3 agents with propensities to steal above 1 are able to thrive by bullying ‘doves’ (agents with negative propensities to steal and defend). This situation is unsustainable and the agent population collapses. This scenario is avoided when the agents’ mental models employ reinforcement learning and  $ha_2$  is not too high<sup>13</sup>: here, agents learn to defend their resources, i.e., there are no doves.

The second finding is that when reinforcement learning is employed alongside habituation (and habituation is not too strong<sup>14</sup>), the latter leads the agents’ propensities to steal to bifurcate into two ‘strategy groups’. In the first group, the agents’ propensities to steal are negative and their propensities to defend exceed 1 (we refer to these as ‘passive-aggressive’ agents). The second group has propensities to steal and defend above 1 (referred to as ‘Al Capone’ agents).

<sup>12</sup>This is when mental models change as a result of reinforcement learning and habituation, and  $ha_2 \geq 0.02$  (approximately), and when reinforcement learning is ‘switched off’.

<sup>13</sup>Approximately,  $ha_2 < 0.02$ .

<sup>14</sup>Again,  $ha_2 < 0.02$ .

We saw a weaker version of this in the default simulations when higher (lower) propensities to steal tend to stay higher (lower), i.e., there was serial correlation (see Fig. 6.3 above). Fig. 6.6 below is equivalent to Fig. 6.3, and is typical of simulations when  $ha_2 = 0.01$ . In these simulations<sup>15</sup>, all the Al Capone agents die because of debilitating fight costs, leaving only agents who respect property.




Figure 6.6: A time series of the agents' propensities to steal over the first 400 rounds of a typical simulation when the agents' propensities change as a result of reinforcement learning and habituation (here,  $ha_2 = 0.01$ ). The red lines depict 'Al Capone' agents, whose propensities increase above 1, and who all died. The blue lines depict passive-aggressive agents (who all survived).

Moreover, in these simulations, property rights emerge more quickly than equivalent simulations without habituation but more agents die in the process.

The third noteworthy finding is the same as that observed in the first model, for low values of  $ha$ . In this situation, the surviving agents co-learn to respect property (the contribution of reinforcement learning is greater than habituation in earlier rounds); however, habituation dominates in the later parts of simulations.

Figures 6.7 and 6.8 below are equivalent to figures 6.4 and 6.5 above but they show the contributions of reinforcement learning and habituation (both positive and negative) to a representative agent's propensity to steal in a typical simulation when  $ha_2 = 0.005$ . Fig. 6.7 shows data for the first 100 rounds and Fig. 6.8 for 1,000 rounds.

#### 6.4.4 The Division of Labour and Knowledge

The results of the default simulations based on the first model point to an explanation for a 'paradox' regarding markets and specialisation in the economics literature.

<sup>15</sup>Here, the agents starting resources are increased to 1,000 units of each resource, which helps us analyse the mechanisms at work.

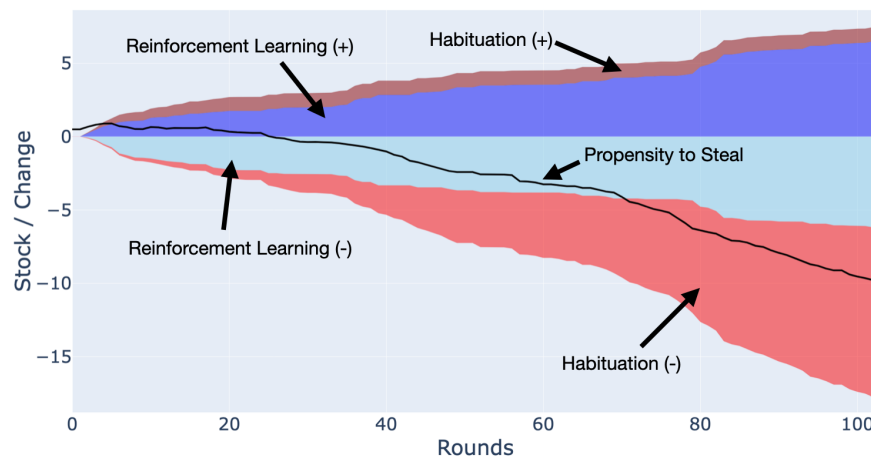


Figure 6.7: A time series of a passive-aggressive agent's propensity to steal (the black line) over the first 100 rounds of a typical simulation when propensities change as a result of reinforcement learning and habituation (here  $ha_2 = 0.005$ ). The y-axis measures the level of the agent's propensity to steal and the contributions to this propensity. The four 'areas' denote contributions to the agent's propensity to steal over time: the upper (dark red) area shows positive contributions by habituation; the second (dark blue) area shows positive contributions by reinforcement learning; the third (light blue) represents negative contributions by reinforcement learning; and the lower (light red) area represents the negative contribution by habituation. This chart shows how the agent's propensity to steal declines below zero (this occurred in Round 26) mainly as a result of reinforcement learning.

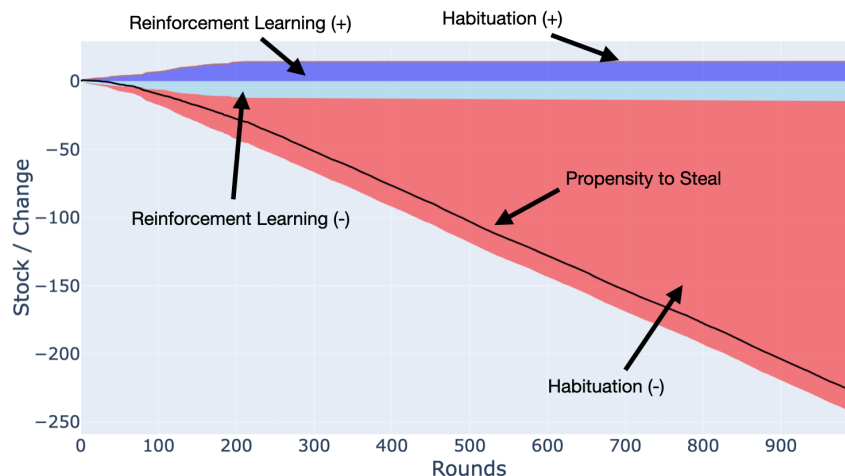


Figure 6.8: The same time series as shown in Fig. 6.7 above but for 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the agent's propensity to steal over time.

The paradox is this: it should be clear that markets are necessary when agents are fully specialised because they provide a medium of exchange. But how do markets form prior to specialisation, if agents are generalists?

In the default simulations a market emerges when generalists trade small amounts of resources, e.g., when they have a small surplus of one resource while other agents have a

small deficit. These modest volumes of transactions early in the simulations are sufficient for a market to emerge. Once this occurs, the agents adapt to the existence of a market by specializing. The volume of transactions then increases significantly.

Note that this description is not presented here as *the* solution to the paradox; rather it is *one* plausible explanation.

It is worth noting, also, that the market changes from being a helpful but non-essential part of the economy to one which is necessary for agent survival (this is when the agents are fully specialised but still need to consume both resources to survive).

Furthermore, we can tie this division of labour to Hayek's division of knowledge. Recall from Chapter 4 that Hayek argued that markets are a more effective means to ensure societies make maximal use of devolved, heterogenous knowledge than systems based on centralised control.

The default simulations based on the first model point to a two-way relationship between markets and knowledge: markets ensure the resources produced by specialists are efficiently allocated; but they also indicate that markets can catalyse heterogeneous knowledge.

In fact, this relationship can be viewed within the wider notion of social structures (here markets and division of labour) *co-evolving*.

### 6.4.5 Liberal Legislation Experiments

In Chapter 12 below we take four scenarios in which we know property rights do not emerge endogenously and examine whether a legal rule might help to enable property rights. As stated in the Introduction, a legal rule here is simply one that is known to all the agents, and is known to be enforced by a third party.

The scenarios are where: (i) agents only acquiesce when other agents attempt to steal their resources; (ii) fight costs are very low; (iii) the outcome of a conflict is a function of the agents' fighting skills (each determined by how many fights each agent has been involved in); and (iv) the outcome of a conflict is a function of total resource holdings and there is one wealthy agent. The latter two scenarios are designed to incorporate forms of power into the simulations.

In all four cases we find there is at least one legal rule that gives rise to property rights.

One helpful way to interpret this result is that an effective legal rule changes a non-enabling environment into an enabling one.

Related to this point, it is important to understand that these legal rules work within a complex and co-adapting environment. It would be wrong to think of these rules as

working as if they are mapped directly into the agents' mental models: the agents are influenced by the rule via reinforcement learning.

Moreover, these rules appear to fit Hayek's reference to 'liberal legislation': the property rights enabled by the legal rule help to enable the emergence of (and allocation within) an efficient market.

Finally, an emphasis in this thesis is on bottom-up pragmatism: with this in mind, variations of the 'liberal legislation' experiments were conducted in which the policing authority might be corrupt. We find that corruption can undermine the efficacy of legal rules such that property rights do not emerge.

### 6.4.6 Invisible Hand Explanations: Characteristics

Chapter 4 included a discussion of Ullmann-Margalit's (1978) three characteristics of invisible hand explanations<sup>16</sup>. As stated above, these should:

- be individual-based;
- include the 'normalcy condition'; and
- be surprising.

Here we briefly discuss how the results of the default simulations meet these three conditions.

#### Individual-Based

As mentioned previously, the models were designed to be 'individual-based' so it should not be too surprising that the explanations of institutional emergence set out above meet this characteristic.

There is one additional point worth emphasising. The organic institutions we saw arose via the co-adaptation of the agents' mental models. In the simulations based on the first model, we observed this via symmetry breaking. In those based on the second model, the agents are all simultaneously influenced by the six patterns identified above - these impacted the agents differently according to who they interacted and communicated with, i.e., it depended on each agent's 'social environment'.

We should view this co-adaptation in the context of the discussion of 'semi-permeable agents' in Chapter 2. Individual agents' mental models are open to information that changes these models. This gives rise to behaviour changes which then impact other agents when then impact others, and so on. The social environment is one of continuous

---

<sup>16</sup>See 4.3.2 (p. 107) for a more detailed discussion of these three.

co-adaptation and in the default simulations of both models, the net result is organic institutional emergence.

Overall, therefore, we can say that the models were individual-based but the organic institutions that emerged did so via the co-adaptation of the agents' mental models.

### **Normalcy**

When the simulation results were analysed and the mechanisms of institutional emergence identified, we found that they met this criterion. Let us look at this in more detail.

In the simulations based on the first model, the concept of symmetry breaking seems somewhat abstract but at the individual agent level, the process seems normal. We saw that agents at smaller markets (i.e., with fewer agents and transactions) were more likely to hear about larger markets than agents at larger markets hearing about smaller markets. This means that those at smaller markets were more likely to 'try out' a larger market they have been informed about than vice versa. Over time, the cumulative effect of this is symmetry breaking and the emergence of a single market.

Equivalently, in the simulations based on the second model, all of the patterns listed above appear to be reasonable reactions of the agents to different social conditions. Provided the environment is sufficiently enabling, their combined effect over time leads to the emergence of property rights.

### **Surprise**

This is perhaps the most interesting of the three characteristics. Given that (i) we have identified certain mechanisms through which institutions emerge in the simulations; and (ii) these seem reasonable and normal, can we claim to be surprised?

The answer is a definite yes.

The mechanisms described above were identified after a lot of investigative work which included the generation of information from the simulations and analysis of the results. With hindsight, the emergence of organic institutions looks less surprising but *ex ante* the results were not predictable.

In the first model's default simulations, the agents' mental models were designed merely to allow them to find other agents. The emergence of a single market was very surprising; moreover, we could not predict before each simulation where the single market would be located.

In the second model's default simulations, the agents' mental models were set up so that they simply learned to defend their resources (or acquiesce) and steal (or trade) based on simple reinforcement learning. It is difficult to anticipate all of the various patterns

identified above, let alone how these combine to give rise to property rights, given this simple feedback mechanism.

Note the ‘surprising’ nature of the single markets and property rights are explored in more detail in the last sections of Chapters 7 and 10, respectively.

### 6.4.7 Emergence and Immergence

Here we note briefly that two related phenomena are observed in the simulations of both models: the emergence of some ‘external’ (to each agent) property and the immergence of some immanent behavioural tendency. The concluding chapter contains a more detailed description of symmetry breaking via emergence and immergence, which is at the heart of the organic institutional emergence observed.

In the simulations based on the first model, the ‘outer’ emergent property is a single market which all the agents attend. In a sense, this ‘market’ is a reification of the fact all the agents target the same location on the grid at the start of the interaction phase of each round, which they then attend. However, it seems reasonable to refer to this location as a market<sup>17</sup>.

In terms of the ‘immergent property’ here, this is the dominance (in terms of memory weights) of one location in all the agents’ memories. When mental models include only reinforcement learning, memory decay means each agent completely forgets about other locations. However, when habituation is included, agents can have multiple locations sustained in memory with positive weights: one location dominates all the others such that the probability of visiting the non-dominant locations declines asymptotically to zero.

The ‘outer’ emergent property in simulations based on the second model is less obvious than those based on the first model but we can state that it is the fact of all the agents trading in interactions rather than stealing. The counterpart immergent property is all the agents’ propensities to steal being negative.

Note that the detailed description contained in the conclusion focuses more on the processes of symmetry breaking via emergence and immergence rather than the end-states described above.

Now that we have discussed the background to the models, their design features, and the main results of simulations based on them, let us look in detail at the first model.

---

<sup>17</sup>This reification is an example of Gilbert’s (2002) second type of cognitive emergence, in which a concept is identified.





# Chapter 7

## Market Emergence Model

If you try and take a cat apart to see how it works, the first thing you have on your hands is a nonworking cat.

– Douglas Adams

This chapter presents the first of the two computational models (the Market Emergence Model) developed for this thesis.

An overview of the model was provided in the last chapter. The first section below (7.1) sets out a more detailed description, including the default parameter set<sup>1</sup>, which contains all of the parameters used in the ‘default scenario’.

Prior to looking at the results from this scenario, in sections 7.2 to 7.4 we look at three ‘null scenarios’. This is in order to contextualise the results from simulations which use the default parameter set:

- Section 7.2 contains the results from simulations when the agents are prevented from trading. This helps us identify the carrying capacity of the environment when agents neither trade nor specialise.
- Section 7.3 presents the results of simulations from the second null scenario. Here, agents can trade but they do not store memories of previous transaction locations which might help them to find other agents. The agents are, however, allowed to specialise.

---

<sup>1</sup>As mentioned in the previous chapter, a ‘default’ set of parameters is specified for both models. When a simulation using these parameters is run, this is referred to as the ‘default scenario’.

- Section 7.4 presents the third null scenario - this time the agents can store memories that might be used to find other agents; however, the agents are not allowed to specialise.

In the default scenario, presented in Section 7.5, the agents can trade and use their memories to find other agents, and they are allowed to specialise. Note the agents' mental models change as a result of reinforcement learning only.

Section 7.6 uses the same parameters as those in Section 7.5 but now the agents' mental models can change as a result of habituation as well as reinforcement learning.

The final section of this chapter (7.7) considers whether the markets we observe emerging in the default simulations (and those including habituation) meet the definition of organic institutions mentioned in the Introduction.

Note that the next chapter contains a summary of the exploration of the parameter space and Appendix B contains a more detailed analysis of these results. This exploration is important in investigating the conditions under which markets emerge and when they do not.

## 7.1 Overview of the Model

Recall from the previous chapter that two resources are generated in an environment and agents must consume *both* resources to survive. The objective of each agent is to stay alive, which they do by foraging from 'resource fountains' and, when it is useful, trading with other agents. In this model, agents respect others' property, i.e., they do not steal resources from other agents - they only trade.

### 7.1.1 The Environment and Time

The environment in which the agents are located is made up of two resource 'fountains' (denoted  $A$  and  $B$ ), which follows Holland (2014), and a grid on which agents can trade.

At the beginning of each round both fountains are replenished with their specific resource, starting each round with  $L$  units each ( $L = 50$  in the default scenario: equivalent to 2 units for each of 25 agents at instantiation). All of the fountains' resources that are not collected perish in between rounds, i.e., no resources are carried over from one round to the next.

Each round is split in to two main phases: (i) foraging, and (ii) trading (see Fig. 7.1 below, which is replicated from the previous chapter for convenience).

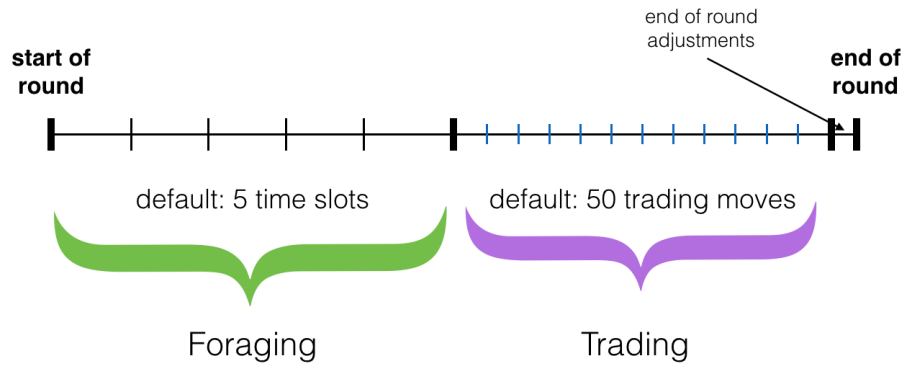


Figure 7.1: A breakdown of each round in the first model. In the first stage of each round, agents forage from resource fountains. In the second stage they try to find each other within a geographic region, in order to trade any resources they might have. At the end of the round, various adjustments are made and some of the agents communicate.

During the foraging phase, agents have five time slots in which they can visit either fountain (they can collect up to one resource unit from the visited fountain per time slot).

After foraging, the agents are placed on a torus on which they can move, i.e., a grid that wrapped around from top-to-bottom and from side-to-side. This torus has dimensions of  $50 \times 50$ .

Each agent has a home location on the grid from which it starts its search for other agents during the trading phase. Agents can see other agents on their current square and adjacent squares, which here means all of the eight Moore's (or 'King's move') squares; however, they can only trade with agents on their current square. The trading phase is split in to 50 time periods: when moving around the grid during these periods, agents can only travel up to 1 grid square at a time.

At the end of each round, i.e., after foraging and trading, the agents consume the resources they hold, update their foraging strategies, and communicate with other agents. Once finished, a new round begins until 1,000 rounds have been completed.

### 7.1.2 The Agents

In the default scenario 25 agents start each simulation. The following is a list of the agents' state variables (individual agents are denoted by the subscript  $i$ ):

- A personal resource array

$$\mathbf{r}_i = [r_i^A, r_i^B]$$

Each element of this array is a stock corresponding to each of the fountain resources (A and B). We can think of these resources as nutrients, both of which are essential for survival. The agents have to maintain a stock of both of these resources by consuming what they have collected or bought during the round (metabolism depletes the agents' stocks by one unit of each resource in between rounds). Once consumed, these resources are 'embodied', i.e., the agents cannot un-consume the resources and trade them.

The resources  $r_i^A$  and  $r_i^B$  are initialised with values drawn from a normal distribution with mean of 50 units and standard deviation of 5.

Agent  $i$  remains alive if both of these resources exceed 0, i.e.:

$$r_i^j > 0 \quad \text{for all } j \in A, B$$

For the agents, therefore, the challenge is to remain alive by sustaining positive values of both  $r_i^A$  and  $r_i^B$ .

- A foraging strategy array

$$\mathbf{h}_i = [h_i^1, h_i^2, h_i^3, h_i^4, h_i^5]$$

The foraging phase of each round is divided in to 5 time slots, which means agents have 5 discrete opportunities to forage from either fountain. This array determines which fountain each agent visits in each time slot. At inception this array is populated randomly, e.g.,  $\mathbf{h}_i = [A, B, B, A, B]$ . Note that the order of this array matters since the quantity of resources at each fountain is finite and could be exhausted before the end of the foraging phase.

- A foraging skill levels array

$$\mathbf{p}_i = [p_i^A, p_i^B]$$

These are probabilities that correspond to levels of skill of the agent in foraging for each resource type. Specifically, each element reflects the probability that an agent will detect and obtain<sup>2</sup> the particular resource associated with a fountain (this is explained further below).

At instantiation:

---

<sup>2</sup>Note that if an agent detects any resource, it always collected it - these words are used interchangeably below.

$$p_i^j = 0.5 \quad \text{for } j \in A, B$$

These skill levels change according to how much time the agent spends foraging for a particular resource: its detection probability for  $j$  will increase if it spends more time foraging for it, i.e., its skill will improve; and vice versa. This process is at the heart of how agents become specialised and is explained further below.

- A basket array

$$\mathbf{b}_i = [b_i^A, b_i^B]$$

This array is used to keep track of the resources successfully collected during the foraging phase of each round, and which might be subsequently traded. The resources remaining in the basket at the end of the round (after foraging and trading) are consumed, i.e., added to  $\mathbf{r}_i$ .

- A memory array

$$\mathbf{m}_i = [m_i^1, m_i^2, m_i^3 \dots]$$

This array records grid locations where agent  $i$  has traded with other agents in previous rounds, and also locations where others have traded that agent  $i$  has been informed about. Memories decay over time, which is explained further below.

- A Home Location on the grid

Each agent is allocated a location on the grid from which it starts the trading phase (a ‘home’). In the default scenario, agents’ homes are evenly spaced on the grid, as shown in Fig. 7.2 below, to maximize sparsity.

An agent can sire children if both its resource holdings exceed 125 units. If this is true of any two randomly paired agents then they will bear a child: 25 units of each resource will be deducted from the parent’s resource arrays, and a child will be instantiated with a personal resource array containing 50 units of each resource. The child will also be given a home location in the grid square furthest away from any other agent (again maximizing the sparsity of the population).




Figure 7.2: The location of agents' homes on the grid. In the default simulations the grid is a torus with dimensions of  $50 \times 50$ . Each of the agents starts the trading phase from its home location. These 'homes' are shown above as black squares. If a new agent is born, its home location on the grid is chosen in such a way as to maximize the sparsity of the population (this is in order to ensure consistency between simulations).

### 7.1.3 Foraging

The resource fountains are not given a geographic location: we assume both fountains are within reach of all the agents. In a sense, the fountains represent work that the agents undertake in order to survive<sup>3</sup>.

In the first time slot of the foraging stage, each agent visits the resource fountain specified by the first cell in its foraging strategy array  $\mathbf{h}_i$ . As a result, at each fountain there is typically a group of agents foraging from that fountain.

These agents form a queue at the fountain (randomly generated), and they then take turns to forage from it. The likelihood of an agent collecting a resource when it is their turn is described below.

If an agent successfully collects one unit of a resource, this is deducted from the fountain's stock so it is not available to other agents, and it is added to the agent's basket array,  $\mathbf{b}_i$ .

After all the agents have foraged from both fountains in the first time slot, the 'clock' ticks forward to the second time slot. The agents move to the fountain designated by the second cells in their foraging strategy arrays.

This process is repeated over the 5 time slots. It is possible that one or both fountains are fully depleted during the foraging phase; however, the agents never know if this happens - they will attempt to collect resources regardless.

<sup>3</sup>When the parameter space is explored, we consider a scenario in which the fountains are given specific locations on the grid and the agents start the trading phase at the fountain where they were located at the end of the foraging phase.

After the foraging part of the round has been completed, agents move on to the trading phase.

#### *Detecting Resources at Fountains*

When foraging for resources, the success of an agent in detecting a resource is dependent on two factors: its skill in detecting that resource (determined by  $\mathbf{p}_i$ ) and the resource's availability at the fountain.

The simplest way to combine these two factors is to multiply an agent's detection probability by the remaining stock of a fountain relative to its maximum stock level. As a result, the probability an agent detects a resource unit at any fountain is:

$$n_i^j = p_i^j \frac{l^j}{L} \quad (7.1)$$

where:

- $n_i^j$  is agent  $i$ 's probability of detecting a resource unit at fountain  $f^j$
- $p_i^j$  is agent  $i$ 's detection skill for resource  $f^j$ , taken from  $\mathbf{p}_i$
- $l^j$  is the stock of resources at fountain  $j$  at the time of foraging
- $L$  is the stock of fountain  $j$  at the beginning of the round (50 by default)

For example, if agent  $i$  reaches Fountain  $A$  when the fountain's resource level is 40 and the agent has a detection skill of 0.5 for that resource then:

$$n_i^A = 0.5 \times \frac{40}{50} = 0.4$$

The agent will have a 0.4 probability of detecting the resource. We can see that an agent is more likely to detect a resource unit if its detection skill is higher and/or if more resources are present at the fountain.

### **7.1.4 Trading**

After the foraging phase is complete, agents are placed on the grid at their individual home locations. See Fig. 7.2 above. Agents then seek out other agents to trade with.

In the default scenario, if agents have no memories (which is true at the beginning of each simulation) they move around the grid in a random walk. This means they choose one of the 9 squares within range (the agent's current square and any of the 8 adjacent squares) with equal probability. During this random walk, if an agent sees another agent

on an adjacent square, it will move on to its square; however, it could not immediately initiate a transaction because agents could either move or trade in each time period, not both.

If agents have memories of previous transactions at the beginning of the interaction phase, they select a target location (from memory) to head towards from their home. In the default scenario, agents use a ‘Roulette Wheel’ approach to choose between different locations in memory, i.e., the probability that an agent chooses a specific location is proportional to its weight in memory<sup>4</sup>.

An important assumption in this model is that agents use their memories of transactions in previous rounds, as signals for where other agents might be in the current round. Put another way, they are used to form expectations about other agents’ locations.

Memories are strengthened by repeated trading at the remembered location but they also decay over time. When an agent transacts at a specific location for the first time, this location is given an initial weight of 1, and each time the agent transacts at this location, its weight in memory is increased by 1. The total weight of each location in memory decays by 20% in between rounds until the weight declines to below 0.05. At this point the grid location is removed from the agent’s memory<sup>5</sup>.

In addition, if the agent is informed about a transaction that had involved two other agents, the weight of this location in memory is increased by 0.5, i.e., half the weight of transactions in which the agent participated<sup>6</sup>.

If an agent heads to a ‘target’ location, it will ignore other agents along the way. Also, once reached, an agent will remain at its target location until the end of the trading phase.

The agents’ memories and the roulette wheel choice algorithm are the main components of the agents’ mental models. Agents use what information they have to make their decision of where to move to but there is no guarantee other agents will be there. In effect, the agents use ‘informed guesswork’ with the only data available.

If two agents meet at the same target location, they will attempt to trade. If multiple agents are at the same location, they are paired randomly and then transact. Each agent can initiate one transaction in each time period of the trading phase (assuming it remains static).

---

<sup>4</sup>For example, if location [3, 15] has a weight of 2 and [35, 17] has a weight of 1 then there is a 2/3 chance the agent will visit [3, 15] and a 1/3 chance it will visit [35, 17].

<sup>5</sup>This means that a single transaction will stay in the agent’s memory for 13 rounds.

<sup>6</sup>A single ‘rumoured’ transaction will stay in memory for 10 rounds.



*Reinforcement Learning*

Here, reinforcement learning is consistent with [Erev and Roth's \(1998\)](#) first and third principles (see Section 1.4.6). The (first) principle of *Law of Effect* should be clear: transactions add to location weights, which increases the likelihood an agent will choose that location again.

The (third) principle is also clear: the Roulette Wheel approach ensures *choice behaviour is probabilistic*.

[Erev and Roth's \(1998\)](#) second principle (*power law of practice*) appears less relevant here because choice is based on relative weights. In principle, location weights could have been designed to accrue by declining amounts over time but this does not seem relevant in this specific context. A choice was made to ensure location weights are proportional to known transactions, with some decay over time due to memory loss.

*Transaction Prices*

In order to determine the prices at which agents transact, the model adopts the process used in [Epstein and Axtell \(1996\)](#): prices are the geometric mean of agents' marginal rates of substitution (MRSs). An MRS is simply the ratio of one resource holding (the combination of the agent's resource stocks and resource basket) relative to another:

$$MRS_i^{AB} = \frac{(r_i^A + b_i^A)}{(r_i^B + b_i^B)}$$

We can think of this as the value an agent attributes to Resource A in units of Resource B. For example, if  $r_i^A + b_i^A = 100$  and  $r_i^B + b_i^B = 50$  then  $MRS_i^{AB} = 2$ . This means the agent values 2 units of Resource A as equivalent to 1 unit of Resource B. From a survival perspective, this makes sense because each agent will attach a greater value to the resource in which it is most deficient (here, Resource B).

The MRS also represents the price above which Agent  $i$  will buy A and sell B and below which it will do the opposite. For example, if the agreed price of Resource A in units of Resource B is 2.5 then Agent  $i$  will prefer to sell B and buy A. By contrast, if the price is 1 then the agent will prefer to sell A in exchange for B.

The price at which the agents trade resources A and B is determined by the geometric mean of the agents' MRSs:

$$Price_{AB} = \left[ \frac{(r_i^A + b_i^A)}{(r_i^B + b_i^B)} \times \frac{(r_k^A + b_k^A)}{(r_k^B + b_k^B)} \right]^{1/2} \quad (7.2)$$

When two agents,  $i$  and  $k$ , interact their MRSs are compared and, if it is beneficial, the agents trade.

The general idea with this approach to pricing is that if agent  $i$  has, say, significantly more of resource A than resource B, then  $MRS_i^{AB}$  will be relatively high. Suppose this agent then meets agent  $k$ , with significantly more of resource B than A ( $MRS_k^{AB}$  is relatively low). The agreed price will be somewhere between  $MRS_i^{AB}$  and  $MRS_k^{AB}$ : the agents will ‘meet’ in the (geometric) middle. This trade is beneficial to both agents because they both sell a resource they have more of in exchange for a resource they have less of.

Note that this trade can only occur if  $i$  and  $k$  have the resources each wants to sell in their basket arrays. In the example,  $i$  and  $k$  will only transact if  $i$  has A and  $k$  has B to sell.

#### *Transaction Quantity*

If the agents have sufficient resources, they will trade an amount that leads to their MRSs becoming equal. If the agents do not have sufficient resources, they will exchange as many resources as possible: their MRSs will converge but not completely.

### 7.1.5 End of Round Adjustments and Communication

At the end of each round, the agents consume the resources they hold in their baskets, their foraging skills are updated ( $\mathbf{p}_i$ ), they adjust their foraging strategies ( $\mathbf{h}_i$ ), and some of the agents communicate. Let us look at each of these in turn.

#### 7.1.5.1 Consumption & Metabolism

The agents add the resources in their foraging baskets ( $\mathbf{b}_i$ ) to their personal resource arrays ( $\mathbf{r}_i$ ); and a metabolism cost of 1 unit for each resource is deducted:

$$\mathbf{r}_i^{end} = \mathbf{r}_i^{start} + \mathbf{b}_i - \mathbf{1}$$

Here,  $\mathbf{r}_i^{end}$  refers to  $i$ 's personal resource array at the end of the round, and  $\mathbf{r}_i^{start}$  at the beginning. The deduction of  $\mathbf{1}$  represents the cost of metabolism.

#### 7.1.5.2 Updating Foraging Skills

The agents' resource detection arrays ( $\mathbf{p}_i$ ) are updated to reflect the agents' experience of foraging in the round: the general idea is that if an agent spent a lot of time foraging for a specific resource during the round, its skill (detection probability) increases, and vice versa.

An adapted logistic equation is used, which allows the detection probabilities to vary between a minimum level of skill ( $\underline{p}$ ) and a maximum level ( $\bar{p}$ ):

$$\Delta p_i^j = \frac{t \cdot [w_i^j - d/x] \cdot (p_i^j - \underline{p})(\bar{p} - p_i^j)}{\bar{p} - \underline{p}} \quad (7.3)$$

There are three parts to this adjusted logistic equation:

1.  $t$  is a speed-of-adjustment parameter.
2.  $w_i^j - d/x$  is a term that ensures the agent's change in skill is positive (negative) if the agent spends more (less) time than average foraging for a specific resource:  $w_i^j$  is the total number of time slots agent  $i$  spends foraging for resource  $j$  during the round;  $d$  is the total number of time slots; and  $x$  is the number of resource fountains. In the default scenario,  $d = 5$  and  $x = 2$ . This means that, in the default scenario, if an agent spends more than 2.5 of its 5 time slots foraging for resource  $j$  then its skill increases, and vice versa.
3.  $\frac{(p_i^j - \underline{p})(\bar{p} - p_i^j)}{\bar{p} - \underline{p}}$  ensures the adjusted logistic equation has a minimum value of ( $\underline{p}$ ) and a maximum value of ( $\bar{p}$ ). In the default scenario, agents' skill probabilities have a floor of 0.2 and a ceiling of 1.0, i.e., the agents can attain skill 'perfection' but they will always have some nominal level of skill even if they stopped foraging for a particular resource.

This logistic equation (which generates a sigmoid curve) is used in the default simulations. There is some evidence that learning from a minimal level of understanding in some field describes a sigmoid curve (e.g., [Leibowitz et al, 2010](#); and [Johnson, 2012](#)). Moreover, the top half of this curve is consistent with [Erev and Roth's \(1998\)](#) second principle of reinforcement learning (*power law of practice*).

Note that a linear approach is considered when the parameter space is explored.

### 7.1.5.3 Updating Foraging Strategies

In updating their strategies, agents are allowed to change one randomly chosen cell of the five cells in their foraging strategy arrays ( $\mathbf{h}_i$ ). In doing this, agents will seek to maximize the acquisition of the resource in which they are most lacking in their personal resource arrays (denoted  $r_{min}$ ), i.e., if they hold less of Resource A than B then they will seek to increase their holdings of A in the next round.

In general, an agent has two ways to increase its holding of  $r_{min}$ : (1) choosing to forage from the counterpart resource fountain in the next round; or (2) choosing to forage for

the resource in which it is most skilled (denoted  $r_{p_{max}}$ )<sup>7</sup> and then attempting to trade  $r_{p_{max}}$  for  $r_{min}$  (if different) in the next round. When  $r_{p_{max}}$  and  $r_{min}$  are different, we refer to this challenge as a *work conundrum*.

Suppose, for example, an agent holds fewer units of Resource A and it must decide whether to change  $h_i^A$  from B to A. Suppose, also, that its detection probability for A ( $p_i^A$ ) is 0.3 and its detection of probability for B ( $p_i^B$ ) is 0.7. The agent faces a conundrum: it could decide to forage from Fountain A; or it could forage from Fountain B (with a higher expected yield), hoping to trade from B to A.

What the agent decides will depend on its detection probabilities, the expected exchange price of the resources, and its expectation of being able to trade (this is explained in Section 7.5.2 below). If it believes the probability of trading will be low then it will be less likely to forage for its specialised resource, and vice versa.

One addition to this process is an error term that reflects the idea that agents do not have perfect information, nor perfect processing power. The agents' expected yields for both resource choices are calculated (accurately) and these are then adjusted by a value drawn from a normal distribution with mean 0 and standard deviation of 0.1, to form the agent's expectations. Here, the higher the standard deviation, the less accurate is the agent's estimated foraging yields.

For the purpose of clarity, let us look at this formally from one agent's point of view.

For Agent  $i$ , the expected foraging yield of Resource  $j$  in any time slot  $t$  is determined by the following equation:

$$E(\text{yield}_i^{j,t}) = p_i^j \cdot \frac{E(l^{j,t})}{L} + \varepsilon^{j,t} \quad (7.4)$$

where:

1.  $E(\text{yield}_i^{j,t})$  denotes Agent  $i$ 's expected yield of resource  $j$  in time slot  $t$ .
2.  $p_i^j$  is  $i$ 's detection skill for resource  $j$  in time slot  $t$ .
3.  $E(l^{j,t})$  refers to the expected stock of resource  $j$  at fountain  $j$  during time slot  $t$ .  
In the model, this is derived from taking the mean of the fountain's beginning and end stocks for each time slot ( $t$ ) in the previous round.
4.  $L$  is the starting stock of all fountains ( $L = 50$  in the default scenario).

---

<sup>7</sup> $p_{min}$  and  $p_{max}$  denote the agent's minimum and maximum foraging skills, respectively (stated as probabilities):  $r_{p_{max}}$  denotes the resource associated with its maximum detection skill.

5.  $\varepsilon_i^{j,t}$  is a variable taken from a normal distribution with mean zero and a standard deviation of 0.1 in the default scenario.

Note that  $p_i^j \cdot \frac{E(j,t)}{L}$  is equivalent to equation 7.1, which defines the probability of detecting a single unit of resource  $j$  in a single time slot. The expected foraging yield is the same as the expected probability of detection in each time slot since agents are attempting to detect a single resource.

This updating of foraging strategies is central to whether or not agents specialise: in the default simulations, we find that once markets formed and the probability of trading increased toward 1 then agents did indeed choose to specialise.

Moreover, we also find that specialisation occurs when an agent enters a positive feedback loop: it forages for the resource for which it had greater skill ( $r_{pmax}$ ), and this typically increased the agent's skill for that resource, making it more likely the agent subsequently chooses to forage for  $r_{pmax}$  when it adjusts its foraging strategy. This is discussed in more detail in Section 7.5.2 below.

#### 7.1.5.4 Communication

At the end of each round, agents communicate with each other: there is a 1% probability that any pair of agents communicate (regardless of location). If two agents communicated, they share information about every transaction they had participated in during that round.

After these end-of-round processes are complete, the next round began until the end of the 1,000th round.

Now that the model has been explained, let us look at the results of three 'null' scenarios before those of the default scenario and the habituation experiments.

## 7.2 Results: Null Scenario - No Trading

As mentioned above, three null experiments were conducted. The first, which is presented here, is the simplest because agents can only forage for resources; they are not allowed to trade.

The rationale for this first null scenario is to evaluate the carrying capacity of the environment when agents cannot trade.

Fig. 7.3 below shows the total population in a typical simulation (extended to 2,000 rounds to be sure of a steady state result): it falls to 15 agents by Round 800 (approximately), after which it stabilizes.




Figure 7.3: Total agent population over 2,000 rounds when the agents are prevented from trading. Over 20 simulations the total population declines to approximately 15-16 agents, which can be viewed as the carrying capacity of the environment when the agents neither trade nor specialise.

Over 20 simulations the mean population in the last 100 of 2,000 rounds was 15.7 (with a standard deviation of 0.7). The carrying capacity of the agents' environment (in the absence of trading) is, therefore, approximately 15-16 agents.

The agents remained generalists throughout these simulations: their foraging strategy arrays were roughly evenly split between the two fountains and their foraging skills remained at approximately 0.5 throughout.

### 7.3 Results: Null Scenario - No Memories with Specialisation

Here we look at the results from the second null scenario, when agents can trade but are restricted to moving around the grid randomly.

Fig. 7.4 below shows a heatmap of the total volume of transactions in each grid square in the last 100 rounds of a typical simulation. It shows a random spread of transactions.

In this scenario, all transactions resulted from agents randomly bumping in to each other.

A useful representation of the (in)efficiency of this null scenario is to compare the actual volume of transactions on the grid with a representation of the market clearing volume of trade. Fig. 7.5 below shows the supply and demand curves for Resource A in a typical




Figure 7.4: Heatmap of transactions in each grid square (last 100 rounds of a typical simulation) when the agents hold no previous transactions in memory. The agents move around the grid in a random walk and interact only when they bump in to each other. This results in a ‘noisy’ dispersion of transactions.

round of a typical simulation; and it shows the actual volume of transactions and the mean transaction price (the red star).

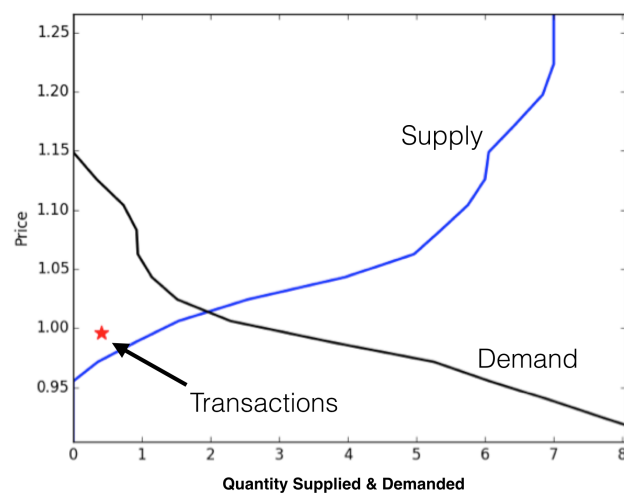


Figure 7.5: Supply and demand curves and total transactions in one round of a typical simulation when the agents have no memories of previous transactions. The supply and demand curves are calculated after foraging and before trading, and are generated by discovering what each agent would have supplied or demanded for each resource over a range of prices; and then aggregating these over all agents. The red star indicates the total volume (measured on the horizontal axis) and the mean price (measured on the vertical axis) of the subsequent transactions in the same round. The mean ratio of actual transactions to the market clearing volume was 18% over 20 simulations.

On average we found that the volume of transactions was about 18% of the market clearing volume: agents are generally left unsatisfied. Many wanted to trade but could

not find other agents to do so.

In these simulations, none of the agents specialise even though they were able to: detection probabilities remained at about 0.5 for all agents. The total agent population declined to the carrying capacity observed in the preceding sub-section. An analysis of the data showed that the transactions that did occur helped to equalize the resources in agents' personal resources arrays (but only very marginally). The lack of specialisation meant the total population did not increase above the carrying capacity.

Note that an additional null experiment was conducted in which the agents could trade but not use memories nor specialise. The results were essentially the same as in this second null experiment. This should not be surprising: in both scenarios the agents walk around the grid randomly and cannot specialise.

## 7.4 Results: Null Scenario - Memories without Specialisation

If agents use their memories of locations in order to choose a target location, we find that, over time, they converge on a single grid square where (essentially) all transactions take place.

Fig. 7.6 below shows the heatmap of transactions during the last 100 rounds of a typical simulation. It is equivalent to Fig. 7.4 above: here, 99.9% of transactions took place on a single grid square<sup>8</sup>.

A single market typically emerged over approximately 200 rounds: initially, several markets emerge as agents bump in to each other and then (in some cases) report these interactions to other agents. The agents' mental models co-adapted and there was also symmetry breaking, which meant larger markets dominated smaller markets until only one market was left. This process is described in more detail Section 7.5.1 below.

Fig. 7.7 below is equivalent to Fig. 7.5 above: it shows how the existence of a market moves the actual volume of transactions much closer to the market clearing volume.

Fig. 7.8 below shows a time series of the actual volume of transactions divided by the market clearing volume of transactions (denoted here as the 'turnover ratio') in a typical simulation. This is the quantity associated with the red star in Fig. 7.5 divided by the market clearing volume (at the intersection of the supply and demand curves). We can see how this ratio rises from zero to approximately 1 over about 200 rounds.

---

<sup>8</sup>Recall that agents do not transact on the way to a target location - they only do so after they reach their target. We relax this assumption when we explore the parameter space in the next chapter.






Figure 7.6: Heatmap of transactions in each grid square (last 100 rounds only) when the agents use their memories of previous transactions to find other agents. Over each simulation they converge on a single grid location in order to trade (a ‘market square’). Note, however that this location differed between simulations: it was an emergent property of each simulation.




Figure 7.7: Supply and demand curves and total transactions in one round of a typical simulation when the agents use their memories of previous transactions to find other agents. See Fig. 7.5 above for a description of the data. On average the ratio of actual transactions to the market clearing volume was 100.0% over 20 simulations.

Agents are not permitted to specialise in this scenario. This means that detection probabilities are kept constant at 0.5. We found that, even though agents could transact efficiently, the overall foraging yield was not sufficient for all 25 agents to survive. We again saw the total agent population decline to the carrying capacity of approximately 15-16 agents.




Figure 7.8: A time series of the turnover ratio over 500 rounds when agents use their memories of previous transactions to find other agents. The turnover ratio is defined as actual transactions divided by the market clearing volume. It was measured and recorded in each round. In this chart we can see how the turnover ratio increased toward 1 as agents find each other and trade. The ratio is maintained at approximately 1 after a single market emerges.

## 7.5 Results: Default Scenario - Memories with Specialisation

In these experiments, agents use their memories to decide on target locations and they are allowed to specialise.

The results are similar to the previous scenario in that a market emerge but, this time, most agents become fully specialised, i.e., the agents' skill for one resource ( $p_{max}$ ) increased to 1 and, perturbations aside, they remained at this resource fountain throughout the foraging phases thereafter.

We found the total agent population increases to approximately 43 agents on average. Typically, 37 agents reached full specialisation and there are on average 5-6 agents born who fail to reach full specialisation. These agents eventually die and are replaced with new, unskilled agents, who also fail to specialise.

A measure that helps us visualise the move from generalists to specialists is shown in Fig. 7.9 below, which was taken from a typical simulation. The chart shows the 'mean specialisation value' over all agents during the first 500 rounds of one simulation (this is explained in Fig. 7.9 - a value of 3 means all the agents are perfect generalists and 5 means they are perfect specialists). In this run, most agents have specialised by approximately Round 200. The oscillations around 4.5 are because most of the

agents achieve a mean specialisation value of approximately 5 and a minority have mean specialisation values of 3 or 4 (those who fail to reach full specialisation and die).




Figure 7.9: Time series of the mean specialisation value in a typical simulation using the default parameter set. For each agent, the ‘specialisation value’ is the number of times an agent forages from the fountain it visits most often (given 2 fountains and 5 foraging opportunities, this value varies between 3 and 5). For the agents as a whole, an average value of 3 would mean all the agents were generalists and a value of 5 would mean they were all specialists (foraging from the same resource fountain in all five time slots during the foraging phase). The data here shows how the agents became more specialised over the first 200 rounds of the simulation, which correlates with the turnover ratio.

This move from generalists to specialists is due to the existence of a market which increases the agents’ expected probabilities of transacting. When their foraging strategy arrays are updated, agents are generally encouraged to select the fountain associated with their highest detection probability. In a sense, when agents are confident of trading, they view the two resources as interchangeable, which means they chose to forage for the resource with the highest expected yield. This is explained further in Section 7.5.2 below.

We can say, therefore, that markets enabled specialisation in these simulations.

It was noteworthy that, prior to specialisation, agents meet only to exchange marginal quantities of resources: as generalists, they forage from both fountains and there is little benefit from transacting (but not none). As agents become specialised, they forage for a single resource but they still need both resources to survive, which means exchange is necessary. As a result, the total average volume traded after a single market emerges (but before agents specialise) was approximately 0.15 resource units per agent per round. By the end of the simulations (when agents had specialised), this increased to 0.54 resource units.

### 7.5.1 The Emergence of the Market

Here we look at how the markets emerge in the default simulations.

As mentioned above, the agents start the interaction phase at their home locations and then walked around the grid randomly.

Eventually, two agents come across each other and attempt to trade. However, the agents would only do so if both held the resources the other wanted to buy. At the beginning of the simulations, when agents are generalists, only about 1 in 4 interactions results in a transaction.

Fig. 7.10 below helps us visualise what happened during the first 50 rounds of a typical simulation. Each line corresponds to the total number of agents targetting a specific square in each of the 50 rounds ('none' is also included: this is the number of agents walking around the grid randomly).




Figure 7.10: A time series of the number of agents targetting six locations on the grid during the first 50 rounds of a typical simulation that used the default parameter set. The chart shows how multiple markets emerged during the first 30 rounds but one of these ([12, 8]) was visited by more agents. Symmetry breaking meant all the agents eventually converged on this location.

If the agents transact then this is recorded in both agents' memories. The first transaction typically takes place within the first five rounds of any simulation.

On average, there were approximately 5 locations on the grid where transactions occurred as a result of agents walking around randomly, and these were typically first observed within the first 15 rounds. We can think of these as *proto-markets*.

In the absence of any other memories, the agents who form these proto-markets revisit them in the rounds after they first appear in memory<sup>9</sup>. It is likely that the two agents who first transact at this location meet again on this square in subsequent rounds: in some rounds they will transact and in others they will not (given the point mentioned above about there being only a 1 in 4 chance of two agents transacting). However, given that single transaction locations remain in memory for 13 rounds, it is likely they will transact again, eventually, and we will observe the location's weight in memory increase as a result.

After these proto-markets are formed, the agents communicate about them. Recall there is a 1% chance that any pair of agents communicate at the end of each round: in the simulation from which Fig. 7.10 is taken, location [12, 8] became a proto-market in Round 2 and was communicated to two other agents at the end of that round. As a result, four agents visited this location in Round 3. None of these agents transacted in that round but all of them transacted in the next round, increasing the weight of this location in their memories.

If we look closer at what happened to these four agents (let us denote them  $i$ ,  $j$ ,  $k$ , and  $l$ ) we observe the phenomenon of co-adaptation. Agents  $i$  and  $j$  formed the proto-market and then communicated its location to  $k$  and  $l$ : the participation of  $k$  and  $l$  meant more transactions occurred at that location than would have otherwise. This had an impact on  $i$  and  $j$ 's mental models, which contained higher weights attributed to this location than would have existed if  $k$  and  $l$  had not visited this location.

We can see from Fig. 7.10 that a number of proto-markets formed but eventually disappeared. For example, location [17, 2] was formed at the same time as [12, 8], in Round 2; however, unlike [12, 8] this location was only communicated to one other agent at the end of Round 3. It is noticeable from Fig. 7.10 that 5 other proto-markets were formed at the same time or after [12, 8] but that none of these survived beyond Round 29.

The reason for this pattern is the existence of symmetry breaking in the simulations. As mentioned above, this means that the markets with more agents tend to cannibalize those with fewer.

To understand this symmetry breaking, consider a situation in which 12 agents visited location Y but 13 visited location X. Symmetry breaking occurred because the probability that an agent at Y would be informed about X was (slightly) higher than vice versa. There was a second influence: it was likely that each agent would have taken part in more transactions at X than at Y so if two agents did communicate, there were probably more transactions at X to communicate about than those at Y.

---

<sup>9</sup>Recall that a memory decay rate of 20% means single transactions remain in agents' memories for 13 rounds.

These two influences together meant that those markets with more agents tended to dominate the smaller ones until only one market was left.

Fig. 7.11 below demonstrates this symmetry breaking particularly well (reproduced from the previous chapter for convenience). This chart is taken from a different simulation in which two markets initially dominate others (at [31, 16] and [22, 30]).




Figure 7.11: A time series of the number of agents targetting different locations on the grid during the first 50 rounds of a typical simulation that uses the default parameter set. This chart demonstrates a scenario when two markets dominate for a period but symmetry breaking meant that, eventually, one dominates. Reproduced from Fig. 6.2 in Chapter 6 for convenience.

In this figure, the number of agents targetting locations [31, 16] and [22, 30] (the two dominant markets) were approximately equal until circa Round 32.

A degree of randomness was built in to the selection of target locations (when agents had more than one location in memory) and in to the communication process. These resulted in perturbations in the agents' choice of target locations.

The final point to make here is to link the formation of a single market square to the concept of emergence in the complexity sciences. The market appears to be an emergent property in the sense that, to paraphrase the definition stated in Section 2.2.5, the agents collectively acquired a qualitatively new property that cannot be understood as the simple addition of their individual contributions. Moreover, the market was clearly useful to the agents because it helped them survive.

## 7.5.2 The Emergence of Specialization

Here we look more closely at how agents came to specialise in the default scenario<sup>10</sup>.

To do this we must first consider in more detail how agents decide between foraging for Resource A or B when they adjust their foraging strategy arrays at the end of each round.

As mentioned in Section 7.1.5 above, the agents are allowed to adjust one of the five cells (randomly chosen) in their foraging strategy arrays ( $\mathbf{h}_i$ ) at the end of each round. When they do this, they choose to maximize the expected yield of the resource they will be most deficient in ( $r_{min}$ ) in the next round.

If the agent has a greater skill (detection probability) in foraging for  $r_{min}$  then the choice is simple: it will visit Fountain  $r_{min}$  during the time slot in question.

Suppose, however, there is a conundrum: the agent is deficient in the resource it is less skilled in foraging for.

If we start with Equation 7.4 and substitute  $r_{min}$  for  $j$ , we find the expected yield of foraging for Resource  $r_{min}$  is:

$$E(\text{yield}_i^{r_{min}, t}) = p_i^{r_{min}} \cdot \frac{E(l^{r_{min}, t})}{L} \quad (7.5)$$

Note we have dropped  $\varepsilon_i$  since the expectation of this would be zero.

The agent will compare this to the expected yield as if it forages for the resource it is more skilled in foraging for ( $r_{pmax}$ ) and then trading this with another agent for  $r_{min}$ . This expected yield is more complicated than  $E(\text{yield}_i^{r_{min}, t})$  above because the agent has to account for the expected price of the transaction and the likelihood of being able to trade.

Let us denote the expected yield of  $r_{min}$  received for the expected foraging yield of  $r_{pmax}$  as  $E(\text{yield}_i^{r_{pmax} \rightarrow r_{min}, t})$  where:

$$E(\text{yield}_i^{r_{pmax} \rightarrow r_{min}, t}) = p_i^{r_{pmax}} \cdot \frac{E(l^{r_{pmax}, t})}{L} \cdot \text{Price}^{r_{min} \leftrightarrow pmax \times 1} \cdot \text{Prob}_{trans} \quad (7.6)$$

The first part of the right hand side of Equation 7.6 ( $p_i^{r_{pmax}} \cdot \frac{E(l^{r_{pmax}, t})}{L}$ ) is the expected foraging yield of Resource  $r_{pmax}$  in time slot  $t$ .

<sup>10</sup>The model is designed with only two resource fountains; however, the principles shown below can extend to more than this.

To find the counterpart expected yield in units of  $r_{min}$  (via exchange), we must multiply this by the expected price of  $r_{min}$  in units of  $r_{pmax}$  (denoted  $\text{Price}^{r_{min} \leftrightarrow r_{pmax} \times 1}$ ) and multiply by the expected probability of transacting,  $\text{Prob}_{trans}$ .

Suppose, for example,

$$\begin{aligned} p_{max} &= 0.7 \\ E(l^{r_{pmax}, t})/L &= 0.8 \\ \text{Price}^{r_{min} \leftrightarrow r_{pmax} \times 1} &= 0.95 \\ \text{Prob}_{trans} &= 0.9 \end{aligned}$$

then:

$$E(\text{yield}_i^{r_{pmax} \rightarrow r_{min}, t}) = 0.7 \times 0.8 \times 0.95 \times 0.9 = 0.4788 \quad (7.7)$$

Now, the agent will choose to forage for  $r_{pmax}$  in time slot  $t$  if:

$$E(\text{yield}_i^{r_{pmax} \rightarrow r_{min}, t}) > E(\text{yield}_i^{r_{min}, t}) \quad (7.8)$$

For example, if  $E(\text{yield}_i^{r_{min}, t}) = 0.4$  the agent will prefer to forage for  $r_{pmax}$  in time slot  $t$  rather than  $r_{min}$  because the expected yield will be slightly higher.

If we substitute equations 7.5 and 7.6 in to Equation 7.8 and rearrange, we arrive at the following, which is the condition for the agent choosing to forage for  $r_{pmax}$  rather than  $r_{min}$  in time slot  $t$ , i.e., when there is a work conundrum:

$$\frac{p_{r_{pmax}}}{p_{r_{min}}} \cdot \frac{E(l^{r_{pmax}, t})}{E(l^{r_{min}, t})} \cdot \text{Price}^{r_{min} \leftrightarrow r_{pmax} \times 1} \cdot \text{Prob}_{trans} > 1 \quad (7.9)$$

This equation can be interpreted by stating that an agent will be more likely to choose  $r_{pmax}$  in a work conundrum:

1. the more skilled it is at foraging for  $r_{pmax}$  versus  $r_{min}$ ;
2. the higher the expected fountain stock for resource  $r_{pmax}$  relative to the expected stock of  $r_{min}$ , in time slot  $t$ ;
3. the more of  $r_{min}$  it expects to receive in exchange for  $r_{pmax}$ ; and
4. the higher the expected probability of transacting.



It is this last point that helps us understand why the emergence of specialisation was conditional on markets existing: the higher the expected probability of transacting, the more inclined agents were to choose  $r_{p_{max}}$  in a work conundrum rather than  $r_{min}$ .

To help us understand the emergence of specialisation better, it is informative to rearrange Equation 7.9 in order to derive a threshold probability of transacting above which an agent will choose  $r_{p_{max}}$  and below which it would choose  $r_{min}$  in a work conundrum:

$$\text{Prob}_{Threshold} = \frac{p_{r_{min}}}{p_{r_{p_{max}}}} \cdot \frac{E(l^{r_{min}, t})}{E(l^{r_{p_{max}}, t})} \cdot \frac{1}{\text{Price}^{r_{min} \leftrightarrow r_{p_{max}} \times 1}} \quad (7.10)$$

We can generate and track this threshold probability for every agent in each round because we can determine values (or their expectations) for each element on the right hand side of this equation. We can also compare this threshold with the actual expected probability of transacting: if the latter exceeds the former, the agent will specialise in a work conundrum<sup>11</sup>. Fig. 7.12 below includes the time series of this threshold probability (the red line) for one agent over the life of a typical simulation.

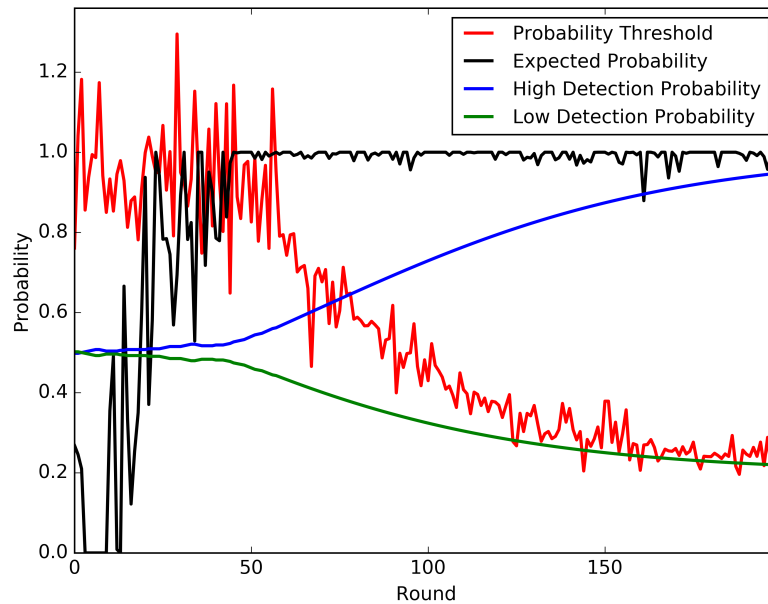


Figure 7.12: The specialisation threshold for an agent over 200 rounds of a typical simulation. The model keeps track of this metric (the red line) for each agent over each round: if the turnover ratio experienced by the agent and its neighbours (the black line) exceeds this value then agents tend to specialise, and vice versa. The blue line shows the foraging skill (as a probability of detection) of the resource in which the agent became specialised, and the green line shows the non-specialised resource.

<sup>11</sup>Note that even when there is no such conundrum,  $\text{Prob}_{Threshold}$  was generated by the model as a ‘thought experiment’.

The chart shows four time series. The red line is the threshold probability which was calculated in each round by applying Equation 7.10. Note that it started the simulation at close to 1 and often exceeded this value.

The black line is the assumed probability of transacting (taken as the mean turnover ratio for the agent and its neighbours whose homes were within 10 grid squares of the agent in the preceding round). This started the simulation at well below 1 but it reached 1 by approximately the 50th round.

The blue line shows what ended up as the agent's highest detection probability at the end of the simulation ( $p_{max}$ ); and the green line is what ended up being the agent's lowest detection probability ( $p_{min}$ ).

In this particular simulation,  $E(l^{r_{min}, t})/E(l^{r_{pmax}, t}) \approx 1$ , i.e., expected fountain levels tended to be about equal and the price remained close to 1 ( $\text{Price}^{r_{min} \leftrightarrow r_{pmax} \times 1} \approx 1$ ), so the variation in the threshold probability was largely due to the variation in detection probabilities  $p_{r_{min}}/p_{r_{pmax}}$  (see Equation 7.10).

When the black line in Fig. 7.12 was above the red line, the agent chose to specialise when faced with a work conundrum. Moreover, when there was no conundrum, this meant  $r_{min}$  was also the resource with the highest detection probability - here the agent would invariably choose  $r_{pmax} = r_{min}$  when it updated its foraging strategy.

In this particular simulation, both the red line and the black line oscillated close to 1 after the market emerged and prior to Round 50 (approximately). As a result, the agent chose to specialise in the face of a conundrum about half of the time, which typically meant it foraged for its higher skill resource in 3 or more of the 5 time slots. This created a bias toward specialisation that meant the detection probability for the higher skill resource increased slightly, and vice versa.

This process contained a positive feedback loop (noted earlier): a greater foraging skill for one resource (combined with a high expected transaction probability) would lead an agent to forage more for that resource on average; and this in turn increased the agent's detection probability for that same resource. Ultimately, agents became fully specialised in one resource. We can see this in Fig. 7.12 above: the blue lined increased more quickly after the black line stabilized at close to 1.

An interesting feature of Fig. 7.12 is the decline in the threshold probability after approximately the 50th round.

Mathematically, we can understand this by examining Equation 7.10: as  $p_i^j$  rose and  $p_i^{r_{min}}$  declined,  $\text{Prob}_{Threshold}$  also declined (the two other factors on the right hand side of Equation 7.10 were approximately 1). In fact, we can see that because  $p_{max}$  increased to 1 and  $p_{r_{min}}$  fell to 0.2,  $p_{r_{min}}/p_{r_{pmax}}$  also fell to 0.2, which we can see in the chart.

Conceptually, this decline in the threshold probability is related to a skill *lock-in* effect. As one foraging skill ( $p_{max}$ ) improved and the other ( $p_{min}$ ) declined, the threshold probability of transacting ‘required’ for the agent to specialise also declined. This lock-in effect meant that when the agent was fully specialised ( $p_{max} = 1$  and  $p_{min} = 0.2$ ), it still made sense for them to choose to forage for  $r_{p_{max}}$  even when the probability of transacting was only slightly above 0.2.

We will make further use of equations 7.9 and 7.10 in Appendix B where we explore the parameter space in detail.

### 7.5.3 Proximity to the General Equilibrium Outcome

An interesting question to consider is how closely the results of the default scenario approximate the General Equilibrium outcome of Arrow and Debreu (1954). This would be where the total volume transacted equals the market clearing quantity; and transaction prices equal the market clearing price.

We can answer this question in two ways: by looking at the turnover ratio and by looking at the mean and standard deviation of market prices. If the turnover ratio was 1.00 and market prices were identical to the market clearing price (with no variation) then the General Equilibrium outcome would have been precisely replicated.

Table 7.1 below shows turnover and price data for 20 simulations which used the default parameter set.

	<b>Default Scenario</b>
<b>Turnover Ratio</b>	1.000 (0.0020)
<b>Mean Price Difference</b>	0.19% (0.14%)
<b>Standard Deviation of Daily Price Changes</b>	1.77% (0.16%)

Table 7.1: Proximity to general equilibrium over the last 100 rounds of 20 default simulations. If the default simulations replicated the general equilibrium model then the turnover ratio would be 1, the mean price difference (mean actual prices divided by the market clearing price) would be zero, and the mean daily standard deviation of price changes would also be zero. These data indicate that the results of the default simulations did indeed approximate the results of the General Equilibrium model.

These data resemble the General Equilibrium outcome: the turnover ratio is very close to 1.00; and there is a small difference between mean market prices and the market clearing price in each round (about 0.19%). Prices varied little: the standard deviation was equal to 1.7% of the mean market price.

Note, however, that these results were contingent upon the emergence of a market institution. Without this institution, the general equilibrium results would not have been approximated.

Moreover, it is worth considering these results in the context of the famous quote from [Williamson \(1975\)](#), that “in the beginning there were markets” (p. 20). The results presented in this chapter reinforce the idea that *markets are institutions* which have to be explained and not assumed (which links to the discussion of pre-existing institutions in Section [1.1.2](#)).

## 7.6 Habituation

Here we discuss the design and results of experiments which explore the impact of habituation on the simulation results.

Broadly speaking, this was done by mimicking [Hodgson and Knudsen’s \(2004\)](#) approach to habituation whereby habits are viewed as a “summation of behaviors in an unbounded set of present and past periods” (p. 35).

Habituation is introduced into the models by adding a value  $ha$  to the weight of any square an agent visits in each round. This value does not decay over time in the way weights associated with transactions do, nor are these weights influenced by any transactions. The parameter  $ha$  simply reflects the idea that an agent is more likely to revisit a location it had previously visited.

With the inclusion of  $ha$ , the agents’ selection of target locations is now determined by two different factors: reinforcement learning and habituation. A roulette wheel approach is still used to select a target location but now the weights of locations are determined by these two factors.

The inclusion of  $ha$  means the model is now similar to [Hodgson and Knudsen’s \(2004\)](#) framing except here the agents’ mental models include reinforcement learning instead of innate dispositions; and there are 2,500 potential options to choose from (locations on the grid) instead of 2.

The main question we are interested in here is about the impact habituation has on the emergence of markets as organic institutions. The idea that habituation can reinforce market as institutions after they have emerged via reinforcement leaning seems intuitively appealing but what about when they are emerging? Section [7.6.1](#) below looks at the results when the agents’ mental models change as a result of reinforcement learning and habituation, for different values of  $ha$ .

Another interesting question is what happens when reinforcement learning is ‘switched off’ and location weights are only determined by habituation. Section [7.6.2](#) looks at these results.

### 7.6.1 Reinforcement Learning and Habituation

The most significant result of the experiments is that habituation tends to interrupt symmetry breaking, which leads to an increase in the number of market institutions existing at any one time. This reduces the allocative efficiency of the markets but it enhances equality because agents are on the whole closer to the market they visit (discussed further below).

A second result is that for lower values of  $ha$ , reinforcement learning dominates the process of markets emergence but, post-emergence, habituation dominates. Agents appear to co-learn to visit specific locations but, over time, their choices look more like habits.

Looking at the results in more detail, at the agent level we observe, unsurprisingly, that agents behave more ‘conservatively’ when  $ha > 0$ , tending to visit locations they had visited in the past rather than trying out new locations suggested by others. We will see below, however, that when the data is analysed, there are some unexpected results, e.g., agents going to the better-attended markets are more likely to survive than those going to less liquid markets.

Table 7.2 below shows the key results for different values of  $ha$ .

$ha$	No. markets	% on one square	Turnover Ratio
0	1.00 (0.00)	100.00 (0.00)	1.00
0.1	1.55 (0.59)	98.90 (4.09)	0.99
0.5	3.65 (1.19)	92.29 (12.02)	0.94
1.0	3.65 (1.28)	78.30 (18.52)	0.93
2.0	4.50 (1.02)	65.01 (14.28)	0.93
4.0	4.75 (1.37)	56.24 (18.24)	0.95
6.0	4.80 (1.21)	53.18 (16.84)	0.96
8.0	4.85 (1.01)	48.86 (11.42)	0.96
10.0	4.70 (1.31)	57.13 (21.27)	0.97
$\infty$	5.15 (1.68)	51.24 (15.79)	0.98
No RL	11.50 (1.16)	20.44 (4.05)	0.98

Table 7.2: A summary of results for ‘habituation’ experiments based on the first model. Data are mean values for the last 100 rounds of 20 simulations, for each value of  $ha$  (standard deviations in parentheses).  $ha$  refers to the habituation parameter added to the weight of an agent’s target location (if there is one), which does not decay over time (note in the default simulations,  $ha = 0$ , which is represented in the top row). ‘No. markets’ refers to the mean number of market locations. ‘% on one square’ refers to the percentage of transactions on the square with the most transactions. The turnover ratio is the volume of actual transactions divided by the would-be volume at the market clearing price. The final row (‘No RL’) refers to a set of experiments when reinforcement learning is ‘switched off’: here, agents only ever visit the location where they first transacted during a random walk.

The table shows that as  $ha \rightarrow 0$ , the results mimic those of the default simulations, when  $ha = 0$  (this is not surprising). Moreover, as  $ha$  increases, the agents increasing

visit the first location they transacted at or heard about (for the whole of their lives). In this latter case, the mean number of market locations over 20 simulation runs was 5.15 (with a standard deviation of 1.68). This seems surprisingly low given that in these simulations, the total population peaks at approximately 43 agents who only ever transact at one location - this is explained below.

In all of the experiments, the turnover ratio was high enough for (surviving) agents to specialise and, eventually, to bear children, i.e., despite there being multiple markets, the agents sold enough of their resources to be sufficiently confident in specialising. To help visualise this, Fig. 7.13 below shows the supply and demand charts in a typical round when  $ha = \infty$ , as well as the turnover / mean price observed.




Figure 7.13: Supply and demand curves and total transactions in Round 999 of a typical simulation when  $ha = \infty$ . See Fig. 7.5 above for a description of the data. On average the ratio of actual transactions to the market clearing volume during rounds 900 - 999 was 0.98 over 20 simulations.

It is tempting to conclude from these results that the introduction of habituation makes little difference in the emergence of the markets; however, these data hide the fact that more agents died in these experiments than in the default simulations. For example, when reinforcement learning was switched off, a mean of 7.3 agents died between rounds 900 and 999 across 20 simulations (with a standard deviation of 2.1) whereas in the default simulations the equivalent numbers were 4.9 and 1.1, respectively.

The specific mechanism that explains these deaths was that, *ceteris paribus*, an agent was more likely to survive if it frequented a market with more agents attending than less well attended markets. There are two issues at play. The first and most obvious is that an agent with an excess of one resource post-foraging is more likely to find someone to trade with at larger markets - this helped it survive.

The second issue was the problem of small markets with an odd number of specialised agents. Consider a situation in which two of the original agents had become specialised (in different resources) and a new, non-specialised agent finds the location where they trade. The new agent would specialise in one of the two resources and, assuming no other agents find the market, this will lead to an imbalance in the resources flowing onto the market which has only 3 agents attending. Suppose now the market has one specialist in Resource A and two specialists in Resource B.

In the mid to late parts of the simulations, the Resource A specialist will not be able to collect enough of this resource for 3 agents to survive (there are approximately 43 agents competing for resources). From the new agent's point of view there will be a swing in price against its specialised resource which is, ultimately, more detrimental than the productivity benefits of specialisation. One of the two Resource B specialists will then perish: it will almost always be the newer agent because the older agent will have accumulated a larger stock of resources over the life of the simulation.

In some simulations we observe the two mature, fully specialised agents also dying at smaller markets because of the volatility caused by new agents and the fact that prices are determined by the agents' *stocks* of resources (not the flow onto the market). The latter point means that the low price of Resource B in the above example can be maintained well after the new agent has died: if the Resource B specialist has low levels of resources at this point, it might also die. If this happens then the Resource A specialist will have nobody to trade with and it too will die (even if the agent switches to foraging for its non-specialised resource, its low level of productivity means it will be unable to survive).

For larger markets, this problem is less likely because they are more liquid and because there is a greater chance of new agents joining that market than smaller markets (the probability that a new agent finds out about a market is proportional to the number of agents transacting at that location).

Hence, agents are more likely to survive if they join the larger markets.

One of the curious implications of this mechanism is that it is a weak form of symmetry breaking: the death of agents at less liquid markets gives 'room' for new agents to be born who are more likely to be informed about the more liquid markets. This is also related to weak downward causation / effects (Campbell, 1974), as discussed in Hodgson and Knudsen (2004). We should note, however, that markets with more than 3 agents (approximately) tended to survive so this symmetry breaking effect was limited.

#### 7.6.1.1 Efficiency and Equality

In the default simulations we observed that agents closer to the single market that emerges benefit from their proximity to that market. They profit from intermediation:

they obtain good prices from new arrivals to the market (especially when agents are specialised) and then sell these to other agents. Given this benefits some agents more than others, the institution should be viewed as a solution to a non-pure coordination situation.

Having a single market which all agents attend, however, maximizes allocative efficiency of the whole system (relative to multiple markets). Approximate speaking, the agents get closest to what they would have bought and sold if they transacted at the market clearing price at one market.

A single market, therefore, is allocatively efficient but it benefits the agents unequally.

When multiple markets emerge due to habituation making agents more conservative in their choice of target location, this reduces allocative efficiency (each market sees different prices) but in general the outcome is more equal because, on the whole, agents live closer to their chosen market.

### 7.6.1.2 Agent Level Data: Reinforcement Learning and Habituation

When  $ha$  is relatively low and the data is disaggregated, we find some interesting results.

The total weight of each target location in an agent's memory is made up of two values: that attributable to (i) reinforcement learning, and (ii) habituation. Fig. 7.14 below shows the contributions of both for a typical agent when  $ha = 0.5$  for the first 200 rounds of a typical simulation. Fig. 7.14 shows the same data for 1,000 rounds. These figures are reproduced from Chapter 6 for convenience.

The agent represented in Fig. 7.14 visited the location depicted from Round 49 onwards.

The figure indicates that the emergence of the market to which the agent travels occurred mostly as a result of reinforcement learning. In Round 100, the weight corresponding to this location was 99.6, which can be decomposed into 68.6 for reinforcement learning and 31.0 for habituation.

As the simulation progressed, however, the reinforcement learning component of the weight stabilized at about 100 and the habituation component continued to increase. This stabilization was a result of agents adding to the reinforcement learning contribution by transacting but this was offset by memory decay (which depreciated the contribution by 20% after each round). Habituation values, however, did not decay.

By the end of Round 1,000, the weight for this location in the agent's memories was 530.6 (112.6 attributable to reinforcement learning and 418 to habituation).

In general, this tells us that for lower values of  $ha$ , market institutions emerge as a result of reinforcement learning (and the co-adaptation of the agents' mental models). Later



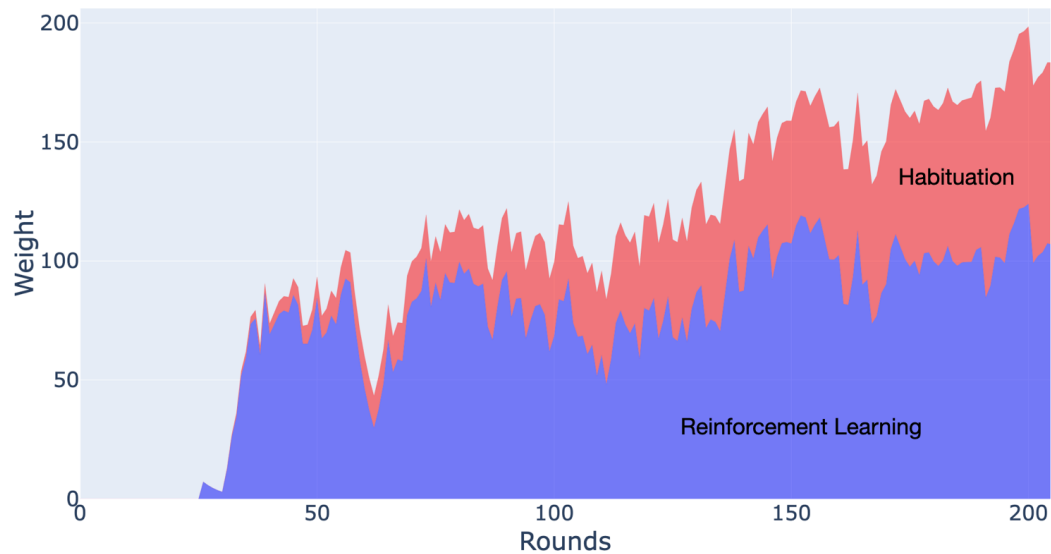


Figure 7.14: A time series of the contributions to a target location’s weight in memory over the first 200 rounds of a typical simulation (for one agent). The two ‘areas’ denote contributions to the total weight over time: the dark blue area shows the contribution of reinforcement learning; and the dark red area shows that of habituation. This chart shows how the agent learned to visit the location depicted as a result of reinforcement learning. Reproduced from Chapter 6 for convenience.

on, however, habituation dominates. This combination is intuitively appealing: agents co-learn to target a specific location and then, over time, this becomes embedded in the agents’ mental models.

## 7.6.2 Reinforcement Learning Turned Off

In the above experiments, agents combine reinforcement learning and habituation. What happens if reinforcement learning is ‘switched off’? This would mean that agents choose target locations as a result of habituation only. In practical terms, agents will wander around the grid in a random walk until they bump into other agents: communication from other agents is viewed as part of learning so this is also switched off.

The bottom row of Table 7.2 above shows the main results. We can see there are many more markets in these simulations (11.50 on average) although the turnover ratio remains high nonetheless.

The main difference between these simulations and those when  $ha = \infty$  (when reinforcement learning is ‘on’ - the penultimate row of Table 7.2) concerns communication. In the former, agents mostly learn about market locations from other agents: they visit a location in the round after learning about it, and always return to it thereafter. In the latter, this cannot happen: agents spend more time walking around randomly and this makes it more likely they will bump into other agents (without knowledge of a

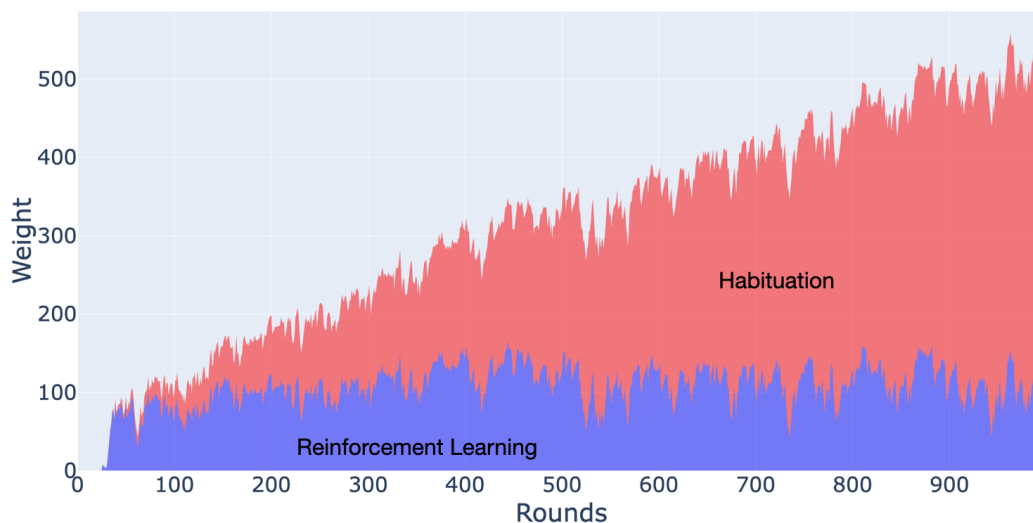


Figure 7.15: The same time series as shown in Fig. 7.14 above but over 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the weight of this location over time. Reproduced from Chapter 6 for convenience.

market) and form their own market. Hence there are significantly more markets when reinforcement learning is switched off.

## 7.7 Discussion of Results

A summary of the results presented in this chapter was provided in Chapter 6 and these will not be repeated here. Moreover, a detailed discussion of how the results correspond to the research questions is left to the concluding chapter. Here we focus on two topics: first, whether the markets observed in the default simulations meet our definition of organic institutions<sup>12</sup> (Section 7.7.1); and, second, whether we should consider the markets seen in the habituation experiments as organic institutions (Section 7.7.2).

### 7.7.1 Organic (Unplanned) Institutions

In the Introduction we adopted the definition of institutions as “durable systems of established and embedded social rules that structure social interactions.” (Hodgson, 2006a, p. 13). Organic institutions meet this definition but are unplanned.

<sup>12</sup>See Section 1.4.1 in the Introduction.

The single market that emerges in the default simulations appears to fit our definition but let us look at this in more detail:

- They are ‘durable’ in part because “they can usefully create stable expectations of the behavior of others” (Hodgson, 2006a, p. 2). This is clear in the simulation results: the single location in memory, which can be interpreted as an expectation, is consistent with other agents targetting the same location, i.e., there is persistent consistency between expectations and agents’ actions.
- The institutions can be viewed as ‘systems’ made up of, *inter alia*, internal (mental) representations and physical locations on the torus.
- They are clearly ‘established’ across and ‘embedded’ within the population of agents via their mental models. A single market can be thought of as an ‘attractor basin’ with steep sides.
- It should also be clear that the markets ‘structure social interactions’ in that they provide a location at which agents can trade.
- The single markets appear consistent with the definition of a ‘rule’, stated in the Introduction. We consider this in more detail below.

In terms of *organic* institutions, the question is whether the single markets are unplanned or not. It should be clear that they are not: agents are endowed with a limited ability to find other agents and to communicate; they could not directly plan to all meet at the same location.

## Rules

Are the market locations consistent with Hodgson’s (2006a) definition of rules?

In the Introduction we adopted the following definition: “The term *rule* is broadly understood as a socially transmitted and customary normative injunction or immanently normative disposition, that in circumstance  $X$  do  $Y$ .” (Hodgson, 2006a, p. 3, emphasis included).

The first question here is whether the markets are **socially transmitted**, which “means that the replication of such rules depends upon a developed social culture and some use of language.” (ibid). In considering this, let us differentiate between emergence and perpetuation of the single markets.

‘Social transmission’ appears to play some role in the emergence of the institutions but this occurs largely through co-adaptive reinforcement learning. The co-adaptive aspect suggests that ‘socially constructed’ is perhaps a more accurate term to describe this emergence than ‘socially transmitted’.

However, when we relax the assumption of agents having infinite lives (subject to resource consumption), which is done in the next chapter, it becomes clear that the replication of the institution occurs through communication from one generation to the next.

The second question is whether the single markets are equivalent to **immanently normative dispositions**<sup>13</sup>. It should be clear that the markets exist “immanently” and as “dispositions” in the agents’ mental models. The question is whether they can be categorised as “normative”.

The Oxford English Dictionary (OED) defines norms as “a standard or pattern of social behaviour that is accepted in or expected of a group.” (OED, 2021, Definition I1b). The reference here to “expected of a group” makes the market locations consistent with “normative” in Hodgson’s definition of rules: when a single market emerges, it becomes the sole location in each agent’s memory and this can be interpreted as an expectation that other agents will target and visit that location. All the agents expect it of all the other agents.

The third question is whether the agents are disposed in such a way that **in circumstances  $X$  do  $Y$**  (and not  $Y^*$ ). This is clearly true when the agents have a single market in memory. The “not  $Y^*$ ” aspect is more questionable but we can consider this pragmatically for a human agent who would be aware of there being 2,500 potential locations to visit even if one dominates its memories.

We conclude, therefore, that the single markets observed in the default simulations meet the definition of organic (unplanned) institutions we have adopted in this thesis fairly comfortably.

### 7.7.2 Organic Institutions with Habituation?

Do the markets observed when habituation is included in the mental models also meet our definition of organic institutions? To consider this, let us start from the ‘habituation only’ results and then add reinforcement learning.

In the ‘habituation only’ simulations we saw that agents can only form ‘markets’ by bumping into other agents, and they always visit the same location for the rest of their lives. Given these points, it is difficult to interpret these markets as either ‘unplanned’ or ‘surprising’. This means they are not organic institutions (and they do not meet Ullmann-Margalit’s (1978) third ‘surprise’ characteristic of invisible hand explanations).

When we add reinforcement learning we observe symmetry breaking; and when we reduce  $ha$  from  $\infty$  to zero, we observe more symmetry breaking. This is born out by the results

---

<sup>13</sup>This phrase seems more relevant than “customary normative injunction” in Hodgson’s definition of rules.

in Table 7.2 above: even when  $ha = \infty$  we see that in the last 100 rounds of simulations, slightly more than half of transactions occur on a single grid square on average. The reason for this appears to be the weak form of symmetry breaking identified above (when agents at smaller markets are more likely to die and new agents are more likely to be told of larger markets).

As  $ha \rightarrow 0$ , we are more likely to observe the symmetry breaking identified in the default simulations, when agents at smaller markets are more likely to switch their ‘allegiance’ than those at larger markets.

Overall, therefore, it appears the institutions we observe that are ‘unplanned’ and ‘surprising’ are this way because of symmetry breaking, which itself results mainly from reinforcement learning. However, this is not fully inconsistent with habituation: we observe reinforcement learning and symmetry breaking when  $ha$  is weak.



# Chapter 8

## Market Emergence Model: Exploring the Parameter Space

Negative results are just what I want. They're just as valuable to me as positive results. I can never find the thing that does the job best until I find the ones that don't.

– Thomas A. Edison

Fifteen parameters were adjusted in order to understand: (1) their impact on the default simulation results presented in the previous chapter; and (2) the conditions under which markets emerge and when they do not. A summary of the results is provided immediately below and more detail is included in Appendix B.

### 8.1 Summary of the Results

- **Memory decay:** in the default simulations, the weights of locations in agents' memories declined by 20% between rounds ( $m_{dec} = 0.2$ ). **The default simulation results were replicated if  $m_{dec} < 0.95$ .** For very low values (approximately  $m_{dec} < 0.02$ ), symmetry breaking was slowed because agents retained memories of transaction locations for longer<sup>1</sup>; and for very high values (approximately  $m_{dec} > 0.95$ ), agents were too forgetful so the results of the second null experiment (Section 7.3) were replicated (agents wandered around the grid randomly).

---

<sup>1</sup>This diluted the impact of information about any larger markets.

- **Population density:** there were  $50^2$  grid squares and 25 initial agents in the default simulations, which meant these simulations started with a density of  $10^2$  grid squares per agent. This density was adjusted by changing the grid dimensions while keeping the number of initial agents fixed at 25 (the time periods in the trading phase were scaled *pro rata* so agents could still access the whole grid). **The default simulation results were replicated if this density was approximately less than  $100^2$  grid squares per agent.** For higher values, the agents rarely found each other during their random walks.
- **Probability of Communicating:** the probability that any pair of agents communicated at the end of each round in the default simulations was 1%. **For the results of the default simulations to be replicated this probability had to exceed zero.** If communication was eliminated, markets still emerged, agents specialised, and the population increased to approximately 43 agents; however, approximately 10 markets emerged on average because symmetry breaking was prevented by the lack of communication.
- **Transacting on the Way to Target:** in the default simulations, when agents selected a target location from memory at the beginning of the trading phase, they would transact only after reaching that target. **Allowing agents to trade on the way to this target made no material difference to the results.** Across 20 simulations we observed that a mean of 83% of transactions occurred on one square in the last 100 rounds, with the remainder split between nearby squares. In fact, the market was more like a small area on the grid: as most of the agents converged on the ‘main’ market square, they would observe other agents and sometimes transact<sup>2</sup>.
- **Speed of Foraging Skill Change:** the speed of adjustment parameter ( $t$ ) in Equation 7.3 was set at 0.01 in the default simulations. **For the default simulations results to be replicated,  $t > 0$ .** If  $t = 0$  the third null scenario (Section 7.4) was replicated because agents could not specialise and the population stabilized at approximately 15-16 agents. Positive values of  $t$  meant agents specialised eventually and the population ultimately reached (approximately) 43 agents.
- **Randomizing Agents’ Home Locations:** in the default simulations, agents’ home locations were positioned on the grid to maximize sparsity. **If agent home locations were chosen randomly this made no difference to the results.** The only marginal change was that markets tended to emerge slightly more quickly if homes happened to be clustered: this meant agents were more likely to bump in to each during their random walks. This also meant the locations of the emergent markets were marginally more predictable than if home sparsity was maximized.

---

<sup>2</sup>Also, a small proportion of the agents would typically target these ‘satellite’ squares in each round.



- **Accuracy of Foraging Yield Estimates.** In the default simulations Equation 7.4 was used to estimate foraging yields in the next round when agents adjusted their foraging strategies at the end of each round. In this equation,  $\epsilon$  was an error term, drawn from a normal distribution with mean of zero and standard deviation ( $\sigma$ ) of 0.1. **We found that the simulation results were replicated if  $\sigma < 0.7$**  (approximately). For higher values, changes to the agents' foraging strategies were in effect randomized so they never specialised.
- **Weight of Other Agents' Transaction Locations in Memory ( $\beta$ ).** In the default simulations,  $\beta = 0.5$ . **For these simulation results to be replicated,  $\beta > 0.06$ .** When  $\beta = 0$ , the results mimicked the simulations when there was no communication between agents (both scenarios meant other agents' interactions were ignored). For low values of  $\beta$ , symmetry breaking was slowed down (a single market emerged if  $\beta > 0.06$ , approximately).
- **Travel Distance:** in the default simulations, agents could travel 25 grid squares during the trading phase, which meant they could access the whole torus. There were nine grid squares between neighbouring initial agents which meant this travel distance had to be at least five grid squares for them to have any possibility of meeting at all. We found that **the default simulation results were replicated if this travel distance was five grid squares or more; however, fewer markets emerged the larger this distance was.**
- **Selection of Grid Targets from Memory - Winner Takes All:** In the default simulations, agents used a 'Roulette Wheel' approach to select target locations from memory (Section 7.1.4). **An alternative 'Winner Takes All' algorithm<sup>3</sup> had no impact on the results** other than very slightly weakening the process of symmetry breaking (agents remained more loyal to the first proto-markets they visited).
- **Starting Resources:** in the default simulations, agents were endowed with two resources, each drawn from a normal distribution with mean 50 and standard deviation of 5. **The mean starting resource value had to exceed approximately 15 units for the default simulation results to be replicated.** Agent populations tended to survive even with lower starting values but typically the agent population would plateau at less than 43 agents (children were more likely to die before specialising if they were born with fewer resources).
- **Fountain Resources ( $L$ ):** The resource fountains began each round with a stock of 50 units each ( $L = 50$ ). **The default simulation results were replicated if  $L > 7$ ,** albeit with agent populations that plateaued at approximately  $0.85 \times L$ .

---

<sup>3</sup>This is when the location with the highest weight in memory was selected rather than selection probability being proportional to weight.

- **Geographically Locating the Resource Fountains:** in the default simulations, the resource fountains were not located on the grid. Here, the fountains were placed on the grid in random locations and agents started the trading phase at the last fountain they foraged from. **This had no impact on the results** with the exception that markets emerged more quickly: one was created at each of the fountain locations. Symmetry breaking still occurred so, ultimately, one market prevailed.
- **Limiting Agents' Life Spans:** agents lived forever in the default simulation provided they did not starve to death. When we gave the agents a limited life span (equal for each agent<sup>4</sup>) we found **the default simulation results were replicated if this age limit exceeded 400 rounds** albeit with the agent population plateauing at less than 43 agents. We observe the markets being replicated across generations.
- **Foraging Skill Acquisition: Linear Approach:** in the default simulations an adapted logistic equation (7.3) was used to adjust the agents' foraging skills. **Taking a linear approach to skill change<sup>5</sup> had no impact on the results.**

---

<sup>4</sup>Agents at instantiation began with an age drawn from a uniform distribution with a lower bound of zero and an upper bound of the age limit.

<sup>5</sup>We used  $\Delta p_i^j = t.[w_i^j - d/x]$  instead of Equation 7.3.

# Chapter 9

## Property Rights Model: Introduction

Freedom and Property Rights are inseparable. You can't have one without the other.

– George Washington

In the Market Emergence Model we assume agents respect others' property. In reality, people might not acknowledge resources as belonging to others and/or they might attempt to take them, i.e., property rights might not be recognised or certain.

Furthermore, it is generally understood by economists that property rights are a necessary condition for markets to work (e.g., [North, 2005](#)). Similarly, [Ostrom \(2005\)](#) lists the many conditions required for markets to operate and some of these are about the respect and enforcement of property rights.

In this and the next four chapters we will explore the conditions under which property rights might emerge as organic institutions, including how they might influence the emergence of markets. This work dovetails with the first model described in [Chapter 7](#) but, here, we adjust the model so agents can steal or trade. Which they choose depends on their propensity to respect other agents' property.

*Overview of Chapters 9 to 12*

This chapter describes the adjustments made to the original model.

The next chapter (10) sets out the results of the null and default simulations<sup>1</sup>.

Chapter 11 explores the parameter space to see whether and how the default simulation results are sensitive to the model's parameters. Ten parameters are adjusted and the simulation results analysed. Note that Chapter 11 contains a summary of the results: Appendix C contains much more detailed analyses.

In Chapter 12 we attempt to apply legal rules to the agents in each of the three scenarios set out in Appendix D. A fourth scenario is added, which is taken from Chapter 11 (this is when the cost of fighting is below the threshold above which property rights emerge). In all four scenarios, property rights never emerge endogenously so the question we consider is whether legal rules might catalyse these rights.

### *Overview of this Chapter*

Several changes are made to the original model from Chapter 7, the most significant of which concerns the process of interaction between two agents on the grid in the interaction phase of each round. Foraging, which takes place in the first part, remains unchanged.

This chapter first discusses the introduction of a 'propensity to steal' and a 'propensity to defend' (Section 9.1)<sup>2</sup>. The following section (9.2) outlines the changes made to bilateral agent interaction. After that, we consider how the new version of the model differs from conventional games in game theory (Section 9.3). Section 9.4 explains how the agents' propensities to steal and defend are adjusted, which is a form of reinforcement learning. In Section 9.5 we examine how agents move around the grid - in this version of the model they avoid agents if they expect to lose out from any interaction. Following this, changes to the end-of-round communication processes are described (Section 9.6). Finally, some alterations to the default parameters in the original model had to be made, which are explained in Section 9.7.

## 9.1 Propensities to Steal and Defend

In the new version of the model, agents are given two additional state variables: a propensity to steal ( $P^S$ ) and a propensity to defend ( $P^D$ ). These propensities change over an agent's lifetime as a result of reinforcement learning. Note that habituation is added to the agents' mental models in experiments set out in the next chapter.

---

<sup>1</sup>As with the original model, we define a set of 'default parameters' that, as stated earlier, are not meant to represent a 'correct' set of parameters.

<sup>2</sup>The latter is the propensity of an agent to defend its resources from theft but for the purposes of brevity we refer simply to a 'propensity to defend'.

In a specific bilateral interaction, agent  $i$ 's probabilities of stealing and defending their resources ( $P_i^{\text{steal}}$  and  $P_i^{\text{defend}}$  respectively) are derived by truncating their corresponding propensities in the following way:

$$P_i^{\text{steal}} = \begin{cases} 1 & \text{when } P_i^S > 1 \\ P_i^S & \text{when } 0 \leq P_i^S \leq 1 \\ 0 & \text{when } P_i^S < 0 \end{cases} \quad (9.1)$$

and, equivalently:

$$P_i^{\text{defend}} = \begin{cases} 1 & \text{when } P_i^D > 1 \\ P_i^D & \text{when } 0 \leq P_i^D \leq 1 \\ 0 & \text{when } P_i^D < 0 \end{cases} \quad (9.2)$$

In terms of the model, the agents' propensities are allowed to vary over time according to their experiences: probabilities of stealing and defending are derived from these propensities only when the agents interact. This truncation does not alter the propensities, only experience does.

Note that in this thesis we define 'property rights' (or 'respect for property') as when an agent has a zero probability of stealing in a bilateral interaction ( $P_i^{\text{steal}} = 0$ ), which is when  $P_i^S \leq 0$ .

#### *Rationale for Using Propensities*

The process by which the agents interact in the new version of the model is explained in more detail in the next sub-section. Here we can state that the agents face various forms of uncertainty, including information about other agents' resource holdings and their propensities to steal and defend. In addition, we assume the agents are cognitively limited relative to the large range of interaction types they face (we will see in the default simulations that the agents faced 363 unique pay-off structures over 20 simulations).

In deciding how the agents handle this uncertainty in the new version of the model, inspiration was taken from the various concepts in the social sciences, including habits and habituation, propensities, heuristics (both as an adjective and a noun), and social norms. The idea of propensities that evolve with experience is a relatively simple approach that respects these various concepts.

Furthermore, the approach fits with the probabilistic framing of the EMIL models discussed in Chapter 5. Also, it is broadly consistent with the use of classifier systems, e.g.,

in [Holland \(1975\)](#), [Holland et al \(1986\)](#), [Marimon et al. \(1990\)](#), [Arthur \(1994\)](#), [Kirman and Vriend \(2000\)](#), [Vriend \(2002\)](#), and [Kirman \(2011\)](#).

As a final note, we should not confuse this propensities approach with the idea of mixed strategies in the context of substantive rationality, which is when agents assign probabilities to strategic choices. In that approach, two agents calculate probabilities that optimise the likely outcome of a game. This is not the same as assigning propensities to agents' behaviour.

Let us now look at the adjustments made to the agents' interaction.

## 9.2 Adjustments to Agent Interaction

As in the original model, each agent selects a grid target at the start of the interaction phase - these are the squares the agents move towards.

However, as we will see below, agents are now able to avoid other agents if they believe an interaction might be detrimental (e.g., if an agent expects to be 'mugged'). For now, however, let us assume two agents interact.

### *Bilateral Interaction*

In an interaction, each agent first decides if it will attempt to steal their counterpart's resources, or attempt to trade. For each agent, this is done by using the agent's probability of stealing ( $P^{\text{steal}}$ )<sup>3</sup>. With probability  $P_i^{\text{steal}}$ , Agent  $i$  will attempt to steal its counterpart's resources. With probability  $1 - P_i^{\text{steal}}$ , the agent will attempt to trade.

Both agents signal to each other their decision.

If both agents wish to trade then they follow exactly the same process as in the original model (Section [7.1.4](#)).

If both agents wish to steal then they enter in to a conflict. This is described in more detail below.

If one agent wishes to steal and the other to trade then this second agent has to decide whether to defend its resources or not. To make this decision, we follow the same process as with the decision to steal but now we use the agent's probability of defending its resources ( $P^{\text{defend}}$ )<sup>4</sup>. Note that if the second agent chooses to acquiesce, we refer to this as a 'mugging'.

---

<sup>3</sup>Recall that this is the agent's propensity to steal truncated by 0 and 1.

<sup>4</sup>Again, this is the agent's propensity bounded by 0 and 1.

This framework creates six different scenarios, which are represented diagrammatically in Fig. 9.1 below. The interaction is similar to a conventional  $2 \times 2$  game but with the addition of choices to be made in the second and third scenarios.

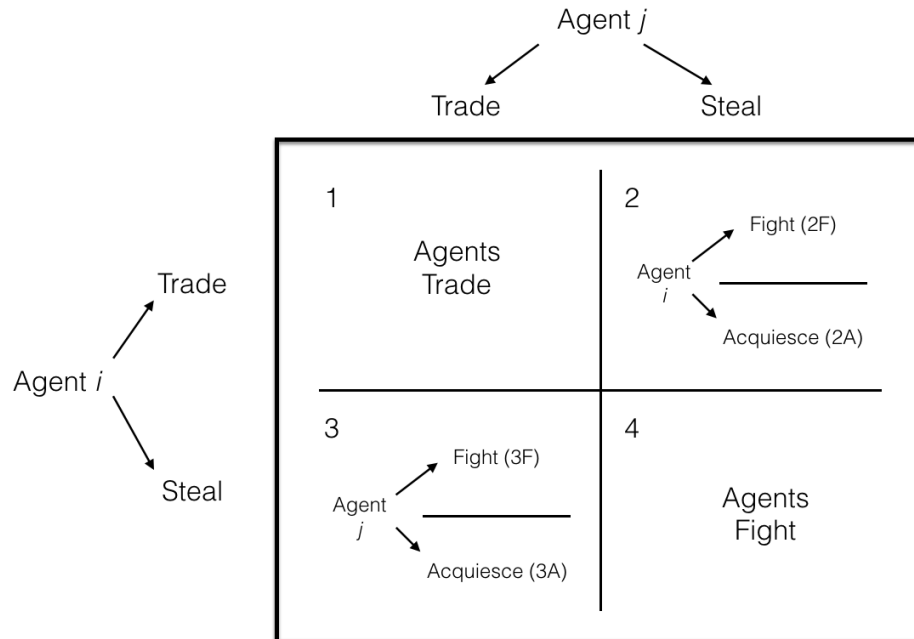


Figure 9.1: Overview of agent interaction in the second model. Agents initially decide whether to steal or trade. If one of them wishes to steal and the other trade then this second agent has to decide whether to defend its resource or not. As a result there are six potential scenarios for each interaction.

At first blush, the idea that agents signal to each other their intention to steal might seem unrealistic. For example, Agent  $i$  might try to 'sucker' Agent  $j$  by stating an intention to trade but then attempt to steal (or vice versa). In this situation, we assume that its counterpart could react immediately to being suckered, which means  $i$ 's actions would simply amount to signalling attempted theft.

The interaction structure described above resembles the Hawk-Dove game. However, this particular game typically has one specific set of pay-offs for the agents: in the simulations run using this model we observed hundreds of unique pay-off types. This is discussed further below. Also, in the classic Hawk-Dove game, agents cannot choose to defend their resources.

The pay-offs in the scenarios are determined by two factors: first, the resources held in the agents' baskets; and, second, by a cost incurred if the agents fight. The winner of any fight receives all of the resources from the loser's basket (scenarios 2F, 3F, and 4), as do muggers in the case of a mugging (scenarios 2A and 3A). In addition, we assume that agents who fight also incur an injury cost that is reflected as a deduction of  $c$  from both resources in the agents' personal resource arrays ( $\mathbf{r}$ ). In the default simulations,

$c = 0.3$  (this is adjusted when we explore the parameter space in the Chapter 11 and Appendix C).

Furthermore, in the new default parameter set, the winner of any fight is determined by the toss of a coin, i.e., each agent has a 50% probability of winning. The winner acquires all of the resources from the loser's resource basket ( $\mathbf{b}$ ).

In Appendix D we will use different mechanisms for deciding a fight outcome, e.g., in one experiment the agents' wealth (from  $\mathbf{r}$ ) is used as a measure of 'power' in conflicts.

In the case of scenarios 2A and 3A (muggings), the acquiescing agent will give all of its resources to the stealing agent but neither agent will incur a fight cost because no fight occurred.

#### *Pay-Offs in Formal Terms*

For the purposes of clarity, let us express the pay-offs and the expected pay-offs for each agent in formal terms.

If we designate the two interacting agents as  $i$  and  $j$  then their personal resource arrays will be  $\mathbf{r}_i$  and  $\mathbf{r}_j$ , and their basket arrays  $\mathbf{b}_i$  and  $\mathbf{b}_j$ . Furthermore, the cost of fighting (for both resources) is denoted as  $\mathbf{c} = [c, c]$ .

For Scenario 1, as mentioned above, we follow the same process as in the original model: the agents agree a price equal to the geometric mean of their marginal rates of substitution, and then they exchange resources at this price (assuming both hold the appropriate resources in their baskets).

For scenarios 2F, 3F, and 4, if Agent  $i$  wins the ensuing fight then its personal resource array is adjusted as follows<sup>5</sup>:

$$\mathbf{r}_i^{new} = \mathbf{r}_i^{old} - \mathbf{c}$$

and it acquires  $j$ 's resources:

$$\mathbf{b}_i^{new} = \mathbf{b}_i^{old} + \mathbf{b}_j^{old}$$

Here, the superscript 'old' refers to values prior to the interaction and 'new' refers to afterwards.

In these same scenarios, if  $i$  loses, the change in its personal resource array is the same as above but it will no longer hold any resources in its basket:

<sup>5</sup>Consistent with matrix algebra notation, bold text refers to arrays / matrices: transformations are done to all the elements in the array or matrix.



$$b_i^{new} = 0$$

More generally, if the probability of either agent winning the fight is 0.5, then, *ex ante*, Agent  $i$ 's expected pay-off (the change in its personal resource and basket arrays) in each of these scenarios will be:

$$E(\text{pay-off}_i^{2F, 3F, 4}) = \frac{b_j^{old} - c}{2} + \frac{-b_i^{old} - c}{2} = \frac{b_j^{old} - b_i^{old}}{2} - c$$

For Scenario 2A, Agent  $j$  is the stealing agent and  $i$  acquiesces. Here, there is no change to either agent's personal resource arrays since there is no fight. Agent  $i$ 's basket array changes in the same way as above:

$$b_i^{new} = 0$$

and  $j$  acquires  $i$ 's resources:

$$b_j^{new} = b_j^{old} + b_i^{old}$$

In terms of expected pay-off, for Agent  $i$  this is simply the loss of its basket (no fight cost is incurred):

$$E(\text{pay-off}_i^{2A}) = -b_i^{old}$$

and for  $j$  it is, equivalently:

$$E(\text{pay-off}_j^{2A}) = b_i^{old}$$

In Scenario 3A,  $j$  acquiesces and  $i$  takes  $j$ 's resources. Again, the agents' personal resource arrays do not change. For  $i$ , its gain is:

$$E(\text{pay-off}_i^{3A}) = b_j^{old}$$

and  $j$  loses:

$$E(\text{pay-off}_j^{3A}) = -b_j^{old}$$

If we know both agents' probabilities of stealing and defending ( $P_i^{\text{steal}}$ ,  $P_i^{\text{defend}}$ ,  $P_j^{\text{steal}}$ ,  $P_j^{\text{defend}}$ ) we can also determine the expected (joint) probabilities of each of the scenarios in Fig. 9.1. This is shown in Fig. 9.2 below.

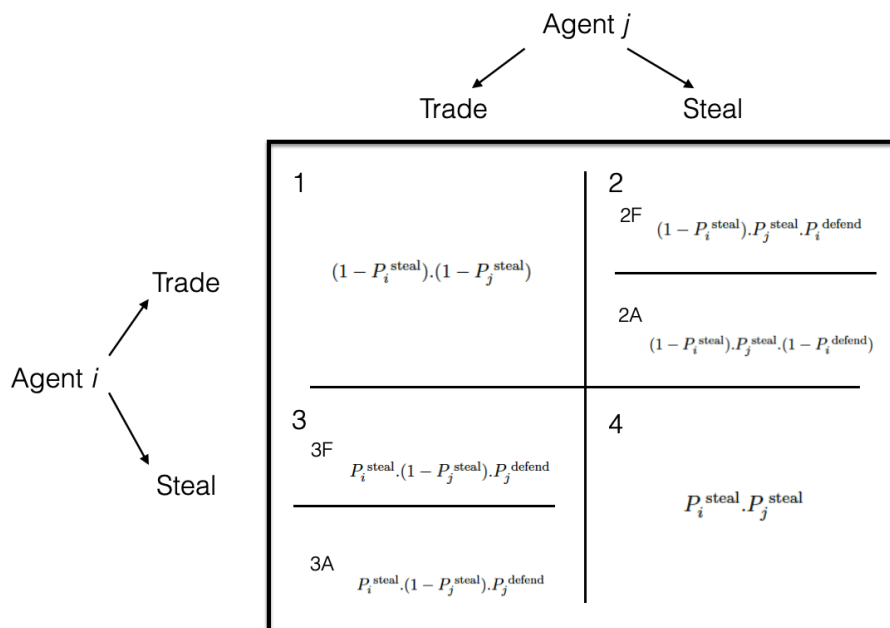


Figure 9.2: Probabilities of each of the six potential scenarios in Fig. 9.1. These are joint probabilities derived from both of the agents' probabilities of stealing and defending.

If we multiply these probabilities by their corresponding expected pay-offs as defined above then we can determine the expected pay-offs from any interaction for both agents<sup>6</sup>. This is useful when agents decide whether to avoid or interact with other agents, which is discussed further below.

### 9.3 Similarity with Conventional Games

As we noted briefly above, the bilateral interaction described above looks similar to the Hawk-Dove game: Scenario 1 appears equivalent to the 'Dove, Dove' outcome whereby the agents co-operate in a mutually advantageous way; and Scenario 4 is equivalent to the 'Hawk, Hawk' outcome.

However, we noted above that in the default scenario the agents interact with a wide variety of resource holdings and this gives rise to hundreds of different game types, only one of which is the Hawk-Dove game<sup>7</sup>. In addition, in the interaction architecture

<sup>6</sup>Note that Fig. 9.2 uses the agents' actual probabilities but the agents themselves have to estimate their counterpart's probabilities. This is described further below.

<sup>7</sup>The 6-scenario game described above could be reduced to a 4-quadrant game if we assume substantive rationality and if the agents use backward induction in scenarios 2 and 3.

described above, agents can defend their resources against would-be thieves, which is not the case in the Hawk-Dove game.

## 9.4 The Evolution of the Agents' Propensities to Steal and Defend

After any interaction, both agents learn from the changes to their basket arrays (through trade or theft) and any fight costs incurred: the net benefit (or cost) of the interaction is used to adjust the propensities to steal and defend of both agents. Consistent with [Erev and Roth \(1998\)](#), if an agent benefits from a particular strategic choice then it will be more likely to repeat this choice again in the future; and vice versa. For example, if an agent successfully steals from its counterpart, its propensity to steal is increased.

### 9.4.1 The Reduced Gain / Loss From Interactions and Adjustments

The immediate problem we face in updating these propensities is that the agents gain and lose multiple resources whereas  $P^S$  and  $P^D$  are scalars. This presents us with an incommensurability problem because we must map changes in arrays on to changes in numbers.

The solution we use to overcome this problem is to employ the agents' marginal rates of substitution (MRSs), which is a measure of relative value, to map from one resource to another. For example, if Agent  $i$  has a personal resource array of  $\mathbf{r}_i = [100, 50]$  then  $MRS_i^{AB} = 2$ , i.e., the agent will view Resource B as equivalent to 2 units of Resource A. These MRSs are used to map agents' gains and losses on to a value equivalent to its minimum resource holding (B, in this example).

We refer to this minimum-resource equivalent gain (loss) as the *reduced gain (loss)* to the agent, denoted by  $v_i$ .

For example, if the same agent has resources in its basket of  $\mathbf{b}_i = [b_i^A, b_i^B] = [1, 1]$  and is 'mugged', its minimum resource-equivalent loss will be:

$$v_i = (-b_i^A \times \frac{1}{MRS_i^{AB}}) - b_i^B = (-1 \times 0.5) - 1 = -1.5$$

The agent loses 1 unit of its minimum resource and 1 unit of the other, which it values as equivalent to 0.5 units of the minimum resource. Its total loss in units of its minimum resource holding is, therefore,  $-1.5$  units. This reduced value is used to adjust the agent's propensities to steal and defend, as described below.

In terms of information to learn from, an agent's own strategic choices, and its gains and losses, are not the only data available to it. There is an argument that Agent  $i$  should also use its counterpart's experience in the interaction to supplement its own learning<sup>8</sup>. This information about their counterparts is available to all agents. For example, if Agent  $i$  attempts to steal from Agent  $j$  who wants to trade but who then fights back, the benefit or loss from  $j$ 's decision is used to adjust  $i$ 's own propensities to steal and defend.

To incorporate Agent  $j$ 's pay-offs (reduced to  $v_j$ ) into  $i$ 's learning, we introduce a weight parameter,  $\beta$ . This is the weight  $i$  associates with  $v_j$  (relative to a weight of 1 it associates with  $v_i$ ). In the default parameter set,  $\beta = 0.5$ .

After  $v_i$  and  $v_j$  are calculated they are transformed in three ways. The resulting value is then used to adjust the agents' propensities to steal and defend. Let us now look at how this transformation is done.

### 9.4.2 Mapping $v_i$ and $v_j$ on to Propensities

The equations that result from the transformation of  $v_i$  and  $v_j$  (and which are used in the model to adjust  $P_i^S$  and  $P_i^D$ ) are shown in Table 9.1 below from the perspective of Agent  $i$  (i.e.,  $\Delta P_i^S$  and  $\Delta P_i^D$ ) in all of the six scenarios in Fig. 9.1. The equations for  $i$ 's counterpart ( $j$ ) are symmetrically the same.

Scenario	Agent $i$
1	$\Delta P_i^S = r \times (-v_i^* - \beta v_j^*)$ $\Delta P_i^D = 0$
2F	$\Delta P_i^S = r \times (-v_i^* + \beta v_j^*)$ $\Delta P_i^D = r \times v_{ii}^{**}$
2A	$\Delta P_i^S = r \times (-v_i^* + \beta v_j^*)$ $\Delta P_i^D = r \times -v_{ii}^{**}$
3F	$\Delta P_i^S = r \times (v_i^* - \beta v_j^*)$ $\Delta P_i^D = r \times \beta v_{ij}^{**}$
3A	$\Delta P_i^S = r \times (v_i^* - \beta v_j^*)$ $\Delta P_i^D = r \times -\beta v_{ij}^{**}$
4	$\Delta P_i^S = r \times (v_i^* + \beta v_j^*)$ $\Delta P_i^D = 0$

Table 9.1: Adjustments to agents' propensities to steal and defend following an interaction.

Notes:

$$v_i^* = \text{sign}(v_i) \times |v_i|^\delta$$

<sup>8</sup>This is equivalent to adding other agents' interaction locations to memory when deciding a target location, which was described in Section 7.1.4.

$$\begin{aligned}
v_j^* &= \text{sign}(v_j) \times |v_j|^\delta \\
v_{ii}^{**} &= \text{sign}(v_i - E_i^{Q2}(v_i)) \times |v_i - E_i^{Q2}(v_i)|^\delta \\
v_{ij}^{**} &= \text{sign}(v_j - E_i^{Q3}(v_j)) \times |v_j - E_i^{Q3}(v_j)|^\delta
\end{aligned}$$

The parameters and variables used in this table are defined and explained below.

There are three parts to this mapping from the agents' reduced pay-offs to changes in their propensities:

1. the use of expected pay-offs ( $E_i^{Q2}(v_i)$ ) when adjusting  $P^D$  (but not  $P^S$ );
2. the transformation of the resulting values by a 'cognitive coarseness' parameter,  $\delta$ , where  $0 \leq \delta \leq 1$ ; and
3. adjustment via a rate of change coefficient ( $r$ ).

Let us look at these three in turn. Section 9.4.2.4 below contains a detailed example.

#### 9.4.2.1 Expected Versus Absolute Pay-offs

At first blush it seems reasonable to use gross reduced pay-offs to determine  $\Delta P_i^S$  and  $\Delta P_i^D$ , to ensure  $P_i^S$  and  $P_i^D$  change monotonically with  $v_i$ . However, there is evidence in the psychology literature that responses to outcomes should be viewed relative to prior expectations of those outcomes, e.g., [Brickman and Campbell \(1971\)](#), [Easterlin \(1974\)](#), and [Rutledge et al \(2014\)](#). This would argue for Agent  $i$ 's propensities responding to  $v_i - E_i(v_i)$  rather than  $v_i$  alone, where  $E_i(v_i)$  is  $i$ 's expectation of  $v_i$ .

Unfortunately, there is no consensus in the literature on this topic, e.g., [Veenhoven \(1991\)](#) distinguished between happiness and contentment, and argues the evidence is that the former did appear to be relative to some benchmark (like expectations) but the latter did not.

The approach taken in this thesis is to consider both techniques, i.e., for  $\Delta P_i^S$  and  $\Delta P_i^D$  to be determined by  $v_i$  and  $v_i - E_i(v_i)$ . This is done by assuming a particular approach in the default simulations and then considering alternative techniques when we explore the parameter space (see Section C.7.3 of Appendix C in particular).

For the default simulations we assume that  $\Delta P_i^S$  is a function of  $v_i$  only, i.e., we use gross reduced values; but for  $\Delta P_i^D$  we assume this is a function of  $v_i - E_i(v_i)$ , i.e., the reduced gain / loss relative to  $i$ 's prior expectation of it. Moreover, we assume the expectation of  $v_i$  is used when adjusting  $P_i^D$  is that pertaining to Scenario 2 rather than the whole interaction (denoted  $E_i^{Q2}(v_i)$ ).

The argument for using relative outcomes when adjusting  $P_i^D$  and absolute outcomes for  $P_i^S$  is that in Scenario 2, neither  $v_i$  nor  $E_i^{Q2}(v_i)$  are influenced by the choices made by  $i$ 's

counterpart ( $j$ ). In that scenario there is no contingency:  $v_i$  is determined by  $i$ 's choice (to acquiesce or defend resources) and  $E_i^{Q2}(v_i)$  is driven by the  $i$ 's probability-weighted expected pay-offs.

By contrast, if we were to adjust  $P_i^S$  by a function of  $v_i$  relative to some expectation of  $v_i$ , we would have to use  $i$ 's expected outcome for the *whole* simulation (because its decision to trade or steal covers all six scenarios<sup>9</sup>). The actual outcome of the interaction ( $v_i$ ) would depend on  $j$ 's choices as would the resulting probability-weighted mean of  $i$ 's outcomes. For  $i$ , therefore, use of  $v_i - E_i(v_i)$  to adjust  $P_i^S$  would involve a considerable amount of uncertainty regarding  $j$ .

This argument is not intended to be a fully compelling rationale for the approach taken in the default simulations but it does point to reasonable ground for using different treatments for  $\Delta P_i^S$  and  $\Delta P_i^D$ . In any case, as mentioned above, we consider alternative approaches when we explore the parameter space.

#### 9.4.2.2 Cognitive Coarseness

We introduce the idea of limited cognition here by adopting a parameter ( $\delta$ ) that influences the impact of  $v_i$  and  $v_i - E_i(v_i)$  on  $\Delta P^S$  and  $\Delta P^D$ , respectively.

More specifically, we make  $\Delta P^S$  a function of  $v_i^\delta$  and  $\Delta P^D$  a function of  $(v_i - E_i(v_i))^\delta$ .

However, we also wish to keep the monotonicity of these values while adjusting their magnitudes by  $\delta$ . Therefore<sup>10</sup>:

$$v_i^* = \text{sign}(v_i) \times |v_i|^\delta$$

$$v_i^{**} = \text{sign}(v_i - E_i^{Q2}(v_i)) \times |v_i - E_i^{Q2}(v_i)|^\delta$$

where  $0 \leq \delta \leq 1$

This transformation means that if  $\delta = 1$ , the precise values of  $v_i$  and  $v_i - E_i^{Q2}(v_i)$  are mapped on to  $\Delta P^S$  and  $\Delta P^D$ ; and as  $\delta \rightarrow 0$ ,  $v_i^*$  and  $v_i^{**}$  tend to +/-1.

We can think of  $\delta = 1$  as when pay-offs have a direct 'fine-grained' (Gell-Mann and Hartle, 2007) impact on propensities; and we can think of  $\delta = 0$  as 'coarse-grained' (or 'lumpy') whereby the agents attach simple values of +/-1 ('good' or 'bad') to outcomes.

In the default parameter set we assume  $\delta = 0.5$ , i.e., we adjust  $v_i$  and  $v_i - E_i^{Q2}(v_i)$  by their square roots while retaining their signs (+ or -). We will look at the impact of changing this parameter when we explore the parameter space.

<sup>9</sup>whereas  $i$ 's decision to defend its resources is only relevant for Scenario 2.

<sup>10</sup>Note that  $\text{sign}(x) = +1$  for any value of  $x > 0$  and, equivalently,  $\text{sign}(x) = -1$  for any value of  $x < 0$  ( $\text{sign}(0)$  is irrelevant here because this value would be multiplied by zero in the two equations).

### 9.4.2.3 Rate of Change Coefficient

The final part of the transformation from  $v_i$  and  $v_j$  to the agents' propensities is the use of a coefficient ( $r$ ). In the default parameter set we fix  $r = 0.01$  but, once again, we will consider the impact of different values of this parameter when we explore the parameter space.

The equations that result from the above three transformations are shown in Fig. 9.1 above.

### 9.4.2.4 Example

For the purposes of clarity, let us look at an example we can apply to these equations.

Consider two agents ( $i$  and  $j$ ) with personal resource and basket arrays as follows:

$$\mathbf{r}_i = [105, 95]$$

$$\mathbf{b}_i = [1, 1]$$

$$\mathbf{r}_j = [95, 105]$$

$$\mathbf{b}_j = [1, 1]$$

Suppose, also, that  $r = 0.01$  (the speed of adjustment variable in the equations above),  $\beta = 0.5$ , and  $P_i^S = P_i^D = P_j^S = P_j^D = 0.5$ .

When we combine these agents' arrays we obtain marginal rates of substitution as follow:

$$MRS_i^{AB} = \frac{r_i^A + b_i^A}{r_i^B + b_i^B} = \frac{106}{96} = 1.104$$

$$MRS_j^{AB} = \frac{r_j^A + b_j^A}{r_j^B + b_j^B} = \frac{96}{106} = 0.906$$

If the agents trade, the price would be exactly 1<sup>11</sup> since each agent's resource holdings is the inverse of the other's.

Table 9.2 below shows the outcomes of the interaction between these two agents in the six scenarios and the corresponding changes to  $P_i^S$  and  $P_i^D$ .

<sup>11</sup>The transaction price is  $[MRS_i^{AB} \times MRS_j^{AB}]^{1/2} = 1$ .

Scenario	$v_i$	$\Delta P_i^S$	$\Delta P_i^D$
1	+0.0943	-0.0014	nil
2F - Agent wins	+1.3340	-0.0257	+0.0160
- Agent loses	-2.4774	+0.0314	-0.0111
2A	-1.9057	+0.0286	+0.0082
3F - Agent wins	+1.3340	+0.0257	-0.0056
- Agent loses	-2.4774	-0.0314	+0.0080
3A	+1.9057	+0.0286	+0.0041
4 - Agent wins	+1.3340	+0.0010	nil
- Agent loses	-2.4774	-0.0181	nil

Table 9.2: An example of changes in an instigating agent's propensities to steal and defend when both agents hold 1 unit of each resource in their baskets. The data is shown for every one of the six scenarios.

In Scenario 1, both agents sell (buy) the resource they have more (less) of, which results in a net positive reduced gain. The (small) gain from trading leads to a (very) small decline in the agents' propensities to steal, i.e., both agents are encouraged to trade again in the future. Neither agent has to decide whether or not to defend its resources so neither  $P_i^D$  nor  $P_j^D$  changed.

For Scenario 2,  $i$  attempts to trade but  $j$  tries to steal its resources. If  $i$  chooses to defend its resources (2F) and wins, the (reduced) gain in resources will exceed the (reduced) cost of fighting: the agent will be encouraged to trade<sup>12</sup> ( $\downarrow P_i^S$ ) and to defend its resources ( $\uparrow P_i^D$ ) in the future. However, if  $i$  chooses to defend its resources and *loses*, the reverse will be true ( $\uparrow P_i^S$  and  $\downarrow P_i^D$ ). The outcome if  $i$  acquiesces would be to encourage  $i$  not to trade in the future ( $\uparrow P_i^S$ ) and to defend its resources  $\uparrow P_i^D$ .

The outcomes in Scenario 3 are identical for those in Scenario 2 for  $i$ 's propensity to steal but with opposite signs:  $i$  tried to steal and not trade so in a sense learning is inverted. In this scenario,  $i$  would learn from  $j$  when it came to defending resources: if  $j$  loses the fight, both agents would see  $\downarrow P^D$  (and vice versa). Both would see their propensities to defend increase if  $j$  acquiesces.

In Scenario 4, both agents attempt to steal. The reduced gain / loss for  $i$  is the same as in scenarios 2F and 3F but the impact of these on the agents' propensities to steal is different: each agent's gains / losses act as a hedge for the other agent. For example, if  $i$  wins the fight,  $v_i$  will be +1.3340 but because  $\beta v_j$  is -1.2387, therefore  $v_i + \beta v_j$  is 0.0953<sup>13</sup>.

<sup>12</sup>or discouraged from stealing.

<sup>13</sup>hence  $r \times (v_i + \beta v_j) = 0.0010$ .



### 9.4.3 Propensities, Learning and Probability Limits

Given the equations in Table 9.1 above, which show how  $P^S$  and  $P^D$  change in the different scenarios, it is possible for the agents' propensities to exceed 1 or to decline below 0. At first blush it might seem simpler to apply a ceiling of 1 and a floor of 0 to both propensities and to use these as the agents' probabilities of stealing or defending their resources.

However, this would be problematic because any further learning at these limits would be asymmetric, e.g., if  $P_i^D = 1$ , Agent  $i$  would no longer learn from interactions that indicated defending its resources was preferable to acquiescing; but it would learn in the opposite situation<sup>14</sup>. This point is relevant for both the propensity to steal and defend at probabilities of 0 and 1.

Two approaches were considered for dealing with this asymmetry: (i) making changes in propensities a function of their levels such that they could not breach 0 or 1, e.g., with an adapted logistic equation<sup>15</sup>; and (ii) allowing the agents' propensities to exceed 1 and to decline below 0 but truncating these propensities to generate probabilities, as described in Section 9.1.

The second approach was chosen because: (i) agents could continue to learn even if the propensities were above 1 or below 0; and (ii) making propensity changes a function of their corresponding levels seems contrived, with no reasonable justification on theoretical grounds.

Next we take a step back and consider how agents move around the grid and how they decide which agents to interact with and which to avoid.

## 9.5 Moving Around the Grid: Interaction and Avoidance

In the first model, agents remain on their target square for the rest of the interaction phase after they reached it. In the second model, we have introduced the idea of agents stealing from each other so it seems reasonable to give the agents an ability to avoid others if they expect an interaction to be detrimental.

<sup>14</sup>If all of the agents' interactions are relatively small and agents learn on average it was preferable to defend their resources then their propensities will bounce along the ceiling of 1 so this asymmetry would hardly matter. However, we found it was often the case that resources would become very concentrated in the interaction phases, which meant that agents' propensities sometimes jumped by 0.2-0.3 in one transaction. Such large changes means this learning asymmetry is significant, e.g.,  $P_i^D$  might bump along the ceiling of 1 but then suddenly decline to 0.7.

<sup>15</sup>As an example of this logistical equation approach, in Scenario 2F the equation for  $\Delta P_i^S$  would change from  $\Delta P_i^S = r \cdot (-v_i + \beta v_j)$  to  $\Delta P_i^S = r \cdot (-v_i + \beta v_j) \cdot P_i^S (1 - P_i^S)$ .

In this section we summarise how this is done.

The key principle applied below is a simple one: agents will interact with those with whom they expect to benefit; and avoid agents they expect to lose from.

As with the original model, each agent could see other agents on its current grid location and also in adjacent (“king’s move”) squares. However, the model is adjusted to allow the agents to estimate the likely pay-offs of all potential interactions<sup>16</sup>.

Fig. 9.3 below depicts a typical situation on the grid. The agent whose turn it is to move (in red) is located in the middle, blue square and the neighbouring squares are shown in grey.

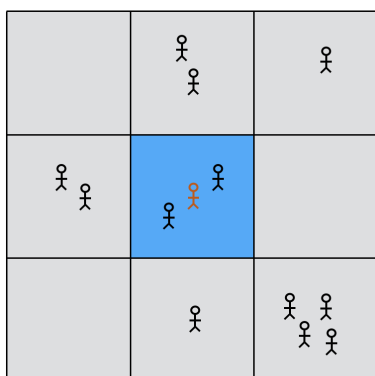


Figure 9.3: A diagram depicting how an agent evaluates interacting or moving on the grid in the second model. The  $3 \times 3$  grid shown here is a subset of the  $50 \times 50$  torus on which agents could move during the interaction phase. The agent (in red) initially evaluates the expected gains and losses from potential interactions on its current square (in blue). If there are no agents on that square or none of the interactions are deemed beneficial, the agent evaluates the gains and losses on the other squares within sight (the grey squares).

Agents take turns to act during the interaction phase (they can transact or move, but not both). When it is Agent  $i$ 's turn, it first looks to see if other agents are co-located on its current square. If this is the case, it evaluates the likely pay-offs from interacting with each of these agents<sup>17</sup> and chooses to interact with the agent associated with the highest expected reduced pay-off (if this is positive). This is a change from the original model where agents are allocated counterparts (located on the same square) randomly, if there is more than one.

<sup>16</sup>Here we assume that the agents have already arrived at their target squares so they can interact.

<sup>17</sup>Note that in forming expected pay-offs, the model calculates an accurate expectation for each of the six scenarios and adds an error term that is drawn from a normal distribution with mean zero and a standard deviation of 0.05 (this is done for each of the six pay-offs). This error represents a form of information uncertainty (it is adjusted when we explore the parameter space in Chapter 11). In addition, agents use their estimates of other agents' propensities to steal and defend in order to form a probability-weighted expected pay-off for each potential interaction.

If the highest expected reduced pay-off on its current square is zero or negative, the agent looks at all the neighbouring squares (it could see all the agents and the resources they hold).

The agent then calculates the mean expected pay-off from all of these neighbouring squares and moves to the square with the highest mean expected pay-off. If all such pay-offs are negative, the agent will move as far away as possible from the square with the lowest pay-off.

In general we saw in the simulations that Agent  $i$  is more likely to interact with another agent if this counterpart has:

- more resources in its basket (this was relevant whether Agent  $i$  is likely to steal or trade);
- a low estimated propensity to steal; and
- a low estimated propensity to defend.

These principles were valued regardless of Agent  $i$ 's resource holdings and its propensities to steal and defend. We see something similar in multi-agent iterated Prisoners' Dilemma games where both co-operators and defectors prefer to interact with co-operators.

The last point to note here is that in each of the time periods during the interaction phase (there were 50 in the default simulations), the order the agents take to act is randomized at the beginning of each time period<sup>18</sup>. This introduces variation, allowing some realistic interactions<sup>19</sup>.

To summarise this section, agents are able to form expectations about pay-offs from interacting with other agents in their immediate vicinity. If agents are co-located, the acting agent would interact with the agent from which it expected to benefit the most. Otherwise, the agent will look to its neighbouring squares and move to (away from) squares where the expected pay-off is most positive (negative).

---

<sup>18</sup>For example, Agent  $i$  might be the fifth agent to act in Time Period 4 but the twelfth in Time Period 5 and then the first in Time Period 6.

<sup>19</sup>This variation means, for example, that an aggressive 'mugging' agent (with high  $P^S$  and  $P^D$ ) might catch a non-aggressive, trading agent (with low  $P^S$  and  $P^D$ ) in an adjacent square if it acts twice in a row relative to the stationary non-aggressive agent. Equivalently, a non-aggressive agent might escape from an aggressive, mugging agent on its square if it acts twice in a row (it would be out of sight if it moves away twice).

## 9.6 End of Round Communications & Grid Target Selection

In the original model, the probability that any two agents communicate at the end of the round is 0.01. If they communicate, the agents exchanged information about all the locations of transactions they had been involved in during the round.

In the new default parameter set, the probability that any two agents **communicate** at the end of the round is increased to 0.2. The rationale for this is explained below.

Two further adjustments are made to communication in the model. The first is to add reputations. If two agents communicate then they will tell each other about which agents had traded with them, attempted to steal from them, and which fought back. They will do this regardless of whether or not they would interact with this other agent since they will only be exchanging information about other agents and not themselves. We can think of this as equivalent to gossip. Each agent maintains databases that store this data, and these are used to estimate other agents' propensities to steal and defend.

The second adjustment is that agents only exchange location information about transactions and fights they themselves had been involved in if *both* agents expect a future interaction to be beneficial. Agents will not provide this information to those they do not want to interact with: they will be more coy about information pertaining to themselves.

In terms of **grid target selection**, the agents use the information concerning transaction and fight locations to decide the grid squares they head toward in the interaction phase. The process is similar to that used in the original model (Section 7.1.5) but now the agents incorporate fight information.

In selecting a target, the guiding principle is that agents are attracted to locations where they expect to gain from interacting, either through theft or trade, and are deterred from locations where they expect to lose out.

Locations on the grid are each given weights: grid squares where the agent had transacted saw their weight in memory increased by 1 as did locations where the agent had fought and won. If the agent had fought and lost (or had been mugged) then the weight corresponding to that location fell by 1. Locations of fights the agent had heard about from others saw their weights increase by  $1/2$  if the agent expected to gain from interacting with both the agents who had fought (on average) and the weight was decreased by  $1/2$  if the agent expected to lose.

At the start of the interaction phase, each agent selects all of the squares with positive values and, as with the first model, uses a roulette wheel approach to choose their target (weighted by each square's positive value).

Finally, all of the location weights in memory decay by 20% in between rounds, consistent with the first model (this is true of positive and negative weights).

## 9.7 Changes to the Original Parameters

The introduction of fight costs in to the model means the environment is harsher for the agents in comparison with the original model. As a result, some of the original parameters are adjusted to make the environment less harsh in the new default parameter set (else most of the agents would simply die).

In the early parts of simulations that use the new default parameter set, there were typically fewer transactions than in the early stages of the original default simulations. The probability of communicating at the end of the round was raised to 0.2 to make it easier for agents to learn where other agents might assemble in the next round.

The resource endowments of the population at instantiation are increased from a mean of 50 units of each resource to 200 units; and the agents have to hold at least 300 units of each resource before they could sire children. Each child is born with a mean of 200 resource units, half of which are taken from each parent.

Finally, there are now 2,000 rounds in each simulation (this was extended when necessary), and, as before, there are 20 simulations in each set.

Now that we have described the changes to the original model, we can look at the results from a set of ‘null simulations’ and from simulations that used the new default parameters.



# Chapter 10

## Property Rights Model: Simulations and Discussion

Why don't you knock it off with them negative waves?

– *Kelly's Heroes*, Oddball

In this chapter we first analyse and discuss the results of four ‘null simulations’ (Section 10.1). As with the first model, these are designed to contextualise the results of the simulations that use the new default parameter set. Section 10.2 analyses and discusses the ‘default simulations’.

The third section (10.3) presents the results of experiments in which habituation is added to the agents’ mental models. In the default simulations, these change as a result of reinforcement learning only: here, we add habituation.

Section 10.4 concludes the chapter by considering whether the property rights that emerge fit our definition of organic institutions. This final section also considers the results in the context of generalised Darwinism (as discussed in Chapter 2).

### 10.1 Property Rights Model: Null Simulations

In this section we present the results of four null scenarios.

In the first null experiment (Section 10.1.1) we replicate the original model by setting  $P^S = 0$  for all agents. In this scenario,  $P^D$  is irrelevant because the agents always trade.

We found the results of the simulations that used the original default parameter set are approximately replicated; however, some of the changes to the original parameters, mentioned in the previous chapter, had an impact on the results. This experiment allows us to examine these changes.

In the second null scenario (Section 10.1.2) we do the opposite of the first: we fix  $P^S = 1$  for all agents, i.e., they only ever fight.

In the third null experiment (Section 10.1.3) we apply substantive rationality to all the interactions: the agents evaluate the pay-offs in each game and make their strategic choices. The agents' propensities to steal and defend are irrelevant in these experiments.

For the fourth null experiment (Section 10.1.4),  $P^S = P^D = 0.5$  for all agents, i.e., each agent has a 0.5 probability of trading or stealing and a 0.5 probability of defending its resources or acquiescing.

### 10.1.1 Null Scenario: Replicating the Original Model

Here we fix all the agents' propensities to steal at zero, i.e., none of the agents attempt to steal from each other.

The results are broadly similar to those seen with the original default parameter set. However, some of the modifications mentioned in the previous chapter affected the simulation outcomes. The main difference concerns the markets that emerge, which are now a collection - or area - of neighbouring squares rather than a single square. See Fig. 10.1 below, which shows a heatmap of transactions in the last 100 rounds of a typical simulation.

This change is due to the interaction mechanics described in Section 9.5 whereby agents move to adjacent squares from their target location if they deem it worthwhile. In the first model, agents stay on their target square and wait: if no other agents show up the agent would not trade and the weight of that particular square in memory would decline. The main result of agents looking and moving to adjacent squares is that, in general, agents maintain multiple squares with positive weights in memory.

It is worth noting that symmetry breaking still occurs if multiple market areas emerge: a single *area* ultimately dominated for the reasons identified in Chapter 7.

Despite this dispersion of transactions across a market area, the turnover ratio was sufficiently high (approximately 0.9) over a long enough period for the agents to be induced to specialise. The agents eventually had children and the total population increased to approximately 43 agents as in the first model's default simulations. See Fig. 10.2 below, which shows the total agent population in a typical simulation.






Figure 10.1: Heatmap of transactions in the last 100 rounds of a typical simulation when the agents' propensities to steal are fixed at 0. The results replicate those of the default simulations of the first model except, here, markets are made up of small areas of adjacent squares.




Figure 10.2: Total agent population over 5,000 of an extended (and typical) simulation when the agents' propensities to steal are fixed at 0. The population plateaus at approximately 43 agents, as in the original default simulations.

### 10.1.2 Null Scenario: All Agents Steal

If all the agents try to steal from each other in every interaction ( $P_i^S = 1$  for all  $i$ ) then the agent population always collapses<sup>1</sup>. Fig 10.3 shows the total agent population over a typical simulation.

On average, approximately two agents survived until the end of 2,000 rounds in all of the simulations. This was because the agents learned to avoid each other and this population

<sup>1</sup>Note that  $P_i^D$  is again irrelevant: if all agents steal, none has to choose whether to defend its resources or not.




Figure 10.3: Total agent population over 2,000 rounds of a typical simulation when the agents' propensities to steal are fixed at 1, i.e., the agents only ever steal. The population collapses in all 20 simulations owing to debilitating fight costs.

was well below the (non-specialised) carrying capacity of the environment, which meant resources per capita were abundant. Note, however, that if we extend the simulations long enough, the population increases to approximately 6-7 agents<sup>2</sup>

The two main contributing factors to the agents' demise were: (i) the cost of fighting, which was a persistent drain on agents' resources; and (ii) the fact they would not specialise in this environment (since there was no hope of trading), which meant levels of productivity remained low.

There is one phenomenon worth noting here, which is the occasional clustering of agents during the interaction phase of some of the rounds. This begs the question of why agents would target squares at or close to those of other agents when the agents would only steal from each other.

In these simulations, an agent would have a location in memory with a positive weight if it had stolen resources from another agents at that location. If the other agent had re-acquired these resources in an adjacent squares then this location would be in the counterpart's memory, also with a positive weight.

In the next round the two agents would ignore the squares in which they had lost resources but would consider going to the squares where they had acquired resources. In this next round, the agents' target squares might end up being adjacent to each other, in which case they would be able to see each other and possibly attempt theft once again,

<sup>2</sup>Recall that in the first model the non-specialised carrying capacity of the environment was 15-16 agents: the carrying capacity when the agents only ever fought was lower than this owing to the impact of fight costs when the agents did interact.

potentially maintaining adjacent locations with positive weights in their memories. This phenomenon led to some clustering of the agents.

### 10.1.3 Null Scenario: Substantive Rationality Approach

In this section we consider what happened in the simulations when the agents adopt full substantive rationality.

Recall from Fig. 9.1 (page 259) that the agents have potentially two decisions to make in any interaction: (i) whether to trade or steal; and (ii) whether to defend their resources or acquiesce.

In order to make these decisions under substantive rationality, agents first use backward induction to determine whether they would acquiesce or defend their resources in scenarios 2 and 3. They would choose to defend if the pay-off exceeds that of acquiescing, and vice versa<sup>3</sup>. Knowing their strategies in scenarios 2 and 3 means each interaction is reduced to a  $2 \times 2$  ‘game’. Substantive rationality is then applied to determine whether the agents attempt to trade or steal. If the agents’ choices are indeterminable (at least one does not have a dominant strategy) then they do not interact.

Fig. 10.4 below shows a time series of the number of interactions in the six different scenarios shown in Fig. 9.1 above over the first 80 rounds of a typical simulation. We can see that in the vast majority of interactions, one or both of the agents attempted to steal (all scenarios except 1). Transactions make up only approximately 1% of all interactions prior to the first agent dying in the simulation shown (in Round 22).

Looking at the data in Fig. 10.4, it seems surprising we observed interactions in Scenario 2A, which meant the instigating agent deliberately allows itself to be mugged (approximately 7% of all interactions). This occurred when both the instigating agent and its counterpart held no resources but pay-off errors meant the instigator expected to benefit, i.e., it made a mistake.

It appears even more perplexing that an instigating agent would find itself in Scenario 2F (approximately 0.2% of interactions). This occurred when the instigating agent wanted to trade (and expected the same of its counterpart) but the other agent’s ‘reading’ of the pay-offs (which were different to the instigator because it made different errors) meant it wished to steal. This was also a mistake by the instigating agent.

The fight costs incurred by the agents (mostly in scenarios 3F, and 4) were sufficient to ensure the agent population declined rapidly. This can be seen in Fig. 10.5 below which illustrates the total agent population in the simulation shown in Fig. 10.4. The

---

<sup>3</sup>In the (very rare) event that these two pay-offs were identical, the agent chose one of the strategies randomly.

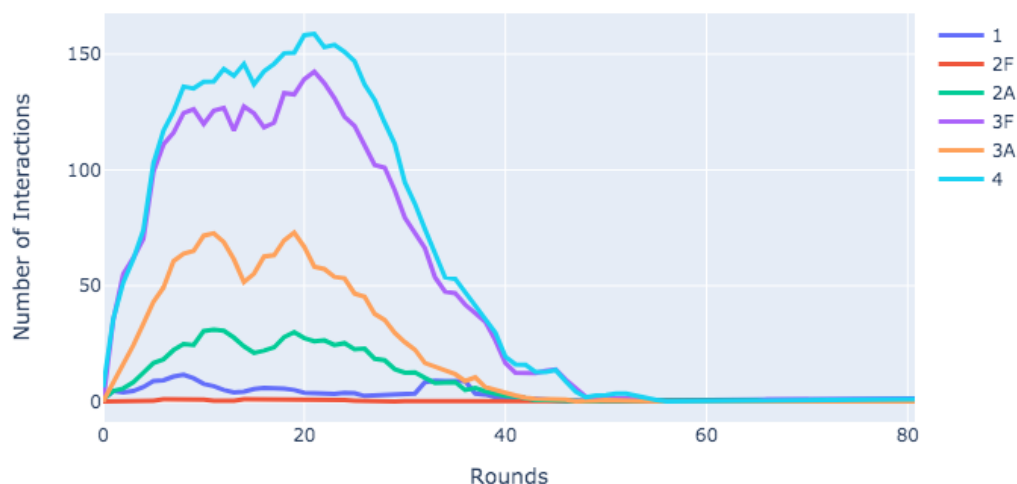


Figure 10.4: Agent interactions by scenario (5-round moving averages) when decisions are made via substantive rationality. Note that ‘1’ refers to transactions; ‘2F’ refers to interactions when the instigating agent attempts to trade but defends its resources when the counterpart tries to steal; ‘2A’ is the same as 2F but the instigating agent acquiesces; ‘3F’ is when the instigating agent attempts to steal and its counterpart fights back; ‘3A’ is the same as 3F but the counterpart acquiesces; and ‘4’ is when the agents both attempt to steal from each other. Substantive rationality leads agents to attempt theft from each other most of the time.

population collapsed to only two agents by Round 49 - this was well below the carrying capacity of the environment and the population began to increase from approximately Round 1,300 (given enough time it would rise to approximately 6-7 agents for the same reasons identified in the previous sub-section).

One additional point worth noting here is that approximately 340 unique interactions were observed across 20 simulations<sup>4</sup>. This variety of game types, in addition to the ability of the agents to defend their resources, is one of the reasons this model differs from the Hawk-Dove game, which typically uses a single pay-off structure.

Moreover, the model compares every considered and enacted interaction with all of the classic  $2 \times 2$  games in the game theory literature (the Prisoners’ Dilemma, Matching Pennies, Stag-Hunt, Hawk-Dove, etc.) and keeps a tally of each. Of the 2.82 million games considered in 20 simulations, 13.0% had the same pay-off structure as the Hawk-Dove game; and of the 0.34 million games enacted, 4.3% had this structure.

<sup>4</sup>A unique interaction type is defined here by the ranking of the pay-offs in the  $2 \times 2$  reduced form game, i.e., post-backward induction.




Figure 10.5: Total agent population when decisions were made via substantive rationality. The data are taken from the same simulation depicted in Fig. 10.4. The fact that most agents stole meant the population collapsed.

Overall we can see that the results were approximately similar to the second null model, when agents only ever stole. The agent population collapsed and the surviving agents never specialised.

### 10.1.3.1 Resource Concentration Effect

It is worth looking in some detail at a phenomenon we saw in these simulations because it is also observed in the default simulations described below. A ‘resource concentration effect’ is seen when the agents become clustered during the interaction phase: they start the round with relatively dispersed resources (immediately after foraging and before interacting) but as this phase of the round progresses, the resources become concentrated in fewer hands.

Fig. 10.6 below illustrates this phenomenon during the interaction phase of Round 30 in a typical simulation.

An analysis of the data revealed that this concentration of resources was due to the combination of two (related) factors: first, the proximity of agents to each other; and, second, a bias of the agents toward interacting with others who had more resources than them (and avoiding agents with fewer resources).

The **proximity of agents** to each other despite the high proportion of fights and muggings was explained in Section 10.1.2 above.




Figure 10.6: Agents' resource holdings at the start, middle, and end of the interaction phase of Round 30 of a typical simulation when agents made decisions via substantive rationality. In these charts, the agents' resource holdings were put in ranked order at the beginning of the interaction phase (the first time period), in the middle (the 25th time period) and at the end (the 50th time period). The charts show how the agents' resource holdings became more concentrated over this phase of the round.

In terms of the **bias to interact with agents holding more resource**, recall from Section 9.5 that an agent chooses to interact with other agents according to which interaction is most beneficial to the agent. It is general preferable, *ceteris paribus*, to interact with an agent holding more resources (whether stealing or trading).

The data shows that in the first 40 rounds or so, agents on the whole chose to steal from those who had more resources than them. There were a few transactions and these typically took place early on during the interaction phases (when holdings were more balanced). However, we observed a positive feedback effect: as resources became more concentrated in fewer hands, agents with a few resources (or none) would attempt to steal from those with more. This resulted in resources becoming more concentrated.

For example, in the middle sub-plot of Fig. 10.6 we can see that one agent held 12 units of resources (two agents held 4 units). In this situation, all the agents within sight of the most 'wealthy' agent would attempt to steal its resources (including those who held 4 units) because the gain was easily worth the cost of fighting.

This concentration of resources explains the time series shown in Fig. 10.4 above. Approximately 92% of interactions were in scenarios 3F, 3A, or 4: here, the instigating agents attempted to steal from other agents<sup>5</sup>.

<sup>5</sup>In general, if the instigating agent has no resources and the counterpart has 1 unit, the counterpart would acquiesce (3A) because the expected (reduced) loss of defending its resources will be approximately 1.1 units: there is a 50% chance of losing 0.6 units via fight costs if it wins the fight, and a 50% chance of losing 1.6 units if it loses. This means it is preferable not to fight in this situation. If the counterpart has more than 1 resource unit then it is preferable for it to

### 10.1.4 Null Scenario: Fixing Agents' Propensities at 0.5

In these simulations, we use the agents' propensities to determine whether they attempt to trade or steal, and whether they would acquiesce or defend their resources. All the agents' propensities to steal and defend are fixed at 0.5.

An important point to note here is that while the agents' decisions were determined by their propensities, they did have discretionary power over *who* they interacted with. In general we (again) found that instigating agents typically chose to interact with counterparts who had more resources than them: this was beneficial to all agents whichever scenario they ended up in.

This bias led to the same concentration effect noted above: resources began the interaction phase of each round relatively dispersed but subsequently they became more concentrated among the agents.

Fig. 10.6 above, which shows this concentration effect for when agents made choices via substantive rationality, also included a gini coefficient for the agents' resource holdings at the beginning (0.287), middle (0.604), and end (0.870) of the interaction phase of a typical round. Fig. 10.7 below shows a time series of this gini coefficient (for resource holdings at the start and end of the interaction phases of each round) for the first 200 rounds of a typical simulation when the agents' propensities were fixed at 0.5. It shows how the coefficient started the interaction phase of each round at approximately 0.3 in the first 100 rounds but ended much higher (as high as 0.8 between rounds 10 and 70).

This phenomenon meant that, in effect, the agents with no or few resources persistently interacted with those who held more resources, in the hope of eventually acquiring all their resources. By contrast, agents with resources generally attempted to evade other agents.

Ultimately, the results of these simulations were similar to the second null scenario when we fixed  $P^S = 1$ : the agent population eventually collapsed and any surviving agents never specialised. In these simulations, the agents typically lived a little longer on average due to lower fight costs.

As a final note, variations of this experiment were run to reveal the propensity to steal at and below which agents would specialise and bear children. We saw in the first null experiment that agents specialise and bear children if their propensities to steal were fixed at zero and we saw that the population collapsed when these propensities were fixed at 0.5. This begs the question of the threshold at which the agents were comfortable specialising.

---

defend its resources. Interactions in Scenario 3A made up approximately 17% of all interactions and typically occurred earlier in the interaction phases.




Figure 10.7: Time series of the gini coefficient of agents' resources holdings over the first 200 rounds of a typical simulation when all the agents' propensities were fixed at 0.5. In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. The chart shows that resources became more concentrated in approximately the first 100 rounds (when the dark blue line was higher than the light blue line) but not thereafter (when most agents had died and those remaining generally avoided each other).

If we fix all the agents' propensities to defend at 1 we find that the propensity to steal threshold at and below which agents specialise is approximately 0.05. In these simulations we found that some of the agents died in the first 500 rounds or so but those who survived specialised and had children. However, the drain on the agents' resources due to any fighting approximately balanced the productivity gains from specialisation such that a total population of approximately 25 agents was maintained.

When we fix the agents' propensities to defend at 0, the propensity to steal threshold below which agents specialise and bear children is approximately 0.15. In these simulations the environment is less harsh for the agents in the sense the total cost of fighting is lower (agents always acquiesce in scenarios 2 and 3). In effect, the agents 'accept' a higher propensity to steal before they specialise.

These simulations tell us that the agents' propensities to steal had to decline to a relatively low level (irrespective of their propensities to defend) in order for the turnover ratio to be sufficiently high to encourage specialisation.

Let us now look at the results of simulations that use the default parameter set, which is when the agents' propensities are allowed to evolve via reinforcement learning, given their experiences in interactions.



## 10.2 Property Rights Model: Default Parameter Set

The results for simulations that use the new default parameter set are described and analysed in this section. In these simulations the agents' initial propensities to steal and defend are both drawn from a normal distribution with a mean of 0.5 and a standard deviation of 0.1; and both are allowed to vary in the way described in Section 9.4<sup>6</sup>.

The results of these simulations can be summarised by the following four points:

- the agents' propensities to defend increased on the whole as they learned it was generally preferable not to acquiesce when other agents attempted to steal from them (we refer to this as *defence of property*);
- when the agents' propensities to defend were below approximately 0.8 on average, the agents' propensities to steal *increased*;
- when their propensities to defend were above approximately 0.8, the agents' propensities to steal *decreased* toward and below 0; and
- when all the surviving agents' propensities to steal were negative (which looked like the institution of *property rights*), the agents specialised and, ultimately, the results of the first model were replicated.

Let us now examine these results in more detail.

### 10.2.1 Propensities to Defend

Fig. 10.8 below illustrates a fan chart of the (living) agents' propensities to defend over the first 250 rounds of a typical simulation. The chart illustrates how the agents' propensities increased immediately after the simulation started (the mean exceeded 1 after Round 28).

The chart shows that, on the whole, *defence of property* emerged across the population. An analysis of the data shows that in most interactions the agents learned it was preferable to defend their resources than to acquiesce<sup>7</sup>.

Fig. 10.9 below shows a time series of the cumulative contributions to the agents' propensities to defend by scenario. The y-axis represents the cumulative net contribution to

---

<sup>6</sup>We examine the effect of different initial propensities when we explore the parameter space.

<sup>7</sup>The only exception to this general rule was when the instigating agent had no resources and its counterpart had a total of approximately 1 resource unit: here, the agents' propensities to defend were driven *lower* on average.




Figure 10.8: Fan chart of the agents' propensities to defend over the first 250 rounds of a typical simulation when the default parameters were used. The black line represents the mean and each red band represents one standard deviation (skew-adjusted) away from the mean. Note the data shows how the agents' propensities were allowed to increase above 1 and to decline below 0 (the probability an agent defended its resources was restricted to between 0 and 1, which are limits shown by two green dashed lines). In this simulation the mean propensity to defend increased above 1 in Round 28.

all the agents' propensities to defend in each of the scenarios shown (1 and 4 are omitted because the agents did not learn anything about defending their resources in those scenarios).




Figure 10.9: Time series of cumulative contributions to the agents' propensities to defend by scenario over 250 rounds in the simulation depicted in Fig. 10.8. The model records changes in the agents' propensities to defend by interaction (scenarios 2F, 2A, 3F, and 3A). These were aggregated in each round by scenario and the cumulative time series for one of these simulations is shown here. The chart shows that agents mostly learned to defend their resources in scenarios 2A and 3A.

There are two main observations: first, in scenarios 2A and 3A agents learned it was preferable to defend their resources; and, second, the contributions of scenarios 2F and 3F varied but their net contributions were positive (albeit less than 2A and 3A).

## 10.2.2 Propensities to Steal

Fig. 10.10 below illustrates a fan chart of the (living) agents' propensities to steal over the first 250 rounds of the simulation shown in figures 10.8 and 10.9.




Figure 10.10: Fan chart of the agents' propensities to steal over the same simulation as that depicted in figures 10.8 and 10.9. The black line represents the mean and each blue band represents one standard deviation (skew-adjusted) away from the mean. Note the data uses the agents' propensities, which were allowed to increase above 1 and to decline below 0 (the probability an agent attempted to steal was restricted to between 0 and 1, which are shown by two green dashed lines). The chart shows how the mean propensity to steal initially increased but then declined to below 0.

We can see from this chart there was a slight tick up in the mean of the agents' propensities in the first few rounds: this mean peaked at 0.65 in Round 9 when the mean propensity to defend was 0.74. This was typical of all 20 simulations.

The relationship between the agents' propensities to defend and their propensities to steal will be examined in more detail when we explore the parameter space in Chapter 11 and Appendix C. For now, it is worth noting that when the agents' propensities to defend were below approximately 0.8 on average, their propensities to steal tended to increase, and vice versa<sup>8</sup>.

To help us understand this better, Fig. 10.11 below shows the net impact of interactions in each scenario on the agents' propensities to steal over the first 250 rounds of the

<sup>8</sup>*In extremis*, if we hold the agents' propensities to defend at zero, the agents' propensities to steal *increase* to above 1; and if we hold them at 1, the agents' propensities to steal *decline* to below 0.

simulation depicted in figures 10.8 to 10.10 above. It is similar to Fig. 10.9 but this chart shows gross, not cumulative, contributions and it includes an aggregated value in black.




Figure 10.11: Time series of contributions to the agents' propensities to steal by scenario over 250 rounds in the simulation depicted in figures 10.8 - 10.10. The model records changes in the agents' propensities to steal by interaction. These were aggregated in each round by scenario and the time series for one of these simulations is shown here. The chart shows that agents mostly learned to steal in scenarios 3A and 3F but they learned to trade in scenarios 2F and 4. Ultimately, the agents learned to respect each other's property.

The small upward movement in the mean propensity to steal at the beginning of the simulations appears to be a result of agents learning that it was preferable to steal after interactions when they attempted to trade but then acquiesced to a stealing counterpart (Fig. 10.11 shows this phenomenon as a sharp rise in the contribution of Scenario 3A in the first few rounds). However, this did not last for very long: as the agents learned it was preferable to defend their resources (fewer counterparts acquiesced) they also learned it was preferable to trade (not to steal) on the whole.

A more detailed analysis of the data indicates there were a total of six different 'patterns' acting on the agents' propensities to steal. These were discussed in Chapter 6 and are summarised here:

- The first is described above: a movement up in the agents' propensities to steal when their propensities to defend were below approximately 0.8.
- The second is the reverse of the first: after the agents' propensities to defend exceed 0.8 on average, the agents learn it is preferable to trade than to steal, i.e., their propensities to steal decline. This is mainly due to fight costs, which on the whole exceed the benefits of theft.

- The third pattern is observed when the rate of decline of agents' propensities to steal decelerates because agents fight less.
- The fourth pattern is observed more as propensities to steal decline: the benefits of transacting further reduces these propensities.
- The fifth pattern was unexpected. When one agent has a positive propensity to steal and all the other agents' propensities are negative, this single agent benefits disproportionately from the resource concentration effect, which is 'centred' around this agent. This encourages it to steal so its propensity to steal increases. Ultimately, other agents learn from this agent, which leads to an increase in total fight costs.
- The final pattern is when all the agents' propensities to steal are negative, leading to lock-in below zero. Here, agents only trade (which they can only benefit from), so their propensities become more negative.

To help illustrate the net effect of these six patterns, Fig. 10.12 below shows the agents' propensities to steal over the first 200 rounds of a typical simulation. The chart shows four sets of data: the propensities to steal of the agents who survive until the end of the simulation (the blue 'cloud'), the mean of these propensities (the blue line), the propensities to steal of the eleven agents who die (the red 'cloud'), and the mean of these (the red line). It is replicated from Chapter 6 for convenience.

The chart shows how (on average) the propensities to steal of the agents who survived were sustained below those of agents who died (the last one died in Round 102). Note that the mean propensities to steal of the agents who died (the red line) declined roughly in parallel with the mean propensities of those who survived (the blue line).

Let us refer to the agents with relatively higher propensities to steal as 'Al Capone' agents, and those with relatively lower propensities as 'passive-aggressive' agents. The latter name reflects the idea that, ultimately, these agents prefer to trade but will defend their resources if necessary.

When we introduce habituation in to the model (see Section 10.3 below), the heterogeneity depicted in Fig. 10.12 is strengthened such that the agents bifurcate into two distinct strategies: those with propensities to steal above 1 and those with propensities below zero.

The net impact of these six patterns and the fact that agents with positive propensities to steal tended to die before those with lower and/or negative propensities mean that property rights eventually emerge.

There are two final points to note here. The first is that the resource concentration effect was seen in these default simulations. Fig. 10.13 below shows the time series of




Figure 10.12: The ‘cloud’ of agents’ propensities to steal over 200 rounds in a typical simulation when the default parameter set was used (copied from Chapter 6 for convenience). The chart shows the propensities to steal of the agents who survived until the end of the simulation as blue ‘dots’. The blue line is the mean of these values. The red dots represent the propensities to steal of the agents who died before the end of the simulation, and the red line is the mean of these values (it ends in Round 102 when the last of these agents died). The chart shows how, on the whole, the propensities to steal of the agents who died was sustained, initially, above those who survived. This figure is replicated from Chapter 6 for convenience.

the mean gini coefficient for resource holdings at the beginning and of the interaction phases of rounds, averaged over 20 simulations that used the default parameter set. The data are consistent with resource concentration until approximately Round 200.

One of the most interesting features of this effect was the observation of ‘passive-aggressive theft’ by the passive-aggressive agents. We observed agents with low or negative propensities to steal choosing to interact with agents with much higher propensities to steal. An analysis of the data tells us that these agents expected to attempt to trade, their counterparty to steal, and then to defend their resources. If they win this fight, they get to keep the other agent’s resources.

This tactic was employed by passive-aggressive agents when they had fewer resources than their Al Capone counterpart. The benefit to the instigating agent exceeded the expected fight costs. This phenomenon contributed to resource concentration.

The second point is that it is tempting to expect a population in which property rights have fully emerged to be susceptible to ‘hawks’ or ‘invading defectors’, i.e., agents who steal. However, the trading agents in these simulations should perhaps be thought of




Figure 10.13: Time series of the gini coefficient of agents' resources holdings (means over 20 simulations). In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. Mean values were also recorded over all 20 simulations in a set, which are shown here. The chart shows that resources became more concentrated in approximately the first 200 rounds (when the dark blue line was higher than the light blue line) but the gap between the two declined until approximately Round 250. The decline in resource concentration this demonstrated was a result of the agents' propensities to steal declining.

as 'contingent doves' (rather than 'pure' doves in the Hawk-Dove game): they will first seek to trade with other agents but they will defend their resources if provoked. We will see this phenomenon at work in the next chapter, when we allow children to be born with positive propensities to steal (they always die because other agents defend their property).

### 10.2.3 Fights and Transactions on the Grid

Let us look briefly at the transaction and fight locations in a typical simulation. Three figures are presented below, which show (for the same simulation) heatmaps of:

1. fight locations during the first 100 rounds (Fig. 10.14);
2. transaction locations also during the first 100 rounds (Fig. 10.15); and
3. transaction locations during the last 100 rounds (Fig. 10.16).

No heatmap for fights in the last 100 rounds is shown because there were no fights.

Figures 10.14 and 10.15 show that agents converged on the same area on the grid to trade and steal in the first 100 rounds.

At first blush it seems odd that agents converged on an area where they might lose their resources and/or incur fight costs. However, as mentioned previously, this was




Figure 10.14: Heatmap of fight locations in the first 100 rounds of a typical simulation when the default parameter set is used. The chart shows how agents congregated on the same area and fought.




Figure 10.15: Heatmap of transaction locations in the first 100 rounds of a typical simulation when the default parameter set is used. The chart shows how agents congregated in the same location shown in Fig. 10.14 above, i.e., where fights took place.

due to the process of target selection at the start of the interaction phases: agents with positive location weights in memory within this area were insensitive to negative weights for adjacent squares, which meant agents tended to congregate. Note we increase this sensitivity to neighbouring squares with negative weights when we explore the parameter space.

Fig. 10.16 shows a single market area at the end of the simulation.






Figure 10.16: Heatmap of transaction locations in the last 100 rounds of the same simulation depicted in figures 10.14 and 10.15. The chart shows how agents congregated in the same location shown in those charts. However, agents only ever traded at this stage of the simulation.

### 10.2.4 Specialisation & Child Birth

After the Al Capone agents died and the propensities to steal of all the surviving agent declined to below zero, the simulations essentially replicated the outcome of default simulations that used the first model: agents specialised and bore children until the total population reached approximately 43 agents.

Fig. 10.17 below shows the turnover ratio from the same simulation from which figures 10.14 - 10.16 are taken<sup>9</sup>.

The turnover ratio shown in Fig. 10.17 increases progressively to almost 1 over 200 rounds as the agents' propensities to steal declined toward and below zero and agents with positive propensities died off.

This rise in the turnover ratio catalysed specialisation among the surviving agents.

In the default simulations the children were born with propensities to steal and defend equal to the mean of their parents' propensities. This meant children were only ever born with negative propensities to steal.

Fig. 10.18 below shows the total agent population in the same simulation as that depicted in figures 10.14 - 10.17 above, over the first 1,000 rounds. The first child was born in Round 403. The total population continued to increase after Round 1,000 and plateaued at 43 agents.

<sup>9</sup>Recall that the turnover ratio is the actual volume of transactions divided by the market clearing volume.




Figure 10.17: Time series of the turnover ratio in the first 250 rounds of the same simulation depicted in figures 10.14 - 10.16. The ratio increased progressively toward 1 as the Al Capone agents died off.




Figure 10.18: Time series of the total agent population from the same simulation depicted in figures 10.14 - 10.17 above. The population initially declined to 15 agents but property rights emerged among these agents who met at the same market, specialised and then bore children.

Now that we have discussed the results of the new default simulations, let us consider whether habituation has a material impact on the simulation results.

### 10.3 Habituation Experiments

In this section we discuss the results of experiments in which habituation plays a role in the model.

Consistent with the equivalent experiments based on the first model, this is done by applying a habituation parameter (here,  $ha_2$ ) to the agents' propensities to steal and defend: if an agent chooses to steal (trade),  $ha_2$  is added to (deducted from) its propensity to steal; and if an agent chooses to defend its resources (acquiesce),  $ha_2$  is added to (deducted from) its propensity to defend.

Two sets of experiments are discussed below: those in which reinforcement learning is 'switched off' such that propensities change only as a result of habituation; and those in which reinforcement learning and habituation both change the agents' propensities.

As with the experiments based on the first model, the interesting question here is what impact habituation has on the emergence of the institution. We can imagine that after property rights emerge across the population, habituation serves to reinforce the institution (agents' propensities to steal would be more negative) but what about when property rights are emerging?<sup>10</sup>

Note that for these experiments the agents' starting resource values were increased to 1,000 units of each resource and the agents were prevented from siring children. In addition, 20 simulations were run for each value of  $ha_2$ , each for 1,000 rounds. These parameters allowed us to focus on the underlying mechanisms at play and to minimize run time.

### Main Results

The results of both sets of experiments provide us with two key observations.

The first, drawn from both experiments, is that *reinforcement learning is necessary* to guarantee that property rights emerge. More specifically, this is to avoid a scenario in which one or two hawkish agents (with positive propensities to steal) are allowed to thrive by exploiting a large group of 'doves'.

The second observation, which is taken from the second set of experiments, is that habituation *catalyses the emergence of property rights* by strengthening the strategy bifurcation mechanism. In these experiments there is a much clearer emergence of Al Capone and passive-aggressive agents (with more extreme propensities): the former die off more quickly as a result of debilitating fight costs.

---

<sup>10</sup>This is a particularly relevant question in light of the general movement of the agents' propensities to steal in the default simulations. These propensities tend to move up early on in the simulations, when the agents' propensities to defend are relatively low; and then decline when these propensities are high. Furthermore, when the data was disaggregated and analysed, we observed a relatively weak bifurcation of 'strategies' (denoted 'Al Capone' and 'passive-aggressive'): does habituation effect this split? These questions are addressed in Section 10.3.2 below.

### 10.3.1 Reinforcement Learning ‘Switched Off’

In these experiments, agents do not learn anything from their interactions. Their propensities to steal and defend change only as a result of  $ha_2$  being added to or deducted these propensities, as described above.

We find there are two simple mechanisms at play here. First, there is a positive feedback effect in which propensities above 0.5 tend to move higher, which increases the probability that it moves higher in the future; and vice versa. We can think of this as inherent bifurcation.

This process leads to a four-way split in the agents’ propensities (creating three ‘types’ or ‘strategies’), determined by the movement in the agents’ two propensities: (i) both the propensities to steal and defend are negative (denoted ‘doves’); (ii) a negative propensity to steal and a propensity to defend above 1 (the ‘passive-aggressive’ agents we have seen before); and (iii) agents with propensities to steal above 1 (‘hawks’ who can have propensities to defend above 1 or below zero - these are irrelevant, ultimately, because these agents never have to choose between acquiescing or defending their resources).

The speed at which these strategies emerge depends on  $ha_2$  and the number of interactions an agent is involved in. If  $ha_2 > 1$ , the agents’ strategies are determined by the first time they choose to steal / trade and defend / acquiesce<sup>11</sup>. For smaller value of  $ha_2$ , the agents’ propensities change more slowly. Fig. 10.19 below helps us visualise this (it shows the agents’ propensities to steal over the first 500 rounds of a typical simulation when  $ha_2 = 0.0001$ ).

The second mechanism at play is lock-in when an agent’s propensity to steal exceeds 1 or falls below 0, which is more relevant when  $ha_2 < 1$ . When this happens, the agent’s decision (to steal / trade or defend / acquiesce) is certain and its propensity will carry on increasing above 1 or declining below 0. We see propensities exceed 1 or decline below zero in the default simulations but, there, agents can learn in a way that brings their propensities back over either threshold. For the simulations discussed here, once an agent’s propensity exceeds 1 or declines below zero, this is locked in.

Table 10.1 below shows some key results for selected values of  $ha_2$ . The table shows the total population of agents at the end of Round 1,000 and the percentage of 20 simulations (for each value of  $ha_2$ ) in which property rights emerged by the same stage.

When we analyse the data, we make three broad observations.

---

<sup>11</sup>The agents’ starting propensities are drawn from a normal distribution with mean 0.5 and standard deviation of 0.1, which means all the agents’ propensities will almost always start between 0 and 1. If  $ha_2 = 1$ , therefore, the agent’s propensities will move above 1 or below zero immediately after their first decision.




Figure 10.19: A time series of the agents' propensities to steal over the first 400 rounds of a typical simulation when  $ha_2 = 0.0001$  and when there was no reinforcement learning. Here, the agents' propensities changed too slowly for property rights to emerge.

$ha_2$	Population	% Property Rights
0	2.9 (0.7)	0
0.00001	3.0 (0.7)	0
0.0001	6.4 (3.3)	50
0.0002	7.3 (4.0)	40
0.0005	11.6 (3.0)	85
0.001	13.4 (2.2)	95
0.005	12.6 (3.0)	80
0.01	12.1 (2.8)	85
2.0	12.0 (2.2)	80

Table 10.1: Main results for simulations when there was no reinforcement learning, for selected values of  $ha_2$ . The middle column shows the total population of agents at the end of 1,000 rounds, and the right hand column shows the percentage of simulations in which property rights emerged (defined as when the propensities to steal of all the surviving agents at the end of the simulation were negative).

The first is related to the 'null' scenario, when  $ha_2 = 0$ , and for very small values of  $ha_2$ . When  $ha_2 = 0$ , there is neither reinforcement learning nor habituation so the agents' propensities never change. The agent population always collapses due to debilitating fight costs and, clearly, property rights can never emerge. We see the same for very low values of  $ha_2$ .

The second observation is that as  $ha_2$  is increased from very small values, property rights are more likely to emerge and the mean total population of agents increases as a result. Let us consider this in more detail by examining the results for when  $ha_2 = 0.005$ , which are typical of this category.

In 17 of the 20 simulations in this set, the agents' strategies split relatively evenly (about half the agents become hawks, about 1/4 doves, and about 1/4 passive-aggressive). Fig. 10.20 below shows the agents' propensities to steal in a typical simulation (first 400 rounds only). Here, 13 hawks (shown in red) emerge alongside 6 doves and 6 passive-aggressive agents (both are shown in blue). All of the hawks die by Round 300 because of debilitating fight costs, and the rest of the agents survive.




Figure 10.20: A time series of the agents' propensities to steal over the first 400 rounds of a typical simulation when the agents' propensities changed as a result of reinforcement learning and habituation (here,  $ha_2 = 0.005$ ). The red lines depict 'hawks', whose propensities increased above 1, and who all died. The blue lines depict 'doves' and passive-aggressive agents (who all survived).

Fig. 10.21 shows a time series of the total stock of resources held by agents over 1,000 rounds from the same simulation. The three agent groups can be seen clearly: all of the hawks' total resources reached zero; the lower group of 6 lines above zero represent the passive-aggressive agents; and the higher group of 6 lines are the doves.

In these simulations, an efficient market emerged after the hawks died off, and the agents specialised.

It is tempting to conclude from this that we have essentially replicated the outcome of default simulations without any reinforcement learning. After all, property rights emerge across the population. This would be a mistake, however, given the third observation.

The third and final observation is that in simulations when (approximately)  $ha_2 \geq 0.0005$  we observe scenarios in which property rights do not emerge. This is due to the 'split' of agents between the three strategies: the problem is observed when there is a relatively large number of doves relative to passive-aggressive agents.

To help us understand this better, Fig. 10.22 below shows the agents' propensities to steal in a simulation when 14 hawks emerged alongside 8 doves and 3 passive-aggressive




Figure 10.21: A time series of the agents' total resources over 1,000 rounds (taken from the same simulation as that depicted in Fig. 10.20 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows that all the 'hawks' died; all the passive-aggressive agents (the lower group of lines above zero) survived; and the 'doves' (the higher group of lines above zero) also survived.

agents ( $ha_2 = 0.01$ ). Fig. 10.23 shows the the agents' total resources in the same simulation.




Figure 10.22: A time series of the agents' propensities to steal of an atypical simulation when the agents' propensities changed as a result of reinforcement learning and habituation (here,  $ha_2 = 0.005$ ). The red lines depict 'hawks', whose propensities increased above 1, all of whom died with one exception. The blue lines depict doves passive-aggressive agents. This pattern is observed when a relatively large number of 'doves' are seen relative to passive-aggressive agents. The surviving hawk thrives in this environment.

The significant point here is that a lone hawk (sometimes two) can thrive in this environment because there are several doves to bully, i.e., for these hawks, the theft of




Figure 10.23: A time series of the agents' total resources over 1,000 rounds (taken from the same simulation as that depicted in Fig. 10.22 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows that all the 'hawks' but one died - this surviving hawk was able to bully the doves (after all the other hawks died) allowing its resource stock to increase significantly.

resources from doves exceeds the loss in resources from fight costs. This is not true when more than two hawks are alive. The line in Fig. 10.23 that rises steeply after Round 350 represents a single surviving hawk. The lines declining gently over the 1,000 rounds represent the doves.

There was some variation in these scenarios. As mentioned above, sometimes 2 hawks survive and thrive. In others, a single hawk thrives by being parasitic on the doves but a single passive-aggressive agent thrives by being parasitic on the surviving hawk! Here, the hawk mugs the doves but the passive aggressive agent shares in this success to some degree by successfully stealing (in a passive aggressive way) from the hawk. The spoils of theft are split between the two on average.

These scenarios in which doves are bullied by one of two hawks are unsustainable: if we extend the simulations beyond 1,000 rounds we see that, eventually, the doves die off.

The critical point to emphasise here is that in the default simulations, reinforcement learning ensured that no doves emerged: agents learned it is best to defend their resources. This scenario, in which 1-2 hawks thrived among doves, was avoided.

Let us now turn to the experiments in which reinforcement learning coexists with habituation.



### 10.3.2 Reinforcement Learning and Habituation

It is helpful to first demonstrate what happens under these conditions when  $ha_2 = 0$ , as a type of null scenario<sup>12</sup>. 20 simulations were run and we found that in 7 of these, property rights emerged successfully by Round 1,000. Fig. 10.24 below shows the ‘cloud’ of the agents’ propensities to steal in a typical simulation. It shows that all the agents’ propensities to steal declined to below zero soon after Round 550.




Figure 10.24: The ‘cloud’ of the agents’ propensities to steal over 1,000 rounds of a simulation when the agents’ propensities change as a result of reinforcement learning but not habituation, i.e.,  $ha_2 = 0$ . Each blue ‘dot’ represents the propensity to steal of a living agent in each round. In this simulation, the agents’ propensities fell below zero after Round 550 (approximately): property rights emerged.

In the other 13 simulations, property rights did not emerge across the population by Round 1,000. Fig. 10.25 below shows the cloud of agents’ propensities to steal in such a simulation: by the end of the simulation, about half the agents had propensities above zero and half below.

The difference between the simulations depicted in figures 10.24 and 10.25 is due to the outcome of the six patterns discussed above. On the one hand, fight costs tend to reduce the agents’ propensities to steal but, on the other, there were forces at play which encouraged propensities to steal higher at lower (and negative) levels. These combined in different ways in these simulations, leading to distinctly different outcomes as shown in figures 10.24 and 10.25 above<sup>13</sup>.

<sup>12</sup>This replicates the default parameter set except, now, agents start with 1,000 units of each resource and can never sire children.

<sup>13</sup>Note that in the default simulations agents had much lower starting resources - this meant agents who sustained relatively high propensities to steal were prone to dying more quickly. In the simulations shown in figures 10.24 and 10.25, these agents survived for longer.




Figure 10.25: The ‘cloud’ of the agents’ propensities to steal over 1,000 rounds under the same conditions as those in Fig. 10.24 above. In this simulation, the agents’ propensities did not fall below zero before the end of the simulation: property rights did not emerge.

We should note that if the simulations are extended and agents allowed to have children, all of the simulations typified by Fig. 10.25 would see property rights and an efficient market emerge, followed by specialisation and childbirth, i.e., the default simulations would be replicated. This is not shown here because we are focused on the impact of habituation on the emergence of property rights.

Table 10.2 below shows some key data from the results for selected values of  $ha_2$ . For each value of  $ha_2$ , the table shows: (i) the mean total population of agents at the end of Round 1,000; (ii) the percentage of 20 simulations in which property rights emerged across all surviving agents by the end of Round 1,000 (this is when all surviving agents had propensities to steal below zero); (iii) the peak propensity to steal (mean across 20 simulations); (iv) the mean number of ‘doves’ observed across the simulations (by Round 1,000); (v) the mean number of ‘passive-aggressive’ agents observed by Round 1,000; and (vi) the mean number of ‘hawks’ observed by Round 1,000. The table’s caption contains a more detailed description of the data.

When we analyse the data we observe two noteworthy phenomena, corresponding to higher and lower values of  $ha_2$ .

The first phenomenon, which is related to higher values of  $ha_2$  (approximately  $0.01 \leq ha_2$ ), is easier to understand because it mimics the results we saw when reinforcement learning was switched off. Higher values of  $ha_2$  dominate changes in the agents’ propensities: reinforcement learning was relatively weaker. The results, therefore, converge on those discussed in the previous sub-section. Understandably, therefore, higher values of

$ha_2$	Pop'n	% Prop Rights	Max Prop Steal	No. Doves	No. PA	No. Hawks
0	24.1 (1.5)	35	0.563	0.0 (0.0)	19.1 (3.8)	0.1 (0.3)
0.0001	23.9 (0.8)	100	0.543	0.0 (0.0)	20.4 (2.0)	0.0 (0.0)
0.001	24.4 (0.6)	100	0.555	0.0 (0.0)	21.5 (1.7)	0.0 (0.0)
0.0025	22.9 (1.0)	100	0.550	0.0 (0.0)	22.8 (0.7)	0.0 (0.0)
0.005	15.5 (1.3)	100	0.597	0.0 (0.0)	15.5 (1.3)	0.0 (0.0)
0.0075	10.8 (2.0)	100	0.643	0.0 (0.0)	10.8 (2.0)	0.0 (0.0)
0.01	10.1 (2.4)	100	0.68	0.0 (0.0)	10.1 (2.4)	0.0 (0.0)
0.02	9.6 (3.0)	95	0.926	0.5 (0.7)	8.9 (3.3)	0.1 (0.7)
0.03	10.1 (2.3)	95	1.184	1.6 (1.3)	8.4 (2.8)	0.1 (0.4)
0.04	10.9 (2.6)	95	0.607	2.5 (1.1)	8.3 (2.6)	0.1 (0.2)
0.05	11.2 (2.4)	90	0.534	3.7 (1.7)	7.3 (3.1)	0.2 (0.6)
0.1	11.7 (3.0)	90	0.569	5.7 (2.1)	5.9 (2.5)	0.1 (0.3)
2.0	12.1 (3.7)	80	0.615	6.5 (2.7)	5.3 (3.3)	0.2 (0.5)

Table 10.2: Main results of simulations in which the agents' propensities changed via reinforcement learning and habituation, for selected values of  $ha_2$  (shown in the left hand column). Data correspond to 20 simulations for each value of  $ha_2$  (standard deviations in parentheses). 'Pop'n' refers to the mean total population of agents at the end of 1,000 rounds; '% Prop Rights' indicates the percentage of 20 simulations in which property rights emerged (this is when all the agents alive at the end of Round 1,000 had negative propensities to steal); 'Max Prop Steal' is the mean propensity to steal at which the agents' propensities peaked in the simulations; 'No. Doves' is the mean number of doves alive at the end of 1,000 rounds; 'No. PA' refers to the mean number of passive-aggressive agents alive at the end of 1,000 rounds; and 'No. Hawks' shows the same data for the number of hawks alive.

$ha_2$  are associated with doves emerging and an increase in the likelihood of the scenario in which one or two hawks survive, bully the doves, and property rights do not emerge.

The second phenomenon is when small values of  $ha_2$  help to catalyse property rights (here, habituation is weak relative to reinforcement learning). We can see this most clearly in Table 10.2 when  $ha_2$  is increased from 0 to 0.0001.

Fig. 10.26 illustrates what happens in these simulations. It shows the agents' propensities to steal in the first 400 rounds of a typical simulation when  $ha_2 = 0.01$ . The chart shows how the agents split into two different strategies (Al Capone and passive aggressive agents) in a much more extreme way than in the default simulations: Al Capone agents saw their propensities to steal increase (significantly) above 1 and, equivalently, the propensities of passive aggressive agents decline (well) below 0.

This bifurcation of the agents' propensities to steal looks similar to that seen when reinforcement learning was switched off; however, the results are significantly different with respect to the agents' propensities to defend. In these simulations, reinforcement learning meant no doves emerged: all of the agents with negative propensities to steal had propensities to defend well above 1.




Figure 10.26: A time series of the agents' propensities to steal over the first 400 rounds of a typical simulation when the agents' propensities change as a result of reinforcement learning and habituation (here,  $ha_2 = 0.01$ ). The red lines depict 'hawks', whose propensities increase above 1, and who all died. The blue lines depict passive-aggressive agents (who all survived). In these simulations there were no doves.

The Al Capone agents always died off because fight costs were debilitating. All of the passive aggressive agents survived in all of the simulations: after the Al Capone agents had died off, these agents thrived in an environment in which property rights were established, and specialised. Fig. 10.27 below shows a time series of the agents' total resources in the same simulation as that shown in Fig. 10.26 but over 1,000 rounds: the passive-aggressive agents' resources decline while the Al Capone agents are alive but then increase after they die.




Figure 10.27: A time series of the agents' total resources over 1,000 rounds (taken from the same simulation as that depicted in Fig. 10.26 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows that all the 'hawks' died and all the passive-aggressive agents survived.

### Agent-Level Data: Reinforcement Learning and Habituation

The model records the contributions of reinforcement learning and habituation to changes in the agents' propensities to steal. These data show some particularly interesting results when habituation is relatively low so let us examine them in more detail.

Three figures are shown below. The first, Fig. 10.28, shows a passive-aggressive agent's propensity to steal over the first 100 rounds of a typical simulation when  $ha_2 = 0.005$ , and four different contributions to that propensity over the same period (when reinforcement learning contributed positively and negatively, and when habituation contributed positively and negatively). See the figure's caption for a more detailed description of the data. The figure is reproduced from Chapter 6 for convenience.

The second figure (10.29) below shows the same data but over 1,000 rounds; and the third figure (10.30) below shows the equivalent data for a 'hawk' while it was alive, taken from the same simulation as the previous two figures.

In Fig. 10.28, the agent's propensity to steal started at 0.5 and declined to below zero in Round 26. The contributions data show us that the net contribution of reinforcement learning over the first 26 rounds was -0.61 and the net contribution of habituation was +0.03 over the same period. This tells us that this agent came to respect others' property because of reinforcement learning and despite a positive contribution (albeit small) from habituation. This same pattern was seen for all passive-aggressive agents and doves.

Fig. 10.29 shows that the contributions from reinforcement learning stabilized after the last hawk died (in Round 235). The reason for this is that the feedback from transactions to propensities to steal was much smaller for transactions than fights. The contribution from habituation was also larger than transactions so that by the end of the simulation, habituation dominated reinforcement learning in terms of contributions. The agent's propensity to steal was -228.8 at the end of the final round: the net contribution of reinforcement learning was -0.6 and the net contribution of habituation was -228.2.

For passive-aggressive agents and doves, therefore, we observe an interesting result: property rights emerge because of reinforcement learning but in terms of the doves' mental models, habituation comes to dominate over time. Put another way, the 'rule' of property rights emerges through reasoning and learning but, eventually, it resembles a habit.

Fig. 10.30 shows the equivalent data for a hawk in the 235 rounds it was alive. Here, the agent's propensity to steal starts at 0.5 and exceeds 1 by the end of Round 18. The net contribution of reinforcement learning over this period is -0.5 and the net contribution of habituation +1.18.




Figure 10.28: A time series of a passive-aggressive agent’s propensity to steal (the black line) over the first 100 rounds of a typical simulation when propensities change as a result of reinforcement learning and habituation (here  $ha_2 = 0.005$ ). The four ‘areas’ denote contributions to the agent’s propensity to steal over time: the upper (dark red) area shows positive contributions by habituation; the second (dark blue) area shows positive contributions by reinforcement learning; the third (light blue) represents negative contributions by reinforcement learning; and the lower (light red) area represents the negative contribution by habituation. This chart shows how the agent’s propensity to steal declines below zero (this occurred in Round 26) mainly as a result of reinforcement learning.

It is noteworthy that, *ceteris paribus*, reinforcement learning on its own would have seen this agent’s propensity to steal decline to zero. This particular agent’s experience in the simulation, however, meant that its propensity to steal rose more quickly than the passive aggressive agent depicted in figures 10.28 and 10.29, resulting in a persistent positive contribution by habituation.

By the time the hawk died at the end of Round 235, its propensity to steal was 7.41 (the contribution of reinforcement learning was -10.44 and habituation +17.35).

## 10.4 Discussion of Results

As with the results from the Market Emergence Model, here we focus on whether the property rights we observed in the default simulations meet the definition of institutions adopted in this thesis (Section 10.4.1); and whether those we observe in the habituation experiments do so also (Section 10.4.2). Here, a third section is added which considers whether property rights emerged in our simulations via generalized Darwinism (10.4.3).




Figure 10.29: The same time series as shown in Fig. 10.28 above but for 1,000 rounds. See the above figure for descriptions of the data. This chart shows how habituation came to dominate the agent's propensity to steal over time.




Figure 10.30: A time series of a hawk's propensity to steal (the black line) over the 235 rounds it was alive. The data were taken from a typical simulation when propensities changed as a result of reinforcement learning and habituation (here  $ha_2 = 0.005$ ). See Fig. 10.28 above for a description of the data. The chart shows that the agent's propensity to steal increased above 1 because the net impact of habituation exceeded that of reinforcement learning.

### 10.4.1 Organic (Unplanned) Institutions

The definition adopted in this thesis is repeated here for convenience: institutions are “durable systems of established and embedded social rules that structure social interactions.” (Hodgson, 2006a, p. 13). Organic institutions meet this definition but are unplanned.

Looking at this in more detail:

- The property rights we observe appear durable for at least three reasons. First, there is a consistency between what agents expect of other agents and their actions; second, when all the living agents' propensities are negative, they can no longer become positive (lock-in); and, third, when we allow children to be endowed with positive propensities to steal (see the next chapter), we find that the agents' defence of property protects these property rights. These children either die or they learn to respect property.
- In terms of systems, these property rights appear to be mainly cognitive with no 'artefacts' but they exist within a family of institutions, including language.
- Property rights are both established within and embedded across the population. The lock-in of property rights mentioned above (and the continued decline of propensities to steal over time) ensures this is the case.
- These rights appear to structure social interactions in that the agents are disposed to trade rather than steal others' resources.
- Below we consider whether these property rights ought to be considered rules.

Should we view the emergent property rights as unplanned? It appears so. It is tempting to believe that the decline of propensities to steal below zero is inevitable after agents learn to defend their property; however, this is an oversimplification. There were six patterns observed in addition to selection pressure which in general led to those with higher propensities to steal dying. It seems reasonable to state that the emergence of property rights across the population was both unplanned and surprising.

### Rules

Do the property rights we observed fit Hodgson's definition of rules? Recall that the "term *rule* is broadly understood as a socially transmitted and customary normative injunction or immanently normative disposition, that in circumstance  $X$  do  $Y$ ." (Hodgson, 2006a, p. 3, emphasis included).

In terms of **socially transmitted**, we can say that while emerging, it is better to think of property rights as 'socially constructed' due to co-adaptive reinforcement learning, rather than socially transmitted. We concluded the same for single market institutions in Chapter 7.

However, when we relax the assumption of infinite lives, social transmission seems to take place: 'replication' occurs from one generation to the next. The mechanism here is an assumed one whereby children are born with propensities equal to the mean of their



parents' propensities. However, this represents children learning from their parents, and we can assume this occurs through language, at least in part. Once emerged, therefore, property rights appear to be 'socially transmitted'.

Looking at rules as **immanently normative dispositions**<sup>14</sup>, it appears that the property rights we observe sit comfortably with the words 'immanent' and 'dispositions'. Furthermore, if we accept the Oxford English Dictionary (OED) definition of a norm, stated in Section 7.7.1, property rights also appear normative given the agents form expectations which prove to be consistent with their experience of other agents.

Moreover, the idea of a rule that **in circumstances  $X$  do  $Y$**  (and not  $Y^*$ ) appears consistent with our property rights. Once again, the "not  $Y^*$ " component seems reasonable if we consider this pragmatically in the context of a human agent. An individual would be aware that theft is an option but experience leads him or her to trade.

In all, therefore, the property rights we observed in the default simulations appear consistent with Hodgson's definition of institutions as rules which we have adopted in this thesis.

### 10.4.2 Organic Institutions with Habituation?

Recall that when reinforcement learning is switched off, and for high values of  $ha_2$  when reinforcement learning is used, three strategies emerged and in most simulations we saw the 'hawks' die off, leaving the doves and passive-aggressive agents to thrive.

The property rights that emerged in these simulations appear to meet our definition of institutions, broadly for the reasons discussed above. However, habituation appeared to make them more durable, established, and embedded in the population.

Furthermore, it seems reasonable to categorise these property rights as 'unplanned': the movements in the propensities to steal were driven entirely or mostly by habituation.

Looking at the experiments with reinforcement learning and weak forms of habituation, the property rights that emerged also appear to meet our definition of institutions, also for the reasons discussed above; and they also appear unplanned.

In summary, it appears that the property rights that emerged in the default simulations and in all the habituation experiments all meet the definition of organic institutions adopted in this thesis.

---

<sup>14</sup>Again this phrase seems more relevant than "customary normative injunction" in Hodgson's definition of rules.

### 10.4.3 Generalized Darwinism?

Here we argue that property rights emerged in a way that was consistent with generalized Darwinism as described in Section 2.3.3.5. Most importantly, selection pressure ensured that agents with higher propensities to steal died and those with lower propensities survived.

Let us look at how variation, selection, and durability worked in the simulations.

#### Variation

This occurred through changes in the agents' propensities to steal and defend, which in turn resulted from co-adaptive reinforcement learning (and habituation when this was included in the mental models). Recall the six 'patterns' described above.

Clearly, this variation is not 'blind' in the same way that genetic mutations are but here that does not matter: we stated in Chapter 2 that in generalized Darwinism, variation can result from phenomena like conscious deliberation or reinforcement learning.

Variation was most clear in simulations with both reinforcement learning and a weak form of habituation: two very distinct strategies emerged among the agents, which we referred to as 'Al Capone' and 'passive-aggressive'. Indeed, we saw three strategies emerge when habituation was stronger and when agents' mental models only changed as a result of habituation.

#### Selection

We stated in Chapter 2 (restated here for convenience) that selection is "a process of sifting and preservation of fortuitous adaptations." (Hodgson, 2003a, p. 89).

In both the default simulations and when a weak form of habituation was added to the agents' mental models, Al Capone agents tended to die before passive-aggressive agents. While there were six patterns, or forces, acting on the agents, the debilitating effect of fight costs impacted the Al Capone agents the most. It appears that the passive-aggressive strategy was selected for.

The results of simulations with stronger habituation and when mental models changed only as a result of habituation indicated that the selection of agents who respected others' property was conditional on the proportions of the three strategies observed. In most simulations, the doves and passive-aggressive agents survived and the hawks died. However, one or two hawks could thrive if there were significantly more doves than passive-aggressive agents (until the doves all died).

## Durability

In Chapter 2 we stated that durability ensures “that much of the pattern and variety is passed on from one period to the next.” (Hodgson, 2003a, p. 89).

A number of factors ensured that after property rights emerged (i.e., when all the living agents’ propensities to steal were below zero), they were sustained:

- the lock-in of these propensities when they were all below zero;
- the assumption that children were born with propensities equal to the mean of their parents, i.e., they inherited respect for (and defence of) property;
- the fact that agents would defend their property against any children born with positive propensities to steal, leading to these children dying; and
- consistency between agents’ expectations of other agents’ propensities and their actions.

Now that we have discussed the results of the null and default simulations, and experiments that incorporate habituation, let us look at the results arising from the exploration of the parameter space.



# Chapter 11

## Property Rights Model: An Exploration of the Parameter Space

With great power comes great responsibility.

– Uncle Ben to Peter Parker, Spiderman

In this chapter and Appendix C we adjust various parameters used in the default simulations presented in the last chapter. We also explore what happens to the results when we change our approach in various parts of the model, e.g., adopting substantive rationality when agents are strangers.

As stated previously, we are as interested in knowing the conditions under which property rights emerge as much as when they do not.

Appendix C contains a detailed analysis of the results of ten different experiments, each of which adjusts a single parameter, or multiple parameters, or some structural part of the model.

This chapter contains a summary of that work (Section 11.1).

It is worth stating here that property rights and defence of property emerge in a large volume of the parameter space, i.e., the results presented in the previous chapter are broadly robust.

## 11.1 Summary of Results

- **Agents' starting propensities.** In the default simulations, each agent was born with propensities to steal and defend drawn from a normal distribution with a mean of 0.5 and a standard deviation of 0.1. In these parameter tests, five different 'initial conditions' tests were conducted for the agents' propensities to steal and defend: we found **the results of the default simulations were replicated in all of these tests.**
- **Cost of fighting.** Various tests were run with different fight costs (the default value  $c$  was 0.3 for each resource): we found that **these costs had to be between approximately 0.1 and 0.6 for the results of the default simulations to be replicated.** If fight costs were less than 0.1, the agents were less discouraged from stealing so their propensities to steal failed to decline to below 0. If fight costs exceeded 0.6, the agents became reluctant to interact for fear of incurring the cost of fighting, which meant they failed to learn and their propensities hardly changed.
- **The influence of the counterparty's reduced values on the adjustments of the propensities.** In the default simulations the agents learned from the experience of their counterpart in all interactions (with a weight of  $\beta = 0.5$  applied to counterpart experiences). This weight was adjusted within a range of 0 and 1: **we found these tests had no material impact on the results.**
- **Errors included in agents' estimations of expected pay-offs.** We assumed agents made errors in evaluating each pay-off in potential interactions in the default simulations (accurate pay-offs were adjusted by an error taken from a normal distribution with mean 0 and standard deviation of 0.05). We found that **for the default simulation results to be replicated, the standard deviation of the pay-off error had to be between approximately 0 and 0.6.** If the error was above 0.6, agents were less able to discern between those with whom they should and should not interact. On the whole this led agents to lose more resources when they defended them, which put downward pressure on the agents' propensities to defend. The resulting failure of these propensities to rise above 0.8 (approximately) meant the agents' propensities to steal did not decline.
- **The standard deviation applied to children's propensities to steal and defend.** In the default simulations, children were born with propensities to steal and defend equal to the mean of their parents' propensities. In this parameter test we added an error to these means: the main result was that **the agent population demonstrated resilience to children born with positive propensities to steal** ('black sheep') who typically died within approximately 70 rounds of their

birth. The defence of property (high propensities to defend), which was latent in the agents' mental models was an important factor in this resilience: this was an evolutionary stable strategy.

- **Initial resource endowment.** Agents instantiated at the beginning of the simulations were provided with two resources in their personal resource arrays, each drawn from a normal distribution with mean 200 and standard deviation of 5. These initial resources acted as a 'buffer' that enabled the agents to learn without dying (and for property rights to emerge). In this parameter test we adjusted the mean value: we found **the default simulation results were replicated for values of 180 units or higher**. The lower this value was, the more simulations we observed in which the agent population collapsed to 1-3 agents. When the agents start simulations with more than 200 units of each resource, we find that property rights emerge more slowly than in the default simulations (this is explained in more detail in Section C.6 below).
- **The nature of the feedback from each interaction to the agents' propensities.** Section 9.4 above described how the agents learned from their interaction experiences: their propensities to steal and defend were adjusted by a value that adapted the reduced gain / loss from each interaction. There were three components to this (a rate of change coefficient ( $r$ ), cognitive coarseness ( $\delta$ ), and the use of absolute and relative changes in reduced values). All three of these were adjusted. We found that for the default simulation results to be replicated,  **$r$  had to exceed 0.004 and  $\delta$  had to be less than or equal to 0.9**. In addition, **propensities to defend had to be adjusted by the agents' reduced gain / loss relative to the expected gain / loss** as in the default simulations; and **the propensities to steal had to be adjusted by absolute gains / losses**, also as in the default simulations. This is discussed in more detail in Section C.7 of Appendix C.
- **The reputations architecture.** In these experiments we considered two different approaches to how agents handle strangers, i.e., when they had no reputational information about their (potential) counterpart. In the first approach the agents used substantive rationality to make decisions, and in the second they 'assumed the worst', that any stranger had propensities to steal and defend of 1. **In both experiments we found the impact on the results was negligible**: defence of property and property rights always emerged. In a third set of tests we adjusted the length of the agents' memories for reputational data (which was 20 rounds in the default simulations): we found that memories had to exceed 3 rounds for the default simulations results to be replicated, i.e., **memory was important in institutional emergence**.

- **The weights of grid locations in agents' memories (when choosing a target).** When agents selected a grid target at the beginning of the interaction phase of each round, they used the weights existing in memories for each grid square. In the default simulations these weights were adjusted by +1 for transactions and successful conflicts; and -1 for unsuccessful conflicts and when the agent had been mugged. Here, these crude weights were replaced with the agents' reduced gains / losses. **This had no material impact on the results.**
- **Reducing agent clustering and resource concentration.** One of the features of the default simulation results was the clustering of agents in an area of the grid during the interaction phase (e.g., Fig. 10.14) despite the fact they frequently lost all their resources. This could be viewed as unrealistic. In these experiments we enhanced the ability of the agents to avoid previous fight / mugging locations: we found this slowed down agent interaction but **defence of property and property rights nonetheless emerged**. In a different set of experiments we allowed agents to remove themselves from the grid if they had increased their holdings by more than 2 resource units during the interaction phase: we found this also mitigated but did not eliminate the resource concentration effect and **the main results of the default simulations were nonetheless replicated**. In a third set of experiments we used both of these methods: the interaction of agents was slowed down enough that in 3 of the 20 simulations run for this experiment, defence of property and property rights did not merge and the agent population collapsed to 1-2 agents. This tells us that some concentration of resources was necessary for institutions to emerge.



# Chapter 12

## Property Rights Model: Liberal Legislation

All right, tinkerbelle. You're nicked!

–*Sweeney!* D.I. Jack Regan

In this chapter we examine the relationship between legal rules and the emergence of institutions. In particular, we are interested in whether the former can be used to catalyse the latter.

Four sets of experiments are presented below, three of which introduce legal rules to the experiments developed in Appendix D. In addition, we introduce these rules to a scenario in which the cost of fighting ( $c$ ) was below the range required for property rights to emerge endogenously ( $c < 0.1$ ).

In the first section below (12.1) we look at attempts to use legal rules when the agents acquiesce if they want to trade and their counterpart attempts to steal (so-called ‘Yellow Agents’). This corresponds with Section D.1 of Appendix D.

The second section (12.2) considers the impact of legal rules when fighting costs are ‘too low’. When the parameter space was explored (Chapter 11 and Appendix C) we found that the default simulation results were replicated when the cost of fighting was within a range of approximately 0.1 - 0.6 resource units. Here we assume that the cost of fighting is 0.05.

The final two sections focus on legal rules when power determines the outcome of conflicts. Section 12.3 looks at power via fighting skills, and Section 12.4 includes a wealthy

agent when power results from the agents' aggregated resource holdings. These experiments correspond with sections D.2 and D.3 of Appendix D, respectively.

### Types of Legal Rule

Two legal rules are tested in the experiments reported below.

In the first, legal rules are applied to all interactions with a fine equal to  $\zeta$  for all agents who attempt to steal (without any compensation paid to the 'victims'). Agents are not fined for defending their resources. This legal rule is applied, therefore, to instigating agents in scenarios 3F, 3A, and 4; and to counterpart agents in scenarios 2F, 2A, and 4. Fig. 12.1 below summarises the fight costs incurred and the fines applied to the agents in the six scenarios.

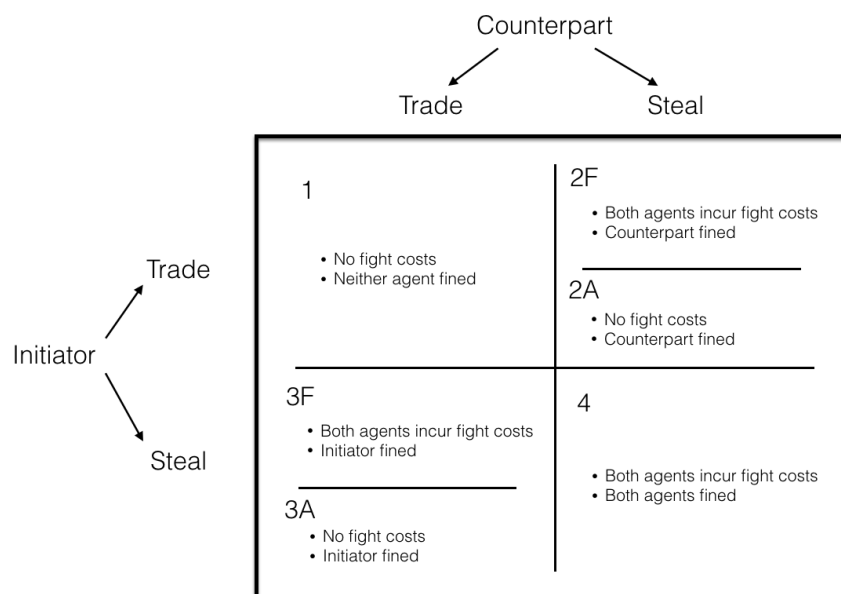


Figure 12.1: A diagram summarising which agents incur fight costs and fines when the legal rules do not include compensation payments to victims.

We assume all the agents are aware of these new parameters when they form expectations of their gains (losses) from any interaction. We also assume they are risk neutral.

The fine is deducted from both resources in the personal resource array of the fined agent in the same way that a fight cost is. For clarity, the reader should assume that a "fine of  $x$ " below refers to  $x$  being deducted from *both* resources in an agent's personal resource array.

A second type of legal rule is also employed: agents are fined as above but now this fine is transferred to any 'victim' as compensation<sup>1</sup>. Fig. 12.2 below summarises the fight

<sup>1</sup>The assumption made in the original model that resources in the agents' personal resource arrays could not be un-consumed or transferred is relaxed in these simulations.

costs incurred, the fines applied, and compensation received by the agents given this type of rule.

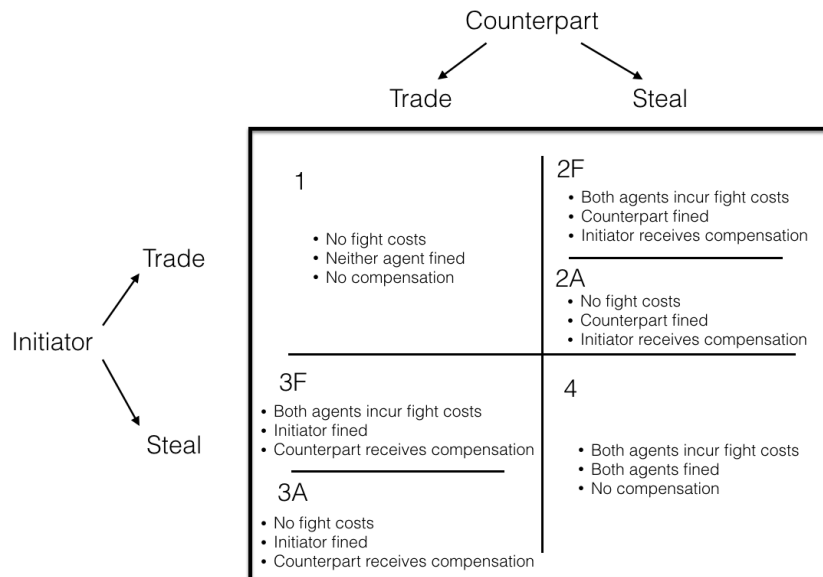


Figure 12.2: A diagram summarising which agents incur fight costs and fines, and which receive compensation, when the legal rules include compensation payments to victims.

### Corruption

We observed that in all four sets of experiments reported below that a rule could be found that meant agent populations always came to respect property rights and never collapsed<sup>2</sup>.

It is tempting to conclude from this that laws and regulations should be used to successfully engineer pro-social institutions. However, this would be naïve because corruption might undermine the efficacy of such rules<sup>3</sup>. In each of the experiments reported below, therefore, we also ran simulations that included corruption. This was done by assuming a probability that the policing authority accepted a bribe from a transgressing agent (which was always less than the fine required by the legal rule).

Moreover, the probability of there being any corruption was assumed to be equal to the median propensity to steal of the agent population (bounded by 0 and 1). This is assumed because we can think of corruption as linked to respect for property: corrupt officials can be thought of as not respecting the agents' property as stated in some legal rule. These mechanics of corruption are explained further in Section 12.1.2 below.

<sup>2</sup>As we shall see, this rule varied between experiments. Also, this result might appear unsurprising given we are 'merely' adjusting pay-offs. We shall see, however, that the impact of legal rules on the population is more complex - they effect the co-adaptive dynamics of whole systems.

<sup>3</sup>Other phenomena might also undermine legal rules, e.g., difficulties / costs in detecting transgressions of the rule. In this chapter we focus on corruption.

It is worth emphasising here that corruption is explored in these experiments with respect to the implementation of a given legal rule and not in the setting of the legal rule.

## 12.1 Yellow Agents

In the Appendix D we found a clear result that when agents always acquiesced, their propensities to steal increased rapidly and plateaued at approximately 0.9. The agent population then collapsed. Can legal rules prevent this outcome by encouraging the emergence of property rights?

### 12.1.1 Legal Rules without Corruption (Yellow Agents)

In the first set of experiments reported in this sub-section we introduced fines to agents who stole and no compensation was paid to ‘victims’. A range of fines were tested in these experiments and we found there was no level of fine that catalysed the emergence of respect for property.

An analysis of the data showed that when fines were less than approximately 1.5 resource units, the fine was too weak to bring about a decline in the agents’ propensities to steal. Fig. 12.3 below shows a fan chart of the agents’ propensities to steal in a typical simulation when the fine was 0.5: the mean propensity of approximately 0.65 was lower than 0.9 observed in Section D.1 but the agent populations nevertheless collapsed here too.

For fines exceeding approximately 1.5 units, agents were reluctant to interact at all. Recall that the initial 25 agents started the simulations with propensities to steal drawn from a normal distribution with a mean of 0.5. When estimating the gains / losses from interaction with potential counterparts the agents knew it was possible they would steal and incur a fine. Therefore, fines exceeding approximately 1.5 units generally discouraged agents from interacting at all.

#### Compensating Victims

The second type of legal rule - when perpetrators were fined and ‘victims’ compensated - was also applied to this ‘Yellow Agents’ scenario. We found that a fine / compensation of 0.4 units or higher was sufficient to catalyse respect for property emerging within the agent population.

Note that in this chapter we refer to the range in which  $\zeta$  catalysed respect for property as the ‘efficacy range’. Here this is  $\zeta \geq 0.4$ .

Fig. 12.4 below shows the ‘cloud’ of agents’ propensities to steal over the first 500 rounds of a typical simulation when the fine / compensation was 0.4 units. Here, only three




Figure 12.3: Fan chart of living agents' propensities to steal when a fine of 0.5 was applied to stealing agents in a 'Yellow Agents' experiment. The black line represents the mean and each band shows one standard deviation (skew-adjusted) away from the mean. This chart shows the fine did not work in that agents did not come to respect each other's property.

agents died and all the surviving agents' propensities had declined to below zero by Round 275.




Figure 12.4: The cloud of agents' propensities to steal over the first 500 rounds of a typical simulation in a 'Yellow Agents' experiment when 0.4 resources units was paid as a fine and compensation. Each dot represents a living agent's propensity to steal in each round. This chart shows that property rights emerged among the surviving agents.

An analysis of the data showed there were two mutually-supporting phenomena: more

agents with resources (and lower propensities to steal) were encouraged to interact because they knew they would be compensated if they traded and their counterpart attempted theft; and the compensation of agents in Scenario 2A generally encouraged propensities to steal lower.

### 12.1.2 Legal Rules with Corruption (Yellow Agents)

For the simulations reported in this sub-section we introduce corruption to the ‘Yellow Agents’ scenario. Here, if an agent attempts to steal then it is possible the agent pays a bribe to the policing authority instead of the full fine as dictated by the legal rule. In addition, if the legal rule includes compensation payments, corruption means compensation was not paid to victims<sup>4</sup>.

We assume that bribes are a fixed proportion of the fine determined by the legal rule, i.e.,

$$B = \lambda\zeta$$

where:

- $B$  is the bribe paid to the policing authority (agents pay this amount in both resources from their personal resource arrays);
- $\zeta$  is the fine prescribed by the legal rule; and
- $\lambda$  is a proportion that determines the bribe as a proportion of the fine ( $0 < \lambda < 1$ ).

If an agent is able to pay a bribe instead of the full fine then we assume they always do. Moreover, an important assumption made in these simulations is that the probability a bribe payment is offered to a stealing agent is equal to the median of the agent population’s propensities to steal.

Let us consider this last point further. There is a good argument that corruption can be linked to property rights: if a legal rule dictates fines (and possibly compensation) then the rule can be interpreted as defining the transfer of resources (property) between agents. Corruption can be viewed here as ignoring what is prescribed by the rule in addition to acquiring resources (the bribe itself).

In these experiments we assume that the policing authority is a representation of the whole population. Here that means the authority is corrupt if agents in the population

---

<sup>4</sup>and victims had no access to any sort of ombudsman scheme.

do not respect others' property, i.e., if the agents' propensities to steal are relatively high; and vice versa.

Corruption is implemented in the model by assuming that after any interaction when a legal rule ought to have been applied, there is a probability the policing authority is corrupt<sup>5</sup>.

The median propensity to steal is used because it is less influenced by outliers and skews in the distribution of the agents' propensities to steal than other first moments. It gives us a propensity not to respect others' property that is 'typical' of the population as a whole.

This approach means that, in a sense, corruption is a cultural phenomenon linked to the population's respect for property.

It is worth highlighting that equating corruption with respect for property introduces an additional 'layer of complexity' to the model because of the co-influence of property rights and corruption. We might expect corruption to undermine the efficacy of legal rules, which in turn might influence the agents' respect for property and, therefore, the prevalence of corruption; etc.

## Results

A range of experiments were run with different values of  $\zeta$  and  $\lambda$ . In all of the experiments, the legal rule included both a fine and compensation.

It seems reasonable to expect the distribution of the initial agents' propensities to steal at instantiation to have an impact on the results because this would determine the median propensity to steal of the whole population at the start of each simulation. Recall that the initial set of agents began each simulation with a propensity to steal drawn from a normal distribution with a mean of 0.5 and a standard deviation of 0.1. The experiments below therefore also explored different mean starting values for the initial agents' propensities (denoted  $\bar{p}^s$ ).

The first noteworthy result was that corruption did not prevent the emergence of respect for property if  $\zeta \geq 1.1$ . This threshold was found by applying the harshest corruption conditions possible in the simulations ( $\lambda = 0$ <sup>6</sup>,  $\bar{p}^s = 1$ , and a mean starting propensity to defend of zero<sup>7</sup>)<sup>8</sup>. However, while respect for property always emerged when  $\zeta \geq 1.1$ , we

<sup>5</sup>A random decimal number is selected from a uniform distribution between 0 and 1. If this random number is lower than the population's median propensity to steal then the policing authority accepts the bribe. If it is higher, the legal rule is applied.

<sup>6</sup>This was a proxy for a trivially small fine.

<sup>7</sup>This was 'harsh' from a corruption point of view because low propensities to defend encouraged propensities to steal (and therefore the incidence of corruption) higher, *ceteris paribus*.

<sup>8</sup>The agents' initial resources were increased to 500 units each to compensate for the harsher conditions.

always observed substantial transitory problems: property rights emerged more slowly than if there had been no corruption, which meant fight costs were higher.

For values of  $\zeta < 1.1$ , corruption undermined the efficacy of the legal rule to the extent that in some of the simulations the agent population never came to respect others' property and the agent population collapsed. Furthermore, the closer  $\zeta$  was to the 0.4 (the bottom of the 'efficacy range'), the more significant was the impact of corruption on the emergence of property rights.

Fig. 12.5 below summarises the results of a range of experiments when  $\zeta = 0.6$ , and when  $\lambda$  was varied between 0 and 0.95 (the Y axis)<sup>9</sup>, and  $\bar{p}^s$  was varied between 0 and 1 (the X axis).




Figure 12.5: A 3-dimensional surface showing the proportion of 'successful' simulations (the vertical axis) in 'Yellow Agents' experiments when the legal rule included fines and compensation of 0.6 resource units. The left-hand horizontal axis, labelled 'Lambda' ( $\lambda$ ), was the bribe stated as a proportion of the fine; and the right-hand horizontal axis, labelled 'Prop. Steal' ( $\bar{p}^s$ ), was the mean starting propensity to steal of the initial agents. Here, success meant the agents population did not collapse and property rights emerged in the surviving population. We can see that success increased monotonically in both  $\lambda$  and  $\bar{p}^s$  when  $\lambda < 0.75$  and  $\bar{p}^s > 0.25$ .

<sup>9</sup>A value of  $\lambda = 1$  was viewed as irrelevant: this would mean a stealing agent would have to pay a bribe equal to the fine.



In Fig. 12.5 the Z axis shows the proportion of 10 simulations that were ‘successful’: here that was defined as when at least 4 agents, all of whom respected others’ property, were alive after 500 rounds.

There are a number of points worth noting. First, if  $\bar{p}^s \leq 0.25$ , the simulations were always successful, irrespective of  $\lambda$ : the simulations started with low levels of corruption and this ensured the agents’ propensities to steal eventually declined. In a sense, the economy started and remained locked in to a situation with low levels of corruption and an efficacious legal rule.

Second, when  $\lambda \geq 0.75$ , the simulations were also always successful, irrespective of  $\bar{p}^s$ . The combination of the fine, the bribe, and any compensation paid to the victims had sufficient ‘bite’ to encourage the agents’ propensities to steal lower. When the data was analysed it was clear that the bribe could be thought of as a *weak fine*, which meant higher values of  $\lambda$  led to what were in effect stronger fines.

Third, for values of  $\lambda < 0.75$  and  $\bar{p}^s > 0.25$ , success increased with higher values of  $\lambda$  (bribes became stronger fines) and decreased with higher values of  $\bar{p}^s$  (the policing authority started more corrupt).

Moreover, an interesting fact to note about higher values of  $\lambda$  was that, in a sense, corrupt officials ultimately undermined their own corruption if  $\lambda$  was set too high: the agents learned to trade rather than steal.

### 12.1.2.1 The Emergence of Successful and Unsuccessful Economies

Some of the simulations that generated the results in Fig. 12.5 showed a phenomenon that is worth considering in more detail: the emergence of successful and unsuccessful economies given the same starting conditions.

To explore this, let us examine a point on the surface of Fig. 12.5, when  $\bar{p}^s = 0.5$  (as in the default simulations) and  $\lambda = 0.25$ . In these simulations, we observed that agents came to respect each other’s property in six of the ten simulations (the agents specialised and bore children). In the other four, they did not and the agent population collapsed.

Note that these experiments went beyond starting with the same set of parameters: the initial agents were instantiated with precisely the same state variables across all the simulations. These included starting resources, foraging strategies, and propensities to steal and defend, which meant the experiments started identically to each other.

Fig. 12.6 below shows the ‘cloud’ of the living agents’ propensities to steal over 500 rounds in a typical simulation when the agents did not come to respect property (an unsuccessful economy). The black line shows the median propensity to steal, i.e., the incidence of corruption.




Figure 12.6: The cloud of agents' propensities to steal over the first 500 rounds of an 'unsuccessful' simulation in a 'Yellow Agents' experiment when fines and compensation of 0.6 resource units were applied and there was corruption ( $\lambda = 0.25$ ). Each dot represents a living agent's propensity to steal in each round and the black line is the median of these propensities. Property rights did not emerge in this simulation.

Fig. 12.7 shows the equivalent data for agents in a successful economy.




Figure 12.7: The cloud of agents' propensities to steal over the first 500 rounds of a 'successful' simulation in a 'Yellow Agents' experiment using the same parameters as Fig. 12.6. Each dot represents a living agent's propensity to steal in each round and the black line is the median of these propensities. Despite the same parameters as Fig. 12.6 being used, here property rights did emerge among the surviving agents.

The reason why the results of the two simulations differed was due to the same positive feedback effect working in two different ways. In Fig. 12.6 the agents on the whole entered a feedback loop that meant increasing propensities to steal gave rise to worsened legal rule efficacy, and further increasing propensities. We can think of the economy as getting

locked in to an unsuccessful state, which includes higher fight costs and a population which never specialised. The agent population collapsed.

By contrast, the agents in Fig. 12.7 were locked in to a ‘successful’ state which meant a relatively low propensity to steal was maintained in the population because the low level of corruption underpinned a legal rule that was increasingly applied. The agents specialised and bore children.

In the context of the complexity sciences, these unsuccessful and successful states resemble ‘basins of attraction’ in the systems’ dynamics. Once a system enters one of these it is very unlikely that it will escape into the other. When initial conditions are near the boundary (known as the ‘separatrix’) between the two basins, stochasticity can determine which way the system tips near the start of a run.

This result, that economies can get locked in to successful or unsuccessful states, correlates with what we observe empirically vis-à-vis economically poor and wealthy countries. Note that this is not a claim that corruption is the only explanatory factor for different productivity levels but these results are consistent with a great deal of empirical evidence concerning the link between corruption and economic wealth, e.g., [Gupta, Davoodi and Alonso-Terme \(1998\)](#).

Moreover, the link between corruption and economic success seen in the above simulations raises a question about policy implications. We might be tempted to conclude that corruption would not be a problem in the long run provided the legal rule was stringent enough, i.e.,  $\zeta$  should exceed 1.1 resource units in the above experiments. However, corruption might also apply to the creation of the legal rule in the first place: it would be wrong to think that policy makers exist outside of the whole system. Put another way, societies in which corruption is endemic might well have corrupt policy makers.

Let us now turn to the second set of experiments.

## 12.2 Low Cost of Fighting

We saw in Chapter 11 that if fight costs are too low (below approximately 0.1), the agents do not learn to respect others’ property. Fig. 12.8 below shows the fan chart of the agents’ propensities to steal over 2,000 rounds of a typical simulation when the cost of fighting was 0.05: it shows these propensities typically declined to -0.1 to 0.5 on average (with a standard deviation of approximately 0.5) in the last 1,000 rounds. In a sense, fight costs act like a discipline on the agents: if they are too weak, property rights do not emerge in the population as a whole.

Note that agents learned to defend their resources when other agents attempted theft irrespective of the cost of fighting. For example, when the cost of fighting was 0.05




Figure 12.8: Fan chart of living agents’ propensities to steal over a typical simulation when the cost of fighting was 0.05 resource units. The blue line represents the mean and each colour band represents one standard deviation (skew-adjusted) away from this mean. Here, the cost of fighting was known to be below the threshold required for property rights to emerge, hence they did not.

units, the agents’ propensity to defend moved above 1 by Round 31 on average (over 20 simulations).

In these simulations we found that agents specialised to some degree (their mean maximum detection probability averaged approximately 0.92 by the end of 2,000 rounds). This is explained in Section C.2.1 of Appendix C<sup>10</sup>. These productivity gains counter-balanced fight costs such that the population typically stabilized at approximately 25 agents.

What happens if we apply a legal rule in simulations when the cost of fighting was 0.05 units? We consider this question now.

### 12.2.1 Legal Rules without Corruption (Low Cost of Fighting)

When we ran the simulations with fines ( $\zeta$ ) only being applied to transgressors (without compensation to victims), we found that the agent population came to respect property rights if  $0.06 \leq \zeta < 1.5$ .

<sup>10</sup>Agents often spent 50-100 rounds without consuming any resources, which meant the resource in which they were most deficient stayed the same for this period. They would then typically focus on foraging for this resource, which improved their foraging skill for it.

It is noteworthy that only a small fine was required in these experiments: agents only needed a small ‘inducement’ to respect others’ property. As mentioned above, agents came to respect each other’s property when fight costs exceed 0.1. If we contrast this with an assumed fight cost of 0.05 we can appreciate why only a modest fine was required to encourage respect for property<sup>11</sup>. This stands in contrast to the ‘Yellow Agents’ experiment above where we found fines alone never enabled the emergence of property rights.

Within the range of  $0.06 \leq \zeta < 1.5$ , all the agents who survived the initial learning phase came to respect each other’s property. Fig. 12.9 below shows the agents’ propensities to steal over the first 500 rounds of a typical simulation when  $\zeta = 0.1$ : property rights emerged comfortably. Note the two Al Capone agents in this simulation (with higher propensities): fines ensured these agents died.




Figure 12.9: The cloud of agents’ propensities to steal over the first 500 rounds of a typical simulation when the cost of fighting was low and a fine of 0.1 resource units was applied without compensation to victims. There was no corruption. Each dot represents a living agent’s propensity to steal in each round and the blue line is the mean of these propensities. The chart shows the fine was sufficient to encourage property rights to emerge.

When  $1.5 \leq \zeta \leq 2$ , whether the whole population came to respect property rights or not depended on the probability of a location being ‘seeded’ by two agents who transacted (this was unlikely but not impossible). For higher fines, this became increasingly unlikely as the risk of paying the fine made interactions non-viable.

For values of  $\zeta > 2$ , agents never interacted even if one had no resources and the other 5: the size of the fine meant it was too big a risk for any agent.

<sup>11</sup>We can think of a fine as a crude proxy for fight costs.

### Compensating Victims

As in the previous section, in these experiments we imposed fines on transgressing agents (equal to  $\zeta$  for each resource) and transferred this amount to the victims.

We found agents came to respect property when  $\zeta \geq 0.06$ . There was no upper bound for the reasons explained in Section 12.1.1.

The results of these simulations were broadly similar to those with legal rules when only fines were applied. When  $\zeta < 0.03$ , the agents did not come to respect property rights (the legal rule was too ‘weak’); and when  $0.03 \leq \zeta < 0.06$ , most of the agents’ propensities to steal declined to below zero but some Al Capone agents were sustained.

Higher values of  $\zeta$  had a different impact on the agents than applying fines only because compensation acted like a ‘carrot’ to the agents in contrast to the ‘stick’ of fines. This was also explained in Section 12.1.1 above. As an example, Fig. 12.10 below shows the ‘cloud’ of living agents’ propensities to steal over the first 200 rounds in a typical simulation when  $\zeta = 3$ . Note how the decline in these propensities to below zero was swift.




Figure 12.10: The cloud of agents’ propensities to steal over the first 500 rounds of a typical simulation when the cost of fighting was low and a fine of 3 resource units was applied along with compensation to victims. There was no corruption. Each dot represents a living agent’s propensity to steal in each round. The chart shows property rights emerged swiftly when this legal rule was applied.

Finally, the agents’ propensity to defend increased overall but none of them exceeded 1 in any simulation. In the default simulations, defence of property acted as a deterrent to other agents against stealing. Here, however, this defence of property was largely

redundant because the legal rules provided the impetus for agents to respect others' property.

### 12.2.2 Legal Rules with Corruption (Low Cost of Fighting)

A range of experiments were designed to explore the impact of corruption on the efficacy of legal rules that were intended to induce property rights within a population when the cost of fighting was 0.05. This involved two sets of experiments, one given a fine only within the efficacy range noted above ( $0.1 \leq \zeta < 1.5$ ) and another with a fine and compensation within the range of  $\zeta \geq 0.06$ .

The results showed that while corruption had a transitory impact (higher fight costs and slightly more deaths on average) on the agent population, it never prevented the eventual emergence of property rights when the legal rules used the above efficacy ranges. This result differs from that noted in the 'Yellow Agents' experiments.

We can understand this result by considering two factors: first, Fig. 12.8 above showed that the agents' propensities to steal tended to decline to approximately 0.25 on average given a fight cost of 0.05 without a legal rule. This meant the median propensity to steal (which determined the incidence of corruption) was also relatively low. Second, as mentioned in the previous section, from the agents' point of view, a bribe can be viewed as a weak fine.

If we combine these two points, we can appreciate how (i) the 'naturally low' propensities to steal meant, eventually, the legal rule was applied most of the time; and (ii) even bribes helped to reduce the agents' propensities below that seen in Fig. 12.8. Hence, corruption never prevented the eventual emergence of property rights.

To illustrate this, Fig. 12.11 below shows the first 500 rounds of a typical simulation when the legal rule included a fine and compensation of 0.1 resource units (the bottom of the 'efficacy range') and under the harshest corruption conditions possible<sup>12</sup>. The black line represents the median propensity to steal of the population - it fell to below zero by Round 113.

We can see in Fig. 12.11 how the propensities to steal of all but one agent had declined to below zero by Round 165. As was typical in these simulations, one agent sustained an Al Capone strategy for a long period (here, until Round 321); however, the efficacy and effectiveness of the legal rule meant this strategy was unsustainable so this agent eventually came to respect others' property.

Let us now consider the impact of legal rules on power.

<sup>12</sup>Recall this was when  $\lambda = 0$ ,  $\bar{p}^s = 1$ , and the mean starting propensity to defend was 0.




Figure 12.11: The cloud of agents' propensities to steal over the first 500 rounds of a typical simulation when the cost of fighting ( $c$ ) was 0.05 and a fine of 0.1 resource units was applied along with compensation to victims. Here, the harshest corruption conditions possible were applied at the beginning of the simulation. Each dot represents a living agent's propensity to steal in each round and the black line is the median of these propensities. The chart shows property rights emerged despite extreme corruption early in the simulation. This corruption was eliminated after the median propensity to steal declined to below zero in Round 113.

### 12.3 Power from Fighting Skill

Recall from Section D.2 that two problems existed when agents were given fighting skills: propensities to steal were generally higher (which meant the population collapsed in about 2/3 of simulations even when no black sheep were born); and the birth of black sheep who became Al Capone agents in the other 1/3 of simulation and when such births were allowed by the model.

From a legal rule perspective, these two represented different challenges because they occurred under different social conditions. We will see that, as a result of this, the efficacy ranges of the legal rules required to encourage property rights to emerge were also different.

Most notably, the agents instantiated at the beginning of each simulation were initially unproductive and they were born in an environment where none of the agents fully respected other agents' property. By contrast, children were born in to an environment in which property rights prevailed, all or most of the agents were fully specialised, and fighting skills were low. Children born with positive propensities to steal (black sheep) therefore faced a very different society to the initial agents.



In the simulations run for this section, black sheep were always allowed, which meant we could explore both of these challenges.

### 12.3.1 Legal Rules without Corruption (Power from Fighting Skill)

The key question here is whether legal rules prevent the emergence of these Al Capone agents in either situation.

When we ran simulations with fines only, we found a narrow range of fines resulted in a population that respected property and that was resilient to black sheep:  $1 \leq \zeta \leq 1.4$ . Above this range, agents generally avoided interacting (which we have seen before) and below it the fine was too weak to prevent black sheep bullying the agent population to the point of collapse.

Fig. 12.12 below shows the ‘cloud’ of agents propensities to steal over 2,000 rounds in a typical simulation when  $\zeta = 1$ . Five black sheep were born in this simulation, all of whom died within approximately 250 rounds of their birth.




Figure 12.12: The cloud of agents’ propensities to steal over a typical simulation when the outcome of fights was determined by relative fighting skills, black sheep were allowed, and a fine of 1 resource unit was applied. Each dot represents a living agent’s propensity to steal in each round. The chart shows five black sheep were born but they all died. In addition, the agent population was resilient to these black sheep: their propensities to steal remained negative and they survived.

Fig. 12.13 below shows the time series of all the living agents’ fighting skills in the same simulation. The initial agents came to respect each other’s property (by Round 192), after which they never fought until the first black sheep was born in Round 630. The agents’ fighting skills had declined to below 1 on average (having previously peaked at approximately 140) by the time the first black sheep was born.

We can see from Fig. 12.13 that the fighting skills of the five black sheep increased quickly after their birth: the fighting skills of the other agents increased too but not as much: fights with the black sheep were distributed across the population.




Figure 12.13: Time series of the agents' fighting skills in the simulation depicted in Fig. 12.12 above. The chart shows how the five black sheep became skilled at fighting; however, the combination of fines and fight costs meant they all died.

From the simulations run in Appendix D (Section D.2) we know that the first black sheep would have bullied the parent population until it collapsed. When a fine of 1 was levied on this agent, however, its Al Capone strategy was unsustainable and it eventually died. This was the case for all five black sheep agents in Fig. 12.12. Note that most of the children born were instantiated with negative propensities to steal (i.e., they were not black sheep) - these survived and the population ultimately reached approximately 43 agents, all of whom respected others' property.

### *Two Challenges*

An analysis of the data showed that whereas a fine of  $1 \leq \zeta \leq 1.4$  was required to prevent black sheep from bullying the population to death, we found that a range of  $0.16 \leq \zeta \leq 1.4$  was required to encourage the initial set of agents to respect each other's property, i.e., a smaller fine sufficed. This was because the social conditions were different between the two challenges, as mentioned above.

Black sheep were born in to a population of agents, all of whom respected others' property, were fully specialised (and therefore productive), and with relatively low fighting skills. This was a very fertile environment for an Al Capone strategy and it meant a much more stringent fine was required to make this strategy non-viable. Indeed, we can say that a fine of  $0.16 \leq \zeta < 1$  would have sufficed to encourage the initial agents to respect property but not the black sheep.

### Compensating Victims

The fine and compensation required to encourage the whole population to respect property (including the surviving initial agents and any black sheep) was  $\zeta \geq 0.5$ . We saw again that there was no upper bound for the reasons discussed in Section 12.1.1.

Fig. 12.14 below shows the ‘cloud’ of the agents’ propensities to steal over 2,000 rounds in a typical simulation when  $\zeta = 0.5$ ; and Fig. 12.15 shows the agents’ fighting skills in the same simulation.




Figure 12.14: The cloud of agents’ propensities to steal over a typical simulation when the outcome of fights was determined by relative fighting skills, black sheep were allowed, and a legal rule applied fines and compensation of 0.5 resource units. Each dot represents a living agent’s propensity to steal in each round. In this simulation, seven black sheep were born: four of them were ‘reformed’ by the legal rule (their propensities to steal declined to below zero) but three of them died (those born with propensities above 2.3).

Seven black sheep were born in this simulation and we can see a clear difference with the simulations when fines only were applied (Fig. 12.12): some of the black sheep were ‘reformed’ such that they came to respect other agents’ property. In the simulation depicted in figures 12.14 and 12.15, the four black sheep born with propensities to steal of below 2.3 were reformed and those with propensities above 2.3 died. The latter did not reform quickly enough before their resource reserves became depleted from fines and fight costs.

Fig. 12.15 shows how the fighting skills of the black sheep who survived initially increased but then declined rapidly after their propensities to steal fell below 0.

These results beg the question of why black sheep were reformed in the simulations when compensation was paid to victims and not when a fine was levied without compensation. After all, black sheep were fined twice as much in the first set of simulations than in the second set, which on its own ought to lead to the opposite result.




Figure 12.15: Time series of the agents' fighting skills in the simulation depicted in Fig. 12.14 above. The chart shows how the black sheep again became skilled at fighting. However, the four 'reformed' sheep saw their skills decline when their propensities to steal fell below zero.

An analysis of the data indicated that it was the behaviour of the other (i.e., non-black sheep) agents that made the difference. Specifically, when compensation was paid to victims the agents who respected others' property and who held no resources (which happened mostly because they had been robbed by the black sheep) used the legal rule to passively-aggressively steal from the black sheep (we saw this phenomenon earlier).

These agents would initiate an interaction and offer to trade (despite having no resources) knowing the black sheep would attempt to steal (assuming its propensity to steal was close to or above 1). The initiator would then either acquiesce or defend their resources - either way it would not lose any resources because it had none. However, a fine would be levied on the black sheep that would then be transferred to the initiators.

This strategy was persistently used against the black sheep to the extent their propensities to steal declined: they learned from their counterparts that it was preferable to trade. Black sheep born with higher propensities to steal died before coming to respect property rights; but those born with lower propensities were 'reformed' quickly enough to survive.

### *Two Challenges*

In the previous sub-section we noted that a less stringent fine was required to encourage the initial set of agents to respect property than it was any black sheep. We found the same for legal rules that also included compensation: if  $\zeta \geq 0.03$  the initial agents came to respect others' property. We can say, therefore, that fines and compensation within the range  $0.03 \leq \zeta < 0.5$  would have led to the initial agents respecting property but not black sheep. Fines and compensation of  $\zeta \geq 0.5$  were required to prevent that.

### 12.3.2 Legal Rules with Corruption (Power from Fighting Skill)

Two experiments were run to test the impact of corruption. First, simulations were designed in which a fine was applied to transgressors (with no compensation to victims) of  $\zeta = 1$ , i.e., at the bottom of the ‘efficacy range’. As was done in the previous sections, we applied a set of parameters that was designed to maximize the impact of corruption.

In the second experiment, simulations were designed in the same way but assuming a fine and compensation was paid of 0.5 (also at the bottom of the efficacy range).

We observed across all simulations in both experiments that property rights emerged among the initial agents who survived the learning phase. These agents then specialised and bore children.

Some of these children were black sheep but the policing authority was never corrupt when these black sheep were born because the median propensity to steal was (significantly) negative. The efficacy of the legal rule was absolute by this stage of the simulations.

In summary, therefore, provided legal rules were designed with fines / compensation within the more stringent efficacy ranges noted above, *corruption did not interfere with the eventual emergence of property rights* nor with ensuring the agent population was resilient to black sheep. However, fighting costs were notably higher than they would have been without corruption.

## 12.4 Power from Accumulated Resources: A Wealthy Agent

In Section D.3 of Appendix D we observed that when a single wealthy agent was included in the agent population (with 8,000 units of each resource), the population always collapsed. The wealthy agent adopted an Al Capone strategy - its propensity to steal increased to well above 1 and it bullied the other agents until most of them died.

### 12.4.1 Legal Rules without Corruption (Power from Resources)

For this experiment we applied legal rules to all the agents in order to explore whether this prevented the wealthy agent from over-exploiting the rest of the population.

As with the ‘Yellow Agents’ experiments, fines on their own were never fully successful in preventing the wealthy agent from exploiting the population. There was a range of

finer (approximately  $0.8 < \zeta < 1.2$ ) in which the wealthy agent came to respect the other agents' property in 50 - 80% of the simulations; but no range existed in which fines were always successful.

### Compensating Victims

When compensation was paid to victims and fines levied on perpetrators, we found that the wealthy agent always came to respect others' property when  $\zeta \geq 0.3$ . Again there was no value of  $\zeta$  above which compensation and fines did not work, for reasons discussed in Section 12.1.1.

Fig. 12.16 below illustrates the 'cloud' of the agents' propensities to steal over the first 500 rounds of a typical simulation when the legal rule included a fine and compensation of 0.4 units. The wealthy agent's propensity is shown as a red line: it came to respect other agents' property.




Figure 12.16: The cloud of the agents' propensities to steal over the first 500 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 8,000 units of each resources) was included in the population. A fine of 0.4 resource units was applied to transgressors alongside compensation to victims. Each blue dot represents a (non-wealthy) living agent's propensity to steal in each round and the red line represents the wealthy agent. In this simulation the legal rule encouraged the wealthy agent to adopt a Passive Aggressive strategy, which meant its propensity to steal became negative quite quickly. Ultimately, property rights emerged across the whole population.

### 12.4.2 Legal Rules with Corruption (Power from Resources)

In the experiments run for this sub-section (which test for various values of  $\lambda$  and  $\bar{p}^s$  within the efficacy ranges identified above) we found that corruption had no impact on the wealthy agent respecting other agents' property eventually for reasons we saw in the last section. When a majority of the agents respected others' property<sup>13</sup> there was no corruption. In many of the simulations corruption did have a short-term effect which meant the wealthy agent's propensity to steal increased to above 1; but it eventually declined to below 0.

To help us understand this, Fig. 12.17 below shows the 'cloud' of the agents' propensities to steal in a typical simulation under the harshest corruption conditions possible, with a legal rule of  $\zeta = 0.5$  (note  $\bar{p}^s = 1$  in this scenario). The red line shows the wealthy agent's propensity to steal and the black line is the median propensity of the whole population.

An analysis of the data indicates that the propensities to steal of the non-wealthy agents declined from approximately 1 because of the cost of fighting (as in the default simulations) and then declined further as the policing authority became less corrupt (the efficacy of the legal rule increased).

The wealthy agent saw its propensity to steal *increase* initially because corruption was rife ( $\bar{p}^s$  started at 1) and the legal rule had no effect.

The propensities to steal of the non-wealthy agents entered a positive feedback loop whereby reduced corruption encouraged these propensities lower (which further reduced corruption). This happened until all of these propensities had declined to below 0 (by Round 189 in Fig. 12.17).

What happened to the wealthy agent's propensity to steal depended on the efficacy of the legal rule. We know from earlier simulations that a fine and compensation of  $\zeta = 0.5$  ought to encourage it to respect other agents' property provided there was no corruption. The black line in Fig. 12.17 shows the median propensities to steal of all the agents (the degree of corruption): we can see that the wealthy agent's propensity to steal generally increased when this median was above approximately 0.5 and fell when it was below this value.

In summary, we can say that this was a second set of experiments in which we observed a majority of agents who respected property rights and who, as a result, applied a credible legal rule that affected those with power. In the case of black sheep in Section 12.3.2

<sup>13</sup>and the median propensity to steal was negative.




Figure 12.17: The cloud of the agents' propensities to steal over the first 500 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 8,000 units of each resources) was included in the population. A fine of 0.5 resource units was applied alongside compensation to victims and the simulation was started with the harshest corruption conditions possible. Each blue dot represents a (non-wealthy) living agent's propensity to steal in each round and the red line represents the wealthy agent. The black line is the agents' median propensity to steal. In this simulation the propensities to steal of the non-wealthy agents were driven lower by fight costs but, initially, the wealthy agent became an Al Capone agent. However, as the efficacy of the legal rule improved, this agent paid more fines and fewer bribes, which ultimately led it to respect property rights.

above, this power was due to fighting skill; and in the experiments discussed above, it was a result of wealth differentials.

It should be noted that transitory problems were observed in both experiments: in the previous sub-section, black sheep fought other agents and were either starved to death by the legal rule or they were 'reformed'; and, in this sub-section, all the agents incurred substantial fighting costs before respect for property emerged across the whole population, including the wealthy agent.

## 12.5 Conclusion

In this section we discuss how the property rights that arise after the imposition of a legal rule appear to meet our definition of institutions.



Much of the content discussed in Section 10.4.1, where we concluded the endogenously emerged property rights meet our definition, is relevant here too. Below we outline pertinent differences only.

For convenience, we repeat the definitions of institutions and rules assumed in this thesis.

Institutions are “durable systems of established and embedded social rules that structure social interactions.” (Hodgson, 2006a, p. 13).

The “term *rule* is broadly understood as a socially transmitted and customary normative injunction or immanently normative disposition, that in circumstance  $X$  do  $Y$ .” (Hodgson, 2006a, p. 3, emphasis included).

The main points are:

- In the property rights that emerged in the above simulations, we can add legal rules to the ‘systems’ that makes up these institutions.
- ‘Social transmission’ still occurred from one generation to the next when we relax the infinite lives assumption and after property rights have emerged.
- Most interestingly, property rights now appear to be both customary normative injunctions and immanently normative dispositions. The former is now relevant because the legal rule represents an injunction which can be interpreted as customary. Indeed, the customary nature of the property rights is seen in both agent interactions and in the lack of corruption.
- Here, property rights cannot be categorised as ‘unplanned’ so they are not ‘organic’. However, it is noteworthy that they emerged via the same co-adaptive reinforcement learning observed in the default simulations but this time with the support of a legal rule.

Finally, the property rights that emerged after the imposition of legal rules seem to sit comfortably with the following quote from Hodgson, which is reproduced from the Introduction for convenience:

For laws to become rules in the sense discussed here, they have to become customary. As discussed later in this essay, there are examples of laws that are widely ignored and have not acquired the customary or dispositional status of a rule. Ignored laws are not rules. For new laws to become rules, they have to be enforced to the point that the avoidance or performance of the behavior in question becomes customary and acquires a normative status (Hodgson, 2006a, p. 6)



# Chapter 13

## Conclusion

I am sorry to have made such a long speech, but I did not have time to write a shorter one.

Winston Churchill

In this concluding chapter we focus on five topics which are aligned with the five sections below.

The first topic answers the first research question. Here we delve more deeply into the mechanisms by which organic institutions emerge / immerse in the simulations based on the two models. A generalised framework is described which contains the different mechanisms seen (to varying degrees) in these simulations but we also note an idiosyncratic mechanism that plays a role in the emergence / immergence of property rights: selection within a generalised Darwinian process (discussed in Section 10.4.3 above).

In this first section we also consider whether the institutions that emerge in the simulations should be considered as forms of spontaneous order.

The second topic answers the second research question regarding ‘liberal legislation’.

In the third section we revisit the question of whether models of institutional emergence can start from an institution-free state of nature or whether pre-existing institutions should be assumed. To that end, we look briefly at how the simulations based on the second model gave rise to four distinct ‘layers’ within a stratified ontology.

The fourth section looks more closely at the procedure of the research that preceded this thesis, specifically how the models were used as tools of investigation.

The fifth and final section discusses potential future research.

## 13.1 First Research Question: Symmetry Breaking via Immergence and Emergence

This section seeks to answer the first research question, which is stated below for convenience:

Can organic institutions emerge spontaneously across a population while also immersing within individuals' mental models via reasoning, learning, and habituation?

In light of the simulation results presented in chapters 7 and 10, the answer to this question seems to be 'yes'. The agents in these simulations were endowed with mental models that allowed them to reason and make decisions under conditions of uncertainty; and we explored versions when changes to these mental models occurred via reinforcement learning and/or habituation. The agents' mental models co-adapted in such a way that organic institutions appeared to immerge and emerge provided their environment was sufficiently enabling.

The rest of this section discusses the mechanisms we observed in more detail. First, we describe a generalised framework of these mechanisms (Section 13.1.1). This is followed by a discussion of different 'downward effects' (Section 13.1.2).

We then look at how closely this generalised framework resembles the simulation results of both models (sections 13.1.3 and 13.1.4). Finally, in Section 13.1.5 we look briefly at how commonality of institutions was achieved in the emergence of markets and property rights.

### 13.1.1 A Generalised Framework of Organic Institutional Emergence

It is important to state two qualifications before proceeding. First, the description below builds on:

- the work of Hodgson in various articles (notably his emphasis on upward and downward 'effects' and the role of habits);
- Hodgson and Knudsen's (2004) traffic convention model and their discussion of upward and downward effects;
- Conte and Castelfranchi's (1995a) concept of *cognitive emergence* as well as the EMIL Project and subsequent work focused on immergence, e.g., Andrighetto, Campenni and Conte (2010), Castellani (2010), and Conte et al (2013); and

- various concepts from the complexity sciences, discussed in Chapter 2.

The added value of the generalised framework below is two-fold: (i) it emphasises the specific roles of both reinforcement learning and habituation, observed in various simulation results discussed above; and symmetry breaking; and (ii) the components appear to form a coherent whole, the totality of which helps to explain (at least in part) organic institutional emergence in the two models' simulations.

It is important to emphasise that symmetry breaking plays a central and cohering role in what follows. We start with an environment prior to the emergence of a new organic institution<sup>1</sup> and show how specific mechanisms give rise to emergence and immergence, resulting in the breaking of symmetry.

The second qualification is that this framework identifies features common to the two models' simulation results: it should not be treated as a complete theory that attempts to explain all organic institutions. It is a step further towards such a theory.

We should add, however, that the results reported in [Hodgson and Knudsen \(2004\)](#) and the four EMIL models (see sections 5.1 and 5.2.3.2, respectively) appear to be broadly consistent with this generalised framework. Symmetry is broken in these simulations too. In fact, below we use Hodgson and Knudsen's results to exemplify some of the mechanisms described.

Fig. 13.1 below includes a visual summary of the generalised framework from the point of view of a single agent. Note that here we focus on the *process* of institutional emergence and comment on the final results at the end.

There are seven 'parts'. The description that follows is fragmented and sequential for reasons of clarity but we can appreciate that agents manifesting these mechanisms can act, learn, and form habits in parallel. Let us briefly list the component parts and then discuss each in turn:

1. **An emerging property** within the external environment (from the point of view of the agent).
2. **Downward effects** which are the influences the emerging property has on the agent.
3. **Upward effects**. These are the impact the agent has on the emerging property.
4. **Positive feedback ('outer emergence')** emphasises that the agents' overall impact on the emerging property is to augment it.

---

<sup>1</sup>As mentioned in the Introduction, we assume pre-existing institutions but focus on the emergence of *new* institutions within a stratified ontology. This is discussed further in Section 13.3 below.

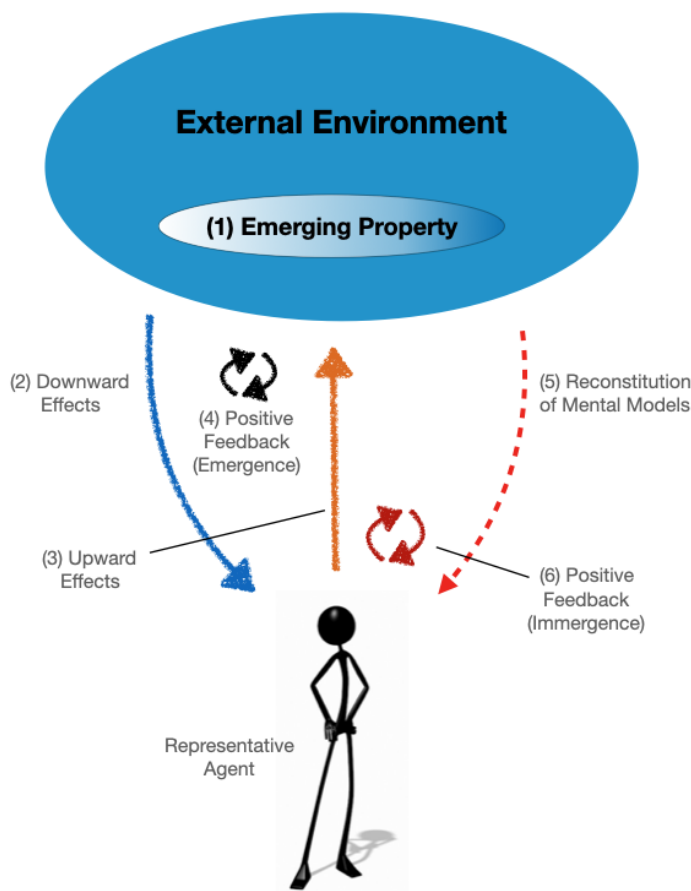


Figure 13.1: An illustration of the generalised framework of organic institutional emergence. This framework contains two ‘learning loops’: the first corresponds to the emergence of a property in the environment; and the second corresponds to the reconstitution of the agents’ mental models. See the main text for a detailed description of the seven features this framework contains.

5. **Reconstitution of Mental Models** is concerned with: (i) the agent learning from the impact of its actions on the external environment (including the emerging property), and (ii) the impact of habituation on its mental model.
6. **Positive feedback (immergence)** emphasises how repeated reconstitution gives rise to an ‘immergent property’ within the agents’ mental models.
7. **An enabling environment** is necessary for the above components to give rise to symmetry breaking.

From these brief descriptions and Fig. 13.1 above we can appreciate that this generalised framework corresponds to “double-loop learning” (Argyris, 1977; and Argyris and Schön, 1978). Interpretations of this mechanism vary slightly but here we refer to the first of

the two loops as when agents make decisions with a given mental model, and the second loop involves a recalibration, or reconstitution, of these mental models<sup>2</sup>.

We should also note that here we distinguish between downward and upward effects for the purposes of clarity but these two should be considered together. This is discussed in more detail below but downward effects correspond to information being input into the agents' mental models, and upward effects are the impact of decisions that arise in part from that information.

Let us now look at each of these mechanisms in detail.

#### 13.1.1.1 An Emerging Property

We start from a state in which there is no emerging property. However, changes in the system, such as perturbations, might give rise to some marginal bias in the external environment. In [Hodgson and Knudsen's \(2004\)](#) simulations, this would be a slight tendency for drivers to drive on one side of the road. This change might also start from the prior emergence of a new stratum in a stratified ontology (such as agents learning to defend their resources).

From an individual agent's point of view, this change occurs 'outside' of its boundaries but it will be represented in its mental model in some way. In fact, as the phenomenon emerges, it might not be recognised as such by the agent as an identifiable concept<sup>3</sup>. The agent might have only an incomplete representation of some local part of that property ([Aoki, 2001](#)).

In the case of [Hodgson and Knudsen's \(2004\)](#) simulations, when a convention is emerging, more than half but less than all of the agents drive on the same side of the road. Each agent will experience this through the information input into its mental model (see Section 5.1.1).

#### 13.1.1.2 Downward Effects

This occurs when the emerging property influences an agent in some way.

There are two parts to this. First, information related to the emerging property will be part of a possibly wider set of information (about, for example, the physical environment and other agents) incorporated into an agent's mental model.

---

<sup>2</sup>We might ask how 'learning' corresponds to a static mental model. Below we look at the example of ants whose mental models are hard-wired: when an ant discovers a food source, it lays down a pheromone trail so other ants can find it. Learning here is a new decision with new information but in the context of a given mental model.

<sup>3</sup>Recall from the Introduction, [Gilbert's \(2002\)](#) two types of cognitive emergence: in the first there is no new cognitive phenomenon whereas in the second, there is.

Second, this information is perhaps best viewed as *influencing* the agent in some way, in the context of its goal(s). Consistent with Hodgson (2011), the emerging property does not fully determine an agent's behaviour.

Note that when an agent makes a decision, its mental model is taken as given. We discuss changes to mental models (reconstitution) separately below.

In the context of Hodgson and Knudsen's (2004) traffic convention, this would mean that if the cars ahead of an agent are (say) driving on the left, the agent would tend to drive on the left given its current mental model.

Note that this description of downward effects is different from Hodgson and Knudsen's (2004) references to strong and weak downward causation / effects. This is discussed further in Section 13.1.2 below.

### 13.1.1.3 Upward Effects

The impact an agent has on the emerging property is referred to as an 'upward effect' here.

As a computational entity (cf Section 2.3.3.1, page 47), an agent will take in information, process it, and reach some decision (even if this means doing nothing). If the agent's decision influences the emerging property in some way, then this is an 'upward effect'. Consistent with previous comments, the agent might not be aware of its impact on the whole property: it simply makes a decision it expects will help it achieve its goal(s).

In the traffic convention example, if an agent chooses to drive on the left because it sees that most the cars ahead of it are doing so, its decision will contribute to the emergence of a convention to drive on the left (and vice versa).

### 13.1.1.4 Positive Feedback ('Outer Emergence')

This component recognises that the combination of upward and downward effects augments the emergent property, i.e., feed back positively to it. This is obvious from the last two items but it is possible for decisions to dampen some phenomenon (negative feedback) so this component clarifies the overall effect.

This is not to suggest that every decision made by every agent will feed back positively to the emergent property. For symmetry breaking to occur, the net effect of these decisions must be positive.

Moreover, this feedback occurs with a given mental model (changes to which are considered below). In principle, 'outer' emergence can occur with fixed mental models. An example would be when ants foraging for food leave a pheromone trail back to the nest that is reinforced by subsequent ants: there is no change in the ants' mental models;



however, this hard-wired behaviour results in self-organisation (Kirman, 1993). Moreover, when the food has all been transported to the nest the trail fades away and the ants return to random wandering.

In the traffic convention model, this positive feedback is clear: an agent who observes people mostly driving on the left will contribute to this emerging convention by also driving on the left given the current state of its mental model.

### 13.1.1.5 Reconstitution of Mental Models

In the simulations discussed in earlier chapters, changes in the agents' mental models can occur through reinforcement learning and habituation.

Reinforcement learning can take a number of forms (including conscious and sub-conscious processes). The critical point is that an agent experiences positive or negative effects (or neither) from its actions from the point of view of its goal(s). There is an enormous literature on this, especially in the fields of machine learning and artificial intelligence; but (as mentioned in the Introduction) this success / failure feedback is the essence of reinforcement learning.

Here, habituation is assumed to take the same form as stated in the Introduction: a tendency to repeat decisions made in the past, which is insensitive to the losses or gains made by these decisions.

In Hodgson and Knudsen's (2004) results, reconstitution takes the form of changes to the habituation variable in the agents' mental models. Interestingly, as noted by the authors, this change was similar to but not the same as reinforcement learning: the habituation variable merely reflects which side of the road a car had driven on in the past.

### 13.1.1.6 Positive Feedback (Immergence)

We can think of immergence as resulting over time from a second positive feedback loop that includes reinforcement learning, habituation, and future actions. This feedback loop is slightly more complicated than the first loop described above: here, the change in the agent's mental model enhances the likelihood that the agent will make the same decision as previously, bringing about the same (or augmented) feedback and a further change in the mental model.

Note that we observed in simulations based on the Market Emergence Model that when habituation is strong relative to reinforcement learning, it can interrupt this positive feedback process.

As stated above, in the traffic convention model, reconstitution occurs through the habituation variable. If an agent chooses to drive on the side of the emerging convention and survives, this variable would change accordingly, increasing the chance that it chooses the same side in the future and repeating the same learning process.

### 13.1.1.7 Enabling Environment

This feature recognises that in practical situations, many of the component parts above have to be enabled by factors not mentioned. These will be context-specific so it is difficult to generalise.

In the traffic convention model, cars have to exist and the agents require a range of cognitive features, including knowing how to drive, what left and right mean, and a whole range of other concepts. Also, agents have to be able to reason in a particular way.

Some of these enabling factors will be pre-existing institutions, which is discussed further in Section 13.3 below.

### 13.1.1.8 Final Results

The above framework is focused on the *process* of institutional emergence / immergence. The end of result will include two main features.

The first feature is an emergent property. In the case of Hodgson and Knudsen's (2004) traffic convention, this is all the drivers driving on the same side of the road. The second feature is an immergent property (agents fully habituated in the case of the traffic convention model).

Note that the corresponding emergent and immergent properties of the simulations based on the two models above were summarised in Section 6.4.7 (p. 211) above and will not be repeated here.

Finally, we have mentioned Hodgson's Klein bottle analogy a few times in this thesis already. The end results of the above framework correspond neatly with this: while the immergent property is not exactly the same as the emergent property, the two are closely associated.

### 13.1.1.9 Mutually Supporting Emergence and Immergence

It is tempting to frame emergent properties as manifestations of their immergent counterparts, i.e., as the outer products of immanent dispositions. This is not unreasonable but it obscures the mutually supporting role the two have while forming (and being maintained).

This happens in two ways. The first and most obvious is via the change in the agents' mental models augmenting the emergent positive feedback loop. In the fourth component above (positive feedback - emergence) we assume a static mental model but changes in the agent's mental model (the fifth component) enhance this feedback process. In the traffic convention example, assuming that a convention is emerging, if an agent's habituation variable becomes more positive or negative (consistent with the emerging convention), it will be more likely to drive on the side of the road of that convention in the future.

The second way is by an enhanced emerging property contributing to immergence. This is because the feedback from the environment changes as the 'outer' property emerges. In the traffic convention example, as a convention emerges, the environment makes it more likely that a decision to drive on the 'conventional' side of the road will be successful.

In a sense, therefore, there is a higher level positive feedback loop between the two positive feedback effects associated with emergence and immergence.

Note that the 'higher' loop identified here is fully consistent with that discussed in the EMIL Project Report and other research that emphasises cognitive emergence and immergence, e.g., [Andrighetto, Campennì and Conte \(2010\)](#), [Castellani \(2010\)](#), and [Conte et al \(2013\)](#), i.e., it is not new.

### 13.1.2 Types of Downward Effects

In previous chapters we noted two types of downward effects: strong and weak, following [Sperry \(1969\)](#) and [Campbell \(1974\)](#), respectively. These were discussed in [Hodgson and Knudsen \(2004\)](#).

Here, strong downward effects are identical to reconstitution in the above framework (for clarity we refer to this as 'reconstitution' but acknowledge that it is a type of downward effect). Weak downward effects correspond with the selection pressure seen in simulations based on the second model, which is discussed further below.

The downward effects described above can be viewed as a third type. Here, this simply refers to the impact an emerging property has on the agents' decisions with a given mental model. This is helpful because it implicitly acknowledges that, in principle at least, inter-group properties can emerge from static mental models. The example of ants' pheromone trails is pertinent.

The main reason for identifying this third type of downward effect is to acknowledge that an emergent property can influence agents' decisions at the point of decision making and not only via reconstitution. This is not to argue that organic institutions always emerge without any counterpart immergence.

Let us now discuss the extent to which the generalised framework described above is relevant for the results of simulations of the two models developed for this thesis.

### 13.1.3 First Model

The framework above is a close fit for the results of simulations based on the first model. The following points are worth noting:

- The emerging properties here are the concentrated transactions at different locations on the torus. These locations are initially seeded by agents bumping into each other during a random walk. The final emergent property is when all transactions occur on one grid square<sup>4</sup>.
- Each agent is aware of a subset of all transactions and these translate into weights in the agents' memories. These mental representations of historical transactions have a downward effect on the agents, i.e., they influence its choice of target location.
- Equivalently, the upward effect is concerned with the impact of an agent's actions (location visited and transactions) on the transactions known to all agents.
- An agent augments the total volume of transactions at some location by: (i) visiting it and transacting; (ii) communicating with other agents about it; and (iii) becoming specialised, which, *ceteris paribus*, increases the volume of transactions.
- Reinforcement learning means that each time an agent visits a location and successfully transacts, its weight in memory increases (offset by memory decay). Also, if habituation is included in mental models, the fact of visiting the location will also increase its weight.
- Immergence takes two different forms here. The first occurs when an agent has multiple locations in memory and, ultimately, one comes to dominate. We can think of this as intra-agent immergence: generally speaking, provided other agents turn up to transact, the location with the highest weight is more likely to see its weight increase, and vice versa.
- The second form of immergence relates to the mechanism described in Chapter 7, when agents at smaller markets are more likely to hear about and visit larger markets than vice versa. This is a form of inter-agent immergence and is essential for system-wide symmetry breaking (the eventual dominance of one market) in these simulations.

---

<sup>4</sup>Or a small number of grid squares when a weak form of habituation is included in the agents' mental models.

- The necessity of an enabling environment for market emergence was made clear in Chapter 8.

In general, we can say that the generalised framework described above fits the first model's simulation results comfortably but there are a few idiosyncratic features worth highlighting, notably the inter-agent immergence mentioned above.

We should also note here how the end result of a single emergent market ensures consistency between each agent's expectations and other agents' actions. The role of institutions as mitigators of uncertainty was an important part of North's work (e.g., [North, 1990](#)) and many researchers have commented on the role of institutions in shaping (and stabilizing) expectations, e.g., [Hayek \(1973\)](#), [North \(1990\)](#), [Ostrom \(1991\)](#), [Aoki \(2001\)](#), [Hodgson \(2006a\)](#), and [Gräbner and Ghorbani \(2019\)](#).

The location weights that agents hold in memory can be interpreted as their expectations of where other agents are likely to be on the torus during the interaction phase. Moreover, when more than one location exists in an agent's memory, this can be interpreted as a state of uncertainty: a 'Roulette Wheel' approach is used to choose between different locations.

In the default simulations, it appears that the immergence of a single location in memory corresponds to an expectation that other agents will go to the same location; and the single emergent market corresponds to the agents going to that location. Put another way, the 'mirroring' of the immergent and emergent properties appears linked to expectation fulfilment.

### 13.1.4 Second Model

The second model's simulation results were 'noisier' than the first model's, and slightly more complicated, but the generalised model is nonetheless relevant for understanding the mechanisms at play.

Note that in what follows we discuss the results from simulations in which a weak form of habituation was included in the agents' mental models. These largely include the results of the default simulations and are more interesting.

- There are three emerging properties relevant to us in these simulations: the group of Al Capone agent; the group of passive-aggressive agents; and resource concentration. The mechanisms described below lead to the latter undermining itself, ultimately; and selection pressures mean the Al Capone agents die off, leaving the passive-aggressive group, who respected property, as the sole surviving emergent property.

- The three mechanisms of (2) downward and (3) upward causation, and (4) positive feedback in the generalised framework above work in the second models' simulations through the resource concentration effect, discussed in Chapter 10<sup>5</sup>. This phenomenon *catalyses* the emergence of the Al Capone and passive-aggressive strategies (explained further below).
- Here, downward causation is the influence of resource concentration on the agents via their mental models, which, again, are assumed to be fixed in this mechanism. Agents nonetheless use these mental models to decide: (i) their target locations on the grid; and (ii) which agents to interact with.
- In terms of upward causation, resource concentration generally leads the agents to choose to interact with others who hold the concentrated resources. This is true of both Al Capone agents and passive-aggressive agents (recall from Chapter 10 how the latter engage in 'passive-aggressive theft'<sup>6</sup>).
- There is, as a result, a clear positive feedback effect in which the agents contributes to the resource concentration effect. But how is this related to the emergence of property rights? We find that resource concentration accelerates the number of interactions between the agents, which leads it to catalyse reconstitution (discussed below). Therefore, while the positive feedback observed in the first model was directly related to the emergence of a market, here, the 'outer' emergence has an indirect impact on the emergence of property rights.
- Reconstitution takes the form of both reinforcement learning and habituation in these simulations. The former includes the six 'patterns', or 'forces', discussed in chapters 6 and 10, which saw each agent's propensity to steal increase or decrease depending on its experiences in interactions. Four of these forces encouraged propensities to steal lower (including the impact of fight costs, which was the most significant) and two of them higher.
- It is clear from the simulation results that the bifurcation between the Al Capone and passive-aggressive strategies arose because of the agents' different experiences of immergence. On the one hand, we saw in Section 7.6 that the effect of habituation on the Al Capone agents overwhelmed that of reinforcement learning. On the other hand, for passive-aggressive agents, reinforcement learning dominated and their propensities to steal declined to - and eventually below - zero. We can think of these two extreme strategies as resulting from two opposing positive feedback loops.

---

<sup>5</sup>See Section 10.1.3.1.

<sup>6</sup>This is described in Section 10.2.2: passive-aggressive agents work out that if they initiate an interaction with an Al Capone agent (holding lots of resources), the latter will try to steal from the former, which gives the passive-aggressive agent a 50% chance of acquiring its resources when defending their own.

- A mechanism that contributed to property rights in these simulations, which was additional to the seven described in the generalised framework above, was selection pressure within a generalised Darwinian process. This was discussed in Section 10.4.3. The Al Capone ‘strategy’ proved to be less advantageous than the passive-aggressive strategy. After these Al Capone agents died, habituation and the benefits of transactions further enhanced emergence, leading to system-wide property rights.
- When this happened, the resource concentration effect, which centres around theft, was replaced with a concentration of transactions in a market.
- As discussed in Chapter 11, property rights only emerged under certain enabling conditions.

The final point to note here is, as in the previous sub-section, about expectations fulfilment.

In Chapter 9 we noted that each agent maintains memories of other agents’ actions such that it could form expectations about their propensities to defend and steal. In the default simulations, these memories employed data from the previous 20 rounds, made up of information gleaned from an agent’s direct experience and that provided by other agents.

When all the living agents respected property rights, they never stole and we observed perfect consistency between expectations and actions<sup>7</sup>.

### 13.1.5 Commonality

In this final sub-section we briefly discuss how markets and property rights became common to agents in their respective simulations.

In Section 4.5.6 we discussed three mechanisms by which institutions might become common within a population: (i) mimicry (after Hayek); (ii) the same rational choice (after Becker); and (iii) population-wide legal rules. Furthermore, in Section 5.1.4 we added a fourth mechanism: symmetry breaking via the co-adaptation of agents’ mental models (after Hodgson and Knudsen, 2004).

It appears that the emergence of a single market fits with the symmetry breaking observed in Hodgson and Knudsen (2004).

---

<sup>7</sup>Note that this was also true in experiments when the agents’ propensities to steal all increased above 1.

For property rights, this symmetry breaking was also observed but only after all the Al Capone agents had died. Put another way, ‘selection’ played a role in the commonality of property rights.

## 13.2 Second Research Question: Liberal Legislation

The second research question asks:

Can ‘liberal legislation’ catalyse institutional emergence when it does not occur endogenously?

As discussed in Chapter 6, the liberal legislation experiments indicate that certain legal rules change a non-enabling environment into an enabling one. This means the answer to the question is clearly ‘yes’. But can we be more specific about the mechanisms at play?

The generalised framework above helps us understand more clearly what happens: perhaps unsurprisingly, legal rules work through the feedback loop between actions and changes in the agents’ mental models, i.e., reconstitution.

We should note, however, that other (indirect) mechanisms are at play that result from the co-adaptation of the agents’ mental models. If an agent’s propensity to steal declines, e.g., after being fined for theft, this changes the social environment for other agents, which in turn changes the feedback on these agents from their actions. For example, the decline in one agent’s propensity to steal will lead to an increase in the number of trades, *ceteris paribus*, which will put more downward pressure on others’ propensities to steal.

If we look more closely at the results of the legal rules experiments, we see that fines either supplement or substitute for fight costs in all of them. This creates a positive feedback loop that had been missing, leading to the emergence of property rights: agents who only trade.

In the context of the above discussion, we should note that the interesting outcome here is not that legal rules can incentivise a change in behaviour (there is an enormous body of work related to this point). The added value is to identify the mechanisms through which legal rules work within a complex and co-adapting environment.

One of the more interesting implications of the mechanisms noted above is that legal rules require that mental models in the second model include reinforcement learning: habituation alone would not be sufficient. Agents have to be sensitive to the feedback from the legal rule.



The final point to make here is to re-emphasise the idea of spontaneous order that is enabled by planning, which was discussed in the Introduction. Legal rules were used to plan property rights in the ‘liberal legislation’ experiments but this planning subsequently enabled the organic institution of markets (as a form of spontaneous order). We can also say that the division of labour was also enabled by this planning.

### 13.3 Pre-Existing Institutions & Stratified Ontologies

Here we discuss briefly how the simulation results presented in this thesis fit into the discussion of pre-existing institutions in the Introduction. Recall Field’s and Hodgson’s criticisms of how some new institutional and game theoretic models of institutional emergence give rise to an infinite regress problem.

In the models developed for this thesis we accepted that pre-existing institutions might help enable new institutions. The most obvious example of this is language: agents were able to communicate. Therefore, the question is not whether organic institutions can emerge from some institution-free state of nature but whether *new* institutions can emerge within a stratified ontology.

We can note, also, that the change between the first and second model fits neatly with Field’s and Hodgson’s criticism. The first model assumes property rights - this is a pre-existing institution. The second model, which relaxes this assumption, can be viewed as a movement ‘down’ in the analysis, to a prior stratum within a stratified ontology.

Indeed, the results presented show organic institutions (and other structures) emerge across multiple strata<sup>8</sup>. In the default simulations based on the second model, defence of property emerges first and this part-enables the emergence of property rights (due to debilitating fight costs). Once established across the whole population, property rights enable the emergence of an efficient market; and this in turn enables agents to specialise. In this description there are four strata, three of which were enabled by a prior social structure.

Furthermore, the legal rule experiments indicate that these rules can play an enabling role when defence of property does not emerge, or is insufficient to counter those with ‘power’. In a sense, legal rules substituted for a stratum.

---

<sup>8</sup>In terms of the relationship between ‘social structures’ and ‘institutions’, we follow [Hodgson \(2006a\)](#) who states that “Social structures include sets of relations that may not be codified in discourse, such as demographic structures in animal species or in human societies before any understanding of demography.” (p. 3). Also, “an institution is a special type of social structure” (p. 4). This approach seems consistent with the use of these terms in the rest of the Institutional Economics literature.

## 13.4 Research Involving Agent-Based Models

In this section we look briefly at how the research reported in this thesis, which integrated agent-based models, proceeded in a practical sense. This is done in order to assist other researchers in similar situations in the future.

We discussed this to some extent in chapters 1 and 6, noting that the research iterated between modelling and theory. The models were *tools of research* and not used merely to support some preconceived theory. This approach was made particularly necessary by the specific focus on ‘surprising’ emergent phenomena in the research.

The metaphor of exploring some unknown terrain is helpful in what follows.

The first component is to have a clear but approximate destination for the research. Identifying a goal is obviously a common part of research but the important point here is that awareness of the destination should be maintained throughout the process because the iteration between theory and model can lead a researcher to lose sight of that destination over time. Furthermore, the results of the research (while iterating) might lead to an appropriate re-consideration of the destination (or new avenues for future research).

The second component is to accept that the focus on ‘surprising’ outcomes means the process is one of ‘trial and error’. Some research efforts will lead to uninteresting results. Adapting the models, however, can be helpfully informed by related research (the models reported in Chapter 5 are examples) and intuition.

At the outset, it is necessary to create a first ‘best attempt’ model which is related to the research destination and other goals. This is the third component. For the research presented here, a rudimentary economy was created in which agents did some work (foraging) and then had to find each other to transact. Subsequent versions arose from that core model.

If the research is being conducted by one individual, it is very helpful to brainstorm with others (from different disciplines, if appropriate). This is true of most research but it seems even more appropriate when dealing with emergent surprises.

Furthermore, if the research is being conducted by a team, it would be beneficial for at least one researcher writing the code to be also immersed in the theory. This would ensure the theory is correctly represented in the models; and it would help with the interpretation of results. Specialist coders can be helpful: the point is that there is value in at least one researcher working on both so that iteration is internalized to some degree.

The final point to make here is to be on the lookout for unintended consequences. A continuous iteration between modelling and theory might generate at least one. In the simulations developed for this thesis, the identification of an explanation for the ‘paradox’ regarding markets and specialisation, discussed in Section 6.4.4, is an example.

## 13.5 Future Research

There are a number of areas in which the research presented in this thesis could be improved or extended. Below we briefly look at six different themes:

1. Simplification;
2. Bottom-up development of the model;
3. Further development of the existing model;
4. Empirical studies;
5. An expansion of the institutional scope and simplification; and
6. Guiding principles and other work

Let us look at these in turn.

### 13.5.1 Simplification

The most obvious piece of future research would simplify the models presented in this thesis.

As mentioned in Chapter 6, [Edmonds and Moss \(2005\)](#) argue for a ‘KIDS’ approach in computational research; however, they note that a researcher should start with a detailed model and then simplify matters as much as possible while maintaining the underlying results. This thesis did the first of these but not the second. With more research time, the models could have been scaled down to focus on the generalised framework identified in Section 13.1 above.

The identification of this framework helpfully points to the key features of a simplified model, which could be developed in the future.

The absence of a simplified model does not mean the final models and simulation results discussed above have no value without this simplification. The two research questions were answered despite this.

### 13.5.2 Bottom-Up Development

Consistent with the approach taken in the complexity sciences, as well as [Field's \(2007\)](#) reference to a lack of “satisfactory microanalytics” (p. 1), a valuable piece of research would be to draw on the cognitive and neuro-sciences, as well as psychology, in the understanding (and computational modelling) of reasoning under conditions of uncertainty. This could also draw further on research concerned with reinforcement learning and habituation.

Examples for inclusion in the framing and computational modelling of agents' mental models would be particular phenomena noted in social psychology. For instance, [Elsenbrioch and Gilbert \(2014\)](#) discuss how conformity, obedience, and compliance all seem to influence pro-social behaviour.

### 13.5.3 Development of the Existing Model

There are several ways in which the underlying model could be developed, even without any changes arising from the two previous sub-sections. For example, we made a simple assumption that agents were given a home location which could not be changed over their lives. A simple adjustment would be to allow agents to move closer to market locations (say, for a ‘resource fee’). This might have an impact on symmetry breaking, and it would reduce the resource concentration effect if agents are allowed to return home (as in Section [C.10.2](#) of Appendix [C](#)).

We could also develop the approach for how children learn in the models. In the second model we assume that children inherit the mean propensities to steal and defend of their parents (with some variation allowed in certain experiments). We could look at different learning processes, notably those related to social constructivism ([Simon, 1987](#)).

### 13.5.4 Empirical Evidence

A potential criticism of the research presented in this thesis is the lack of empirical testing. After all, the tenth principle of Complexity Economics (CE) listed in Section [2.3.3](#) concerns empirical evidence.

In addition, we saw in Chapter [5](#) how [Brown \(1996\)](#), [Duffy and Ochs \(1999\)](#), and [Duffy \(2001\)](#) all contributed (very constructively) to the monetary emergence literature by conducting empirical research.

This thesis has focused on theoretical and modelling matters because the gaps identified in the literature, discussed in the Introduction, were of this nature.

Nonetheless, empirical research could help improve the realism of the models. Two areas would be of particular interest: (i) the conditions under which organic institutions do

and do not emerge; and (ii) the relationship between legal rules and the immergence of organic institutions.

### 13.5.5 Expansion

The research in this thesis has focused on the emergence / immergence of two different organic institutions: markets and property rights. These arose from different mental models and, as a consequence, the nature of these institutions were also very different.

Future research could look at different ‘institutional spaces’. The key ingredients appear to be:

- a ‘problem’ that agents have to deal with under conditions of uncertainty;
- the creation of a ‘space’ in which agents’ mental models can co-adapt (like a geographic area or behavioural propensities);
- forms of interaction which are economically interesting;
- reinforcement learning from agents’ experiences to changes in their mental models; and
- habituation.

If we were to speculate, it is possible that certain institutional *types* could be found within a wider range of experiments. For example, it is possible the market institutions we observed are one of a class of geographic-related institutions; and perhaps property rights are part of a propensity-related class.

### 13.5.6 Guiding Principles & Other Domains

The final point is to emphasise the use of CE and Agent-Based Models (ABMs) in economic research. The research presented in this thesis can be thought of as an example that makes use of both fields.

We saw in Chapter 2 how CE can be viewed as a more generalised form of economics than Neoclassical theory. Moreover, CE seems to provide a set of guiding principles for research in other domains of economics that: (i) are non-ergodic, (ii) where uncertainty prevails, and (iii) where agents’ mental models seem to exist within a stratified ontology.

Moreover, if we accept the argument in Chapter 2 that human cognition is computational in nature (assuming a broad definition of computation) then ABMs appear to be a more appropriate modelling technology for economics than closed-form mathematics. This is not to suggest mathematical approaches have no value - as mentioned a number of

times in this thesis already, [Arthur \(2013\)](#) argues that these can be useful as first-order approximations of economic activity.

On the value of CE and ABMs in economics, Arthur captures it well when he writes that “in the ocean the interesting things happen not at the equilibrium sea level which is seldom realized, they happen on the surface where everpresent disturbances cause further disturbances. That, after all, is where the boats are.” ([Arthur, 2013](#), p. 12).

# Bibliography

- Aguirre, R. and Nyerges, T. (2014) ‘An Agent-Based Model of Public Participation in Sustainability Management’, *Journal of Artificial Societies and Social Simulation*, 17(1).
- Akerlof, G. A (1970) The market for ‘lemons’: Quality uncertainty and the market mechanism, *Quarterly Journal of Economics* Vol. 84(3), pp. 488-500.
- Al-Suwailem, S. (2014) ‘Complexity and endogenous instability’, *Research in International Business and Finance*, Vol. 30, pp. 393-410.
- Apaloo, J., Brown, J., McNickle, G. G., Vincent, T. L. and Vincent, T. L. (2015), ‘ESS versus Nash: solving evolutionary games’, *Evolutionary Ecology Research*, Vol. 16, pp. 293-314.
- Albino, V., Nunzia, C., and Ilaria, G. (2003) ‘Coordination mechanisms based on cooperation and competition within Industrial Districts: An agent-based computational approach’, *Journal of Artificial Societies and Social Simulation* 6(4).
- Alchian, A. A. (1950) ‘Uncertainty, Evolution, and Economic Theory’, *Journal of Political Economy*, 58(2), pp. 211-222.
- Alchian, A. A. (2023) ‘Property Rights’, *Concise Encyclopedia of Economics*, published online at <http://www.econlib.org/library/Enc/PropertyRights.html> (accessed 17 February 2023).
- Aligica, P. D. (2013) *Institutional Diversity and Political Economy: The Ostroms and Beyond*. Oxford: Oxford University Press.
- Allen, D. W. (2000) ‘Transaction Costs’, in Bouckaert, B. and de Geest, G. (eds.) *Encyclopedia of Law and Economics, Volume 1*. Cheltenham: Edward Elgar Publishing, pp. 893-926.
- Alston, L. J., Eggertsson, T., and North, D. C. (eds.) (1996) *Empirical Studies in Institutional Change*. Cambridge (England): Cambridge University Press.

- Ambrosino, A. (2016) 'Heterogeneity and law: toward a cognitive legal theory', *Journal of Institutional Economics*, Vol. 12(2), pp. 417–442.
- Ambrosino, A., Fontana, M., and Gigante, A. (2018) 'Shifting Boundaries in Economics: The Institutional Cognitive Strand and the Future of Institutional Economics', *Journal of Economic Surveys*, Vol. 32(3), pp. 767–791.
- Anderson, P. W., Arrow, K. J., Pines, D. (eds) (1988) *The Economy as An Evolving Complex System*. Boulder, USA: Westview Press.
- Andreoni, J. (1988) 'Privately provided public goods in a large economy: The limits of altruism', *Journal of Public Economics*, Vol. 35, pp. 57–73.
- Andrighetto, G., Campenni, M., Conte, R., and Paolucci, M. (2007) 'On the Emergence of Norms: a Normative Agent Architecture', in *AAAI Symposium, social and organizational aspects of intelligence*, Washington DC.
- Andrighetto, G., and Conte, R. (2010) 'Theoretical Foundations: The Revised EMIL Ontology', in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 2, pp. 9–13.
- Andrighetto, G., Campenni, M., and Conte, R. (2010) 'Making the Theory Explicit: the EMIL-A Architecture', in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 9, pp. 77–88.
- Andrighetto, G., Campenni, M., Cecconi, F., and Conte, R. (2010) 'The Complex Loop of Norm Emergence: A Simulation Model', in Takadama et al. (eds) *Simulating Interacting Agents and Social Phenomena: The Second World Congress, Agent-Based Social Simulations* 7.
- Aoki, M. (2001) *Toward a Comparative Institutional Analysis*. Cambridge, MA: MIT Press.
- Aoki, M. (2007), 'Endogenizing institutions and institutional changes', *Journal of Institutional Economics*, Vol. 3(1), pp. 131. DOI: 10.1017/S1744137406000531.
- Aoki, M. (2011), 'Institutions as cognitive media between strategic interactions and individual beliefs', *Journal of Economic Behavior & Organization*, Vol. 79(1), pp. 20–34. ISSN: 0167-2681, DOI: <https://doi.org/10.1016/j.jebo.2011.01.025>.
- Archibald, G. C., Simon, H. A, Samuelson, P. A. (1963), *Discussion, The American Economic Review - Papers and Proceedings of the Seventy-Fifth Annual Meeting of the American Economic Association*, Vol. 53, No. 2.



- Arnsperger C., & Varoufakis Y. (2006) 'What Is Neoclassical Economics? The three axioms responsible for its theoretical oeuvre, practical irrelevance and, thus, discursive power', *Panaeconomicus*, 53(1), pp. 5-18.
- Archer, M. (2003) *Structure, Agency and the Internal Conversation*. Cambridge (UK): Cambridge University Press.
- Argyris, C. (1977), 'Double loop learning in organizations', *Harvard Business Review*, Vol. 55(5), pp. 115-125.
- Argyris, C., and Schön, D. (1978), *Organizational Learning: A Theory of Action Perspective - Addison-Wesley Od Series v. 1*. Addison-Wesley Publishing Company. ISBN 9780201001747.
- Aristotle (1901) *Aristotle's Posterior Analytics*. Translated by E. S. Bouchier. Oxford (UK): Blackwell.
- Arrow, K. J. (1951) 'Alternative Approaches to the Theory of Choice in Risk-Taking Situations', *Econometrica*, 19(4), pp. 404-37.
- Arrow, K. J. (1994), 'Methodological Individualism and Social Knowledge', *The American Economic Review*, Vol. 84(2), pp. 1-9.
- Arrow, K. J. and Debreu, G. (1954) 'Existence of an equilibrium for a competitive economy', *Econometrica*, 22(3), pp. 265-290.
- Arthur, W. B. (1988) 'Self-Reinforcing Mechanisms in Economics', in Anderson, P. W., Arrow, and K. J., Pines, D. (eds) *The Economy as An Evolving Complex System*. Boulder, USA: Westview Press.
- Arthur, W. B. (1989) 'Competing Technologies, Increasing Returns, and Lock-In by Historical Events', *The Economic Journal*, Vol. 99(1), pp. 116-131.
- Arthur, W. B. (1994) 'Inductive Reasoning and Bounded Rationality (the El Farol Problem)', *American Economic Review (Papers and Proceedings)*, Vol. 84, pp. 406-411.
- Arthur, W. B. (1994b) *Increasing Returns and Path Dependence in the Economy*. Michigan: University of Michigan Press.
- Arthur, W. B. (1999) 'Complexity and the Economy', *Science*, Vol. 284, pp. 107-109.
- Arthur, W. B. (2009) *The Nature of Technology: What It Is and How It Evolves*. London: Allen Lane.
- Arthur, W. B. (2013) 'Complexity Economics: a Different Framework for Economic Thought', *SFI working paper 2013-04-012*.

- Arthur, W. B. (2021) 'Foundations of complexity economics', *Nature Reviews Physics*, Vol. 3, pp. 136–145.
- Arthur, W. B. (forthcoming) 'Algorithms and the Shift in Modern Science'. In Doria, F. A., Wuppuluri, S. (eds.) *Taming Complexity: Life and Work of Gregory Chaitin*. World Scientific Publishing Company.
- Axelrod, R. (1984) *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, R. (1986) 'An evolutionary approach to norms', *American Political Science Review*, 80(4), pp. 1095-1111.
- Axelrod, R. (1997) *The Complexity of Cooperation*. Princeton: Princeton University Press.
- Axtell, R. (2007) 'What economic agents do: How cognition and interaction lead to emergence and complexity', *The Review of Austrian Economics*, Vol. 20, pp. 105–122.
- Baas, N. A. (1994) 'Emergence, hierarchies, and hyperstructures', in Langton, C. G. (ed.) *Artificial life III*. Reading (USA): Addison-Wesley Publishing Co.
- Barry, N. P. (1982) 'The Tradition of Spontaneous Order', *Literature of Liberty*, Vol. V(2), pp. 7-58.
- Barry, N. P. (1994) 'The Road to Freedom - Hayek's Social and Economic Philosophy', in Birner, J., and van Zijp, R. (Eds.) *Hayek, Co-Ordination and Evolution. His Legacy in Philosophy, Politics, Economics and the History of Ideas*, pp. 141-63. London: Routledge.
- Başçi, E. (1999) 'Learning by imitation', *Journal of Economic Dynamics and Control*, Vol. 23, pp. 1569-1585.
- Bates, R. (2014) 'The New Institutionalism', in Galiani, S. and Sened, I. (eds) *Institutions, Property Rights and Economic Growth: The Legacy of Douglass North*. Cambridge and New York: Cambridge University Press.
- Bateson, G. (2002) *Mind and Nature: A Necessary Unity*. Hampton Press, UK ed. edition.
- Becker, G. S. (1992) 'Habits, addictions and traditions', *Kyklos*, Vol 45, pp. 327-346.
- Bergstrom, T., Blume, L., and Varian, H. (1986) 'On the private provision of public goods', *Journal of Public Economics*, Vol. 29, pp. 25–49.
- Bhavnani, R. (2003) 'Adaptive Agents, Political Institutions and Civic Traditions in Modern Italy', *Journal of Artificial Societies and Social Simulation*, 6(4).

- Bicchieri, C. (1993) *Rationality and Coordination*. Cambridge / New York: Cambridge University Press.
- Bicchieri, C. (2006) *The Grammar of Society*. Cambridge / New York: Cambridge University Press.
- Binmore, K. (1987) 'Modeling Rational Players, Part 1', in *Economics and Philosophy*, Vol. 3, pp. 179-214.
- Blaug, M. (2002) 'Ugly currents in modern economics', in Mäki, U., *Fact and Fiction in Economics*. Cambridge (UK): Cambridge University Press, pp. 35–56.
- Black, J., Hashimzade, N. and Myles, G. (2009), *A Dictionary of Economics*, Oxford University Press. ISBN: 9780191726637, DOI: 10.1093/acref/9780199237043.013.1376.
- Blackburn, J. M. (1936) 'Acquisition of Skill: An Analysis of Learning Curves', *IHRB Report*, No. 73.
- Blumer, H. (1969) 'Symbolic Interactionism: Perspective and Method' Berkley, California: University of California Press.
- Boehm, S. (1994) 'Spontaneous Order', in Hodgson, G., Samuels, W. N., and Tool, M. R. (eds), *The Elgar Companion to Evolutionary and Institutional Economics: L-Z*, Aldershot: Edward Elgar, pp. 296-301.
- Boulton, J. G. (2019) *Complexity Science as Relationship, Patterning and Process*. Available at: <https://www.youtube.com/watch?v=EnDkd0KQ1YU> (Accessed 15 August 2019).
- Bowles, S. (2004), *Microeconomics: Behavior, Institutions, and Evolution*. Princeton University Press. ISBN: 9780691091631.
- Bowles, S. and Gintis, H. (2011) *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton: Princeton University Press. <https://doi.org/10.1515/9781400838837>
- Branch, W. A., and Evans, G. W. (2007) 'Model uncertainty and endogenous volatility', *Review of Economic Dynamics*, Vol. 10(2), pp. 207-237.
- Brickman, P. and Campbell, D. T. (1971) *Hedonic Relativism and Planning in the Good Society*, pp. 287–302. Academic Press, London.
- Bromley, D. W. (2006) *Sufficient Reason: Volitional Pragmatism and the Meaning of Economic Institutions*. Princeton, NJ: Princeton University Press.

- Bronk, R. (2013) 'Reflexivity unpacked: performativity, uncertainty and analytical monocultures', *Journal of Economic Methodology*, 20(4).
- Brown, P. M. (1996) 'Experimental evidence on money as a medium of exchange', *Journal of Economic Dynamics and Control*, Vol. 20, pp. 583-600.
- Buchanan, J. B. (1986) *Liberty, Market and State: Political Economy in the 1980s*. Brighton, England: Wheatsheaf Books.
- Bullock, S. (2014) 'Levins and the lure of artificial worlds', *The Monist*, 97(3), pp. 301-320.
- Buzing, P., Eiben, A., and Schut, M. (2005) 'Emerging Communication and Cooperation in Evolving Agent Societies', *Journal of Artificial Societies and Social Simulation*, 8(1).
- Calcagno, P. and Lopez, E. 'Informal Norms Trump Formal Constraints: The Evolution of Fiscal Policy Institutions in the United States', *Journal of Institutional Economics*, 13(1).
- Caldes, J. and Coelho, H. (1999) 'The Origin of Institutions: socio-economic processes, choice, norms and conventions', *Journal of Artificial Societies and Social Simulation*, 2(2).
- Caldwell, B. (2004) *Hayek's Challenge: An Intellectual Biography of F.A. Hayek*. Chicago: The University of Chicago Press.
- Caldwell, B. (2006) 'Popper and Hayek: Who Influenced Whom?' in Jarvie, I., Milford, K., and Miller, D. (eds) *Karl Popper: A Centenary Assessment*. Oxford: Routledge.
- Calvert, R. (1995) 'Rational Actors, Equilibrium, and Social Institutions', in Knight, J. and Sened, I. *Explaining Social Institutions*. Michigan (USA): University of Michigan Press.
- Campbell, D.T. (1974) '“Downward causation” in hierarchically organized biological systems', in Ayala, F.J., Dobzhansky, T. (Eds.), *Studies in the Philosophy of Biology*. London: Macmillan, pp. 179–186.
- Campenni, M., Henn, S., Jazayeri, P., Lotzmann, U., Möhring, M., Oberhausen, M., Mehmet-Hadi, T., and Weyrich, J. (2010) 'Demonstrating the Theory 4: Multi-scenario World - the Sequel', in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 17, pp. 191-199.
- Castelfranchi, C. (1998) 'Simulating with Cognitive Agents: The Importance of Cognitive Emergence', in Sichman, J. S.; Conte, R.; and Gilbert, N. (Eds) *Lecture Notes in*

- Artificial Intelligence 1534: Proceedings of the First International Workshop, MABS '98.*
- Castellani, E. (2010) 'On the meaning of symmetry breaking', in Brading, K. and Castellani, E. (eds.) *Symmetries in physics: philosophical reflections*. Cambridge (UK): Cambridge University Press, pp. 321-334.
- Castellani, B., and Gerrits, L. (2021) *2021 Map of the Complexity Sciences*, published online at [https://www.art-sciencefactory.com/complexity-map\\_feb09.html](https://www.art-sciencefactory.com/complexity-map_feb09.html) (accessed 30 June 2022).
- Chan, S. (2001) *Complex Adaptive Systems*. Available at <http://web.mit.edu/esd.83/www/notebook/Complex%20Adaptive%20Systems.pdf> (Accessed 2 August 2019).
- Cilliers, P. (1998) *Complexity and Postmodernism: Understanding Complex Systems*. New York: Routledge.
- Coase, R. H. (1937) 'The Nature of the Firm', *Econometrica*, Vol. 4, pp. 386-405.
- Coase, R. H. (1960) 'The Problem of Social Cost', *The Journal of Law and Economics*, Vol. 3, pp. 1-44.
- Colander, D. (2000) (ed.) *The Complexity Vision and the Teaching of Economics*. Cheltenham (UK): Edward Elgar Publishing.
- Colander, D., Follmer, H., Haas, A., Goldberg, M. D., Juselius, K., Kirman, A., Lux, T., and Sloth, B. (2009) 'The Financial Crisis and the Systemic Failure of Academic Economics', *University of Copenhagen Dept. of Economics Discussion Paper No. 09-03*.
- Colander, D., and Kupers, R. (2012) *Laissez-Faire Activism: The Complexity Frame for Policy*. Princeton (USA): Princeton University Press.
- Commons, J.R. (1934) 'Institutional Economics. Its Place in Political Economy'. Reprinted by New Brunswick: Transaction Publishers (1990).
- Conte, R., Andrighetto, G., and Campenni, M. (2013) *Minding Norms - Mechanisms and Dynamics of social order in agent societies* (Oxford Series on cognitive models and architectures). Oxford: Oxford University Press.
- Conte, R., and Castelfranchi, C. (1995a) *Cognitive and Social Action*. London: UCL Press.
- Conte, R., and Castelfranchi, C. (1995b) 'Understanding the functions of norms in social groups through simulation', in Gilbert, N., and Conte, R. (eds.) *Artificial Societies: the computer simulation of social life*. London: UCL Press.

- Conte, R., and Paolucci, M. (2004) 'Responsibility for Societies of Agents', *Journal of Artificial Societies and Social Simulation*, Vol. 7(4).
- Conte, R., and Edmunds, B. (2010) Introduction, in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 1, pp. 5-7.
- Cowan, G. A., Pines, D. and Meltzer, D. (eds.) (1999), *Complexity: Metaphors, Models, and Reality*. USA: Perseus Books. ISBN: 0738202320.
- Crockett, S. (2013) 'Price dynamics in general equilibrium experiments', *Journal of Economic Surveys*, Vol. 27, pp. 421–438.
- Dale, G. (2018) "Our World Was Made by Nature": constructions of spontaneous order', *Globalizations* Vol. 15, pp. 924-40.
- Darwin, C. (1988) *The Descent of Man and Selection in Relation to Sex*. London: Encyclopedia Britannica (1871).
- David, P. A. (2007) 'Path Dependence, its Critics, and the Quest for 'Historical Economics'', in Hodgson, G. (ed) *The Evolution of Economic Institutions*, pp. 120-142. Cheltenham: Edward Elgar Publishing.
- Davidson, P. A. (1991) 'Is Probability Theory Relevant for Uncertainty? A Post-Keynesian Perspective', *Journal of Economic Perspectives*, 5(1), pp. 129-43.
- Davidson, P. A. (1996) 'Reality and Economic Theory', *Journal of Post Keynesian Economics*, 18(4), pp. 479-508.
- Denning, P. (2010) 'Ubiquity Symposium "What is Computation?" Opening Statement', Ubiquity, November 2010.
- Denzau, A., and North, D. (1994) 'Shared Mental Models: Ideologies and Institutions', *Kyklos*, Vol. 47.
- Dewey, J. (1925), *Experience and Nature*. Chicago: Open Court Publishing Company.
- Dewey J. (1938) *Logic: The Theory of Inquiry*. New York: Henry Holt and Company.
- Di Paolo, E. (2000) 'Ecological Symmetry Breaking can Favour the Evolution of Altruism in an Action-response Game', *Journal of Theoretical Biology*, Vol. 203, pp. 135-152.
- dos Anjos, P. L., Lotzmann, U., and Pauli, M. (2010) 'Demonstrating the Theory 3: Micro Finance', in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 16, pp. 169-189.
- Dosi, G., Marengo, L., and Fagiolo, G (2003) 'Learning in Evolutionary Environments', *LEM Papers Series 2003/20*, Sant'Anna School of Advanced Studies, Pisa, Italy.

- Duffy, J. (2001) 'Learning to speculate: Experiments with artificial and real agents', *Journal of Economic Dynamics and Control*, Vol. 25, pp. 295-319.
- Duffy, J. and Ochs, J. (1999) 'Emergence of money as a medium of exchange: an experimental study', *American Economic Review*, Vol. 89, pp. 847-877.
- Dugger, W. (1990) 'The New Institutionalism: New But Not Institutional', *Journal of Economic Issues*, XXIV(2), pp. 423-431.
- Easterlin, R. A. (1974) 'Does economic growth improve the human lot? Some empirical evidence' in David, P. A. and Melvin, W. R. (eds.) *Nations and Households in Economic Growth*, p. 98-125. Stanford University Press, Palo Alto, USA.
- Edmonds, B. (2004) 'How formal logic can fail to be useful for modelling or designing MAS', in *Regulated agent-based social systems* (Lecture notes in artificial intelligence, Vol. 2934, pp. 1-15). Berlin: Springer.
- Edmonds, B., and Moss, S. (2005) 'From KISS to KIDS - an 'anti-simplistic' modelling approach', in Davidsson et al. (Eds), *Multi agent based simulation 2004* (Lecture notes in artificial intelligence, Vol. 3415, pp. 130-144). Berlin: Springer.
- Eichner, A. S. and Kregel, J. A. (1975) 'An Essay on Post-Keynesian Theory: A New Paradigm in Economics', *Journal of Economic Literature*, Vol 13(4), pp. 1293-1314.
- Elsenbrioch, C., and Gilbert, N. (2014) *Modelling Norms*. Springer.
- Elsner, W. (2012) 'The Theory of Institutional Change Revisited: The Institutional Dichotomy, Its Dynamic, and Its Policy Implications in a More Formal Analysis', *Journal of Economic Issues*, 46(1), pp. 1-43.
- Elster, J. (1989) *The Cement of Society*. Cambridge: Cambridge University Press.
- Epstein, J. (2012) *Generative Social Science: Studies in Agent-Based Computational Modelling* (Princeton Studies in Complexity). Princeton: Princeton University Press.
- Epstein, J., and Axtell, R. (1996) *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press.
- Erev, I., and Roth, A. E. (1998) 'Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria', *American Economic Review*, Vol. 88, pp. 848-881.
- Estes, W. K. (1950) 'Toward a Statistical Theory of Learning', *Psychological Review*, Vol. 57(2), pp. 94-107.
- Farmer, J. D. (2012) 'Economics Needs to Treat the Economy as a Complex System', *Institute for New Economic Thinking*, 1(4).

- Faundez, J. (2016) 'Douglass North's Theory of Institutions: Lessons for Law and Development', *Hague Journal on the Rule of Law*, Vol. 8(2), pp. 373-419.
- Feldman, M., and Pentland, B. (2003) 'Reconceptualizing organizational routines as a source of flexibility and change', *Administrative Science Quarterly*, 48 (1), pp. 94-118.
- Ferguson, W. (2013) *Collective Action and Exchange: A Game-Theoretic Approach to Contemporary Political Economy*. Stanford University Press.
- Ferguson, A. (1767) *An Essay on the History of Civil Society*. London: T. Cadell, and Edinburgh: W. Creech and J. Bell.
- Ferrari-Filho, F. and Conceição, O. A. C (2005) 'The Concept of Uncertainty in Post Keynesian Theory and in Institutional Economics', *Journal of Economic Issues*, XXXIX(3).
- Field, A. J. (1979) 'On the Explanation of Rules Using Rational Choice Models', *Journal of Economic Issues*, 13(1), pp. 49-72.
- Field, A. J. (1981) 'The Problem with Neoclassical Institutional Economics: A Critique with Special Reference to the North/Thomas model of Pre-1500 Europe', *Explorations in Economic History*, Vol. 18, pp. 174-198.
- Field, A. J. (1984) 'Microeconomics, Norms and Rationality', *Economic Development and Cultural Change*, 32(4), pp. 683-711.
- Field, A. J. (2007) 'Beyond Foraging', *Journal of Institutional Economics*, 3(3).
- Focardi, S. (2015) 'Is economics an empirical science? If not, can it become one?', *Frontiers in Applied Mathematics and Statistics*, Vol. 1, pp. 1-15. Available at <https://www.frontiersin.org/articles/10.3389/fams.2015.00007/full> (Accessed 20 August 2019).
- Forno, A. and Melone, U. (2002) 'A multi-agent simulation platform for modelling perfectly rational and bounded-rational agents in organizations', *Journal of Artificial Societies and Social Simulation*, 2(2).
- Forster, J. (1860) *Oliver Cromwell; Daniel Defoe, Sir Richard Steele, Charles Churchill, Samuel Foote; biographical essays*. London: J. Murray.
- Foxon, T. J. (2002) *Technological and institutional 'lock-in' as a barrier to sustainable innovation*, ICCEPT Working Paper, November 2002. Available at <http://www.iccept.ic.ac.uk> (Accessed 10 October 2019).
- Friedman, M. (1953) *Essays in Positive Economics*. Chicago: University of Chicago Press.



- Gallegati, M., and Kirman, A. (2012) 'Reconstructing Economics: Agent Based Models and Complexity', *Complexity Economics*, Vol. 1, pp. 5-31.
- Gao, D., Deng, X., Zhao, Q., Zhou, H., and Bai, B. (2015) 'Multi-Agent Based Simulation of Organizational Routines on Complex Networks', *Journal of Artificial Societies and Social Simulation*, 13(3).
- Gedeon, P. (2015) 'Spontaneous Order and Social Norms Hayek's Theory of Socio-Cultural Evolution', *Society and Economy*, Vol. 37(1), pp. 1-29.
- Gell-Mann, M., and Hartle, J. B. (2007) 'Quasiclassical coarse graining and thermodynamic entropy', *American Physical Society*, 76(2).
- Gilbert, N. (2002) 'Varieties of Emergence', Paper presented at the Social Agents: Ecology, Exchange, and Evolution Conference Chicago.
- Gomes, O. (2014) 'On the Concept of Endogenous Volatility', in Faggini, M., and Parziale, A. (eds) *Complexity in Economics: Cutting Edge Research. New Economic Windows*. Springer, Cham.
- Gräbner, C. (2016) 'Agent-based computational models - a formal heuristic for institutional pattern modelling?', *Journal of Institutional Economics*, Vol. 12(1), pp. 241-261.
- Gräbner, C. (2018) 'Formal Approaches to Socio-economic Analysis - Past and Perspectives', *Forum for Social Economics*, 47(1), pp. 32-63.
- Gräbner, C. and Ghorbani, A. (2019) *Defining institutions - A review and a synthesis*. ICAE Working Paper Series No. 89 - February 2019. Available at: [https://www.jku.at/fileadmin/gruppen/108/ICAE\\_Working\\_Papers/wp89.pdf](https://www.jku.at/fileadmin/gruppen/108/ICAE_Working_Papers/wp89.pdf) (Accessed 20 January 2020).
- Gräbner, C., and Kapeller, J. (2015) 'New Perspectives on Institutional Pattern Modeling: Systemitism, Complexity, and Agent-Based Modeling', *Journal of Economic Issues*, Vol. 49(2), pp. 433-440.
- Gray, J. (1988) 'Hayek, the Scottish School, and Contemporary Economics', in Winston, G. C., and Teichgraeber III, R. F. (eds) *The Boundaries of Economics*. Cambridge: Cambridge University Press
- Gray, J. (1998) *Hayek on Liberty*, Third Edition. London: Routledge.
- Greif, A. (1989), 'Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders', *The Journal of Economic History*, Vol. 49(4), pp. 857-882. DOI: 10.1017/S0022050700009475.

- Greif, A. (1994), 'Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies', *Journal of Political Economy*, Vol. 102(5), pp. 912-950. DOI: 10.1086/261959, eprint: <https://doi.org/10.1086/261959>.
- Greif, A. (1998) 'Self-Enforcing Political System and Economic Growth: Late Medieval Genoa', in Bates, B., Greif, A., Margaret, L., Rosenthal, J., and Weingast, B. *Analytic Narrative*, pp. 23-63. Princeton University Press.
- Greif, A. (2006) *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade (Political Economy of Institutions and Decisions)*. Cambridge University Press, UK.
- Greif, A., Milgrom, P., and Weingast, B. R. (1994), 'Coordination, Commitment, and Enforcement: The Case of the Merchant Guild', *Journal of Political Economy*, Vol. 102(4), pp. 745-776. ISSN: 00223808,1537534X.
- Gupta, S., Davoodi, H., and Alonso-Terme, R. (1998) *Does Corruption Affect Income Inequality and Poverty?* IMF Working Paper WP/98/76.
- Haken, H. (1987) 'Synergetics: An approach to self organization', in Yates, F. E. (ed.) *Self-organizing systems: The emergence of order*. Berlin: Plenum Press.
- Hamilton, W. H. (1919) 'The Institutional Approach to Economic Theory', *American Economic Review*, Vol. 9, Supplement, pp. 309-318.
- Hamilton, W. H. (1932) 'Institution', in Seligman, R. A. and Johnson, A. (eds) *Encyclopedia of the Social Sciences*, Vol. 8, pp. 84-89. New York: Macmillan.
- Hamowy, R. (1987) *The Scottish Enlightenment and the Theory of Spontaneous Order*. Carbondale : Southern Illinois University Press.
- Hamowy, R. (2003) 'F.A. Hayek and the Common Law', *Cato Journal* Vol. 32, pp. 241-64.
- Hardin, G. (1968) 'The Tragedy of the Commons', *Science*, Vol. 162, pp. 1243-1248.
- Hayek, F. A. (1933) *Monetary Theory and the Trade Cycle*, Translated by N. Kaldor and H. M. Croome. New York: Sentry Press.
- Hayek, F. A. (1937) 'Economics and Knowledge', *Economica*, Vol. 4, pp. 33-54.
- Hayek, F. A. (1945) 'The Use of Knowledge in Society' in *The American Economic Review*, Vol. 35(4), pp. 519-530.
- Hayek, F. A. (1958) 'Freedom, Reason, and Tradition' in *Ethics - An International Journal of Social, Political, and Legal Philosophy*, Vol. LXVIII, pp. 229-245.

- Hayek, F. A. (1960) Letter (27 February) to Karl Popper. Hoover Institute, Hayek Archives, Box 44, Folder 2.
- Hayek, F. A. (1967) *Studies in philosophy, politics, and economics*. Chicago: University of Chicago Press.
- Hayek, F. A. (1973) *Law, legislation and liberty, Vol. I: Rules and Order*. London: Routledge & Kegan Paul.
- Hayek, F. A. (1976a) *The Denationalization of Money*. London: Institute of Economic Affairs.
- Hayek, F. A. (1976b) *Individualism and Economic Order*. London: Routledge & Kegan Paul.
- Hayek, F. A. (1976c) *Law, Legislation, and Liberty, Vol. II: The Mirage of Social Justice*. London: Routledge.
- Hayek, F. A. (1976d) *The Road to Serfdom*. Chicago: University of Chicago Press.
- Hayek, F. A. (1978) *New Studies in Philosophy, Politics, Economics and the History of Ideas*. Chicago: The University of Chicago Press.
- Hayek, F. A. (1979) *Law, legislation and liberty, Vol. III: The Political Order of Free People*. London: Routledge & Kegan Paul.
- Hayek, F. A., and Bartley, W. W. (ed) (1988) *The Fatal Conceit: The Errors of Socialism*. Chicago: The University of Chicago Press. [Note: the authenticity of this as a book written solely by F. A. Hayek is disputed by some researchers. Hayek fell ill in 1985 and it is therefore “not clear how much of the book should be attributed to Hayek and how much to Bartley.” (Caldwell, 2004, p. 317). Bartley was purported to be the book’s editor and not a co-author.]
- Hayek, F. A. (2006) *The Constitution of Liberty*. Oxford: Routledge.
- Hayek, F. A. (2014) ‘The Theory of Complex Phenomena’ in Caldwell, B. (ed) *The Market and Other Orders*. Oxford: Routledge.
- Hayek, F. A. (2018) ‘The Sensory Order’ in Vanberg, V. (Ed.) *The Collected Works of F. A. Hayek*, Vol. 14. Oxford: Routledge.
- Hédoin, C. (2012) ‘Linking institutions to economic performance: the role of macrostructures in micro-explanations’, *Journal of Institutional Economics*, 8(3), pp. 327-349.
- Heiner, R. (1983) ‘The Origin of Predictable Behavior’, *American Economic Review*, 73(4), pp. 560-595.

- Heyward, A. (2015) *Global Politics*. Basingstoke: Palgrave Macmillan.
- Hindriks, F. and Guala, F. (2015), 'Institutions, rules, and equilibria: a unified theory', *Journal of Institutional Economics*, Vol. 11(3), pp. 459-480. DOI: 10.1017/S1744137414000496.
- Hobbes, T. [1651] (2017) *Leviathan*. Penguin Classics.
- Hodgson, G. M. (1988) *Economics and Institutions: A Manifesto for a Modern Institutional Economics*. Cambridge and Philadelphia: Polity Press and University of Pennsylvania Press.
- Hodgson, G. M. (1991) 'Hayek's Theory of Cultural Evolution', *Economics and Philosophy*, Vol. 7, pp. 67-82.
- Hodgson, G. M. (1992) 'Carl Menger's theory of the evolution of money: some problems', *Review of Political Economy*, Vol. 4(4), pp. 396-412.
- Hodgson, G. M. (1994) 'The State, Money, and "Spontaneous Order"', *Critical Review*, Fall 1994, pp. 579-589.
- Hodgson, G. M. (1998) 'The Approach of Institutional Economics', *Journal of Economic Literature*, 36(1), pp. 166-192.
- Hodgson, G. M. (2000a) 'What is the Essence of Institutional Economics', *Journal of Economic Issues*, XXXIV(2), pp. 317-329.
- Hodgson, G. M. (2000b) *Structures and Institutions: Reflections on Institutionalism, Structuration Theory and Critical Realism*. Unpublished mimeo, University of Hertfordshire.
- Hodgson, G. M. (2000c) 'The Concept of Emergence in Social Science: Its History and Importance', *Emergence: Complexity and Organization*, 2(4), pp. 65-77.
- Hodgson G. M. (2001) 'Is social evolution Lamarckian or Darwinian?', in: Laurent, J., Nightingale, J. (eds) *Darwinism and evolutionary economics*, pp. 87-118. Cheltenham: Edward Elgar.
- Hodgson, G. M. (2002a) 'The Evolution of Institutions: An Agenda for Future Theoretical Research', *Constitutional Political Economy*, 13(2), pp. 111-127.
- Hodgson, G. M. (2002b) 'Darwinism in economics: from analogy to ontology', *Journal of Evolutionary Economics*, Vol. 12, pp. 259-281.
- Hodgson, G. M. (2002c) 'Reconstitutive Downward Causation: Social Structure and the Development of Individual Agency' in Fullbrook, E. (ed.) (2002) *Intersubjectivity in Economics: Agents and Structures*, pp. 159-180. London and New York: Routledge.

- Hodgson, G. M. (2003a) 'Darwinism and Institutional Economics', *Journal of Economic Issues*, XXXVII(1), pp. 85-97.
- Hodgson, G. M. (2003b) 'John R. Commons and the Foundations of Institutional Economics', *Journal of Economic Issues*, Vol. XXXVII(3), pp. 547-576.
- Hodgson, G. M. (2003c) 'The hidden persuaders: institutions and individuals in economic theory', *Cambridge Journal of Economics*, Vol. 27, pp. 159-175.
- Hodgson, G. M. (2004a) *The Evolution of Institutional Economics: Agency, structure and Darwinism in American Institutionalism*. New York: Routledge.
- Hodgson, G. M. (2004b) 'Reclaiming habit for institutional economics', *Journal of Economic Psychology*, Vol. 25, pp. 651-660.
- Hodgson, G. M. (2006a) 'What Are Institutions?', *Journal of Economic Issues*, XL(1), pp. 1-25.
- Hodgson, G. M. (2006b) 'Instinct and Habit Before Reason: Comparing the Views of John Dewey, Friedrich Hayek and Thorstein Veblen' in Krecké, E., Krecké, C. and Koppl, R.G. (Eds.) *Cognition and Economics (Advances in Austrian Economics, Vol. 9)*, pp. 109-143. Bingley: Emerald Group Publishing Limited.
- Hodgson, G. M. (2006c) *Economics in the Shadows of Darwin and Marx: Essays on Institutional and Evolutionary Themes*. Cheltenham: Edward Elgar.
- Hodgson, G. M. (2007) 'Meanings of methodological individualism', *Journal of Economic Methodology*, 14(2), pp. 211-226.
- Hodgson, G. M. (2009) 'Institutional Economics into the Twenty-First Century', *Studi e Note di Economia*, Anno XIV, n. 1-2009, pp. 3-26.
- Hodgson, G. M. (2011) *Downward Causation - Some Second Thoughts*. Available at: <https://www.geoffreymhodgson.uk/downward-causation> (Accessed 13 September 2022).
- Hodgson, G. M. (2012) 'On the Limits of Rational Choice Theory', *Economic Thought*, 1(1), pp. 94-108.
- Hodgson, G. M. (2015) 'On Defining Institutions: Rules versus Equilibria', *Journal of Institutional Economics*, Vol. 11, pp. 499-505.
- Hodgson, G. M. (2017) 'Introduction to the Douglass C. North memorial issue', in *Journal of Institutional Economics*, 13(1), pp. 1-23.
- Hodgson, G. M. (2019) 'Taxonomic definitions in social science, with firms, markets and institutions as case studies', in *Journal of Institutional Economics*, 15(2), pp. 1-18.

- Hodgson, G. M., and Knudsen, T. (2004) 'The complex evolution of a simple traffic convention: the functions and implications of habit', *Journal of Economic Behavior & Organization*, Vol. 54, pp. 19–47.
- Hodgson, G. M., and Knudsen, T. (2006) 'Why we need a generalized Darwinism, and why generalized Darwinism is not enough', *Journal of Economic Behavior & Organization*, Vol. 61, pp. 1-19.
- Hodgson, G. M., and Huang, K. (2012) 'Evolutionary game theory and evolutionary economics: are they different species?', *Journal of Evolutionary Economics*, Vol. 22, pp. 345-366.
- Holland, J. (1975) *Adaptation in natural and artificial systems*. Ann Arbor (USA): University of Michigan Press.
- Holland, J. (2014) *Signals and Boundaries: Building Blocks for Complex Adaptive Systems*. Cambridge (USA): MIT Press.
- Holland, J., Holyoak, K., Nisbett, R., and Thagard, P. (1986) *Induction: Processes of Inference, Learning, and Discovery*. Cambridge (USA): MIT Press.
- Hollander, C. D., and Wu, A. S. (2011) 'The Current State of Normative Agent-Based Systems', *Journal of Artificial Societies and Social Simulation*, 14(2).
- Horgan, J. (1995) *From Complexity to Perplexity: Can science achieve a unified theory of complex systems? Even at the Santa Fe Institute, some researchers have their doubts*. Available at: <http://www2.econ.iastate.edu/tesfatsi/hogan.complexperplex.htm> (Accessed: 25 July 2019).
- Horgan, J. (1997) *The End of Science: Facing the Limits of Knowledge in the Twilight of the Science Age*. New York: Broadway Books.
- Howitt, P. and Clower, R. (2000) 'The emergence of economic organization', *Journal of Economic Behavior and Organization*, 41(1), pp. 55–84.
- Hurwicz, L. (1993), 'Toward a framework for analyzing institutions and institutional change', in Bowles, S., Gintis, H., and Gustafsson, B. (eds) *Markets and Democracy: Participation, Accountability and Efficiency*, pp. 51-67. Cambridge University Press. DOI: 10.1017/CBO9780511983580.005.
- Hurwicz, L. (1996), 'Institutions As Families Of Game Forms', *The Japanese Economic Review*, Vol. 47, pp. 113-132.
- Isaac, A. (2008) 'Simulating Evolutionary Games: A Python-Based Introduction', *Journal of Artificial Societies and Social Simulation*, 11(3).

- Jaffe, K. (2002) 'An economic analysis of altruism: who benefits from altruistic acts?', *Journal of Artificial Societies and Social Simulation*, 5(3).
- James, W. (1890) *The Principles of Psychology*, 2 vols, 1st edn. New York and London: Holt and Macmillan.
- Janzen, D. H. (1980) 'When is it co-evolution?', *Evolution*, 34(3), pp. 611-612.
- Jensen, M. C. and Meckling, W. H. (1979) 'Rights and Production Functions: An Application to Labor-Managed Firms and Codetermination', *Journal of Business*, 52(4), p. 469-506.
- Jevons, W. (1875) *Money and the mechanism of exchange*. New York: D. Appleton and Co.
- Johnson, S. (2001) *Emergence: The connected lives of ants, brains, cities and software*. New York: Scribner.
- Johnson, W. (2012) 'Throw Your Life a Curve', *Harvard Business Review*, September 2012.
- Jones, R. (1976) 'The origin and development of media of exchange', *Journal of Political Economy*, Vol. 84(4), pp. 757-775.
- Judd, K. (2006) 'Computationally Intensive Analyses in Economics', in Tesfatsion, L. and Judd, K. (eds.) *Handbook of Computational Economics: Vol. 2 Agent-based Computational Economics*. New York (USA): North-Holland Elsevier, pp. 881-893.
- Kahneman, D. (2012) *Thinking, Fast and Slow*. Penguin Books.
- Katouzian, H. (1980) 'Positive Economics: The Logic of Neo-Classical Economic Theory', in *Ideology and Method in Economics - Macmillan New Studies in Economics*. London: Palgrave.
- Kauffman, S. (1993) *Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.
- Kauffman, S. A. (2008) *Reinventing the Sacred: A New View of Science, Reason, and Religion*. New York (USA): Basic Books.
- Keita, L. (1983), 'The Instrumentalism-Realism Debate: A Case for a Diachronic Phase Theory', *Crítica: Revista Hispanoamericana de Filosofía*, Vol. 15(43), pp. 79-102.
- Kelsen, H. [1949] (2006) *General Theory of Law & State*. New Brunswick (USA): Transaction Publishers.
- Keynes, J. M. (1921) *A Treatise on Probability*. London: Macmillan & Co.

- Keynes, J. M. (1936) *The General Theory of Employment, Interest and Money*. London: Palgrave Macmillan.
- Keynes, J. M. (1937) 'The General Theory of Employment', *The Quarterly Journal of Economics*, 51(2), pp. 209-223.
- Kidland, F. E. and Prescott, E. C (1977) 'Rules Rather Than Discretion: The Inconsistency of Optimal Plans', *Journal of Political Economy*, 85(3), pp. 473-492.
- Kirman, A. (1993) 'Ants, Rationality, and Recruitment', *The Quarterly Journal of Economics*, Vol. 108(1), pp. 137-156.
- Kirman, A. (2011) *Complexity Economics*. New York: Routledge.
- Kirman, A. P., and Vriend, N. J. (2000) 'Learning to Be Loyal. A Study of the Marseille Fish Market.', in Gatti, D. D., Gallegati, M., and Kirman, A. (eds) *Interaction and Market Structure. Lecture Notes in Economics and Mathematical Systems*, Vol. 484. Springer, Berlin, Heidelberg.
- Kiyotaki, N., and Wright, R. (1989) 'On money as a medium of exchange', *Journal of Political Economy*, Vol. 97, pp. 927-954.
- Klein, D. B. (2013) 'Kenneth J. Arrow', *Econ Journal Watch: Scholarly Comments on Academic Economics*, Vol. 10(3).
- Klein, D. B. and Daza, R. (2013) 'Franco Modigliani', *Econ Journal Watch: Scholarly Comments on Academic Economics*, Vol. 10(3).
- Kley, R. (1994) *Hayek's Social and Political Thought*. Oxford: Clarendon Press.
- Knight, F. H. (1921) *Risk, Uncertainty, and Profit*. New York: Houghton Mifflin and Co.
- Knight, J. (1992) *Institutions and Social Conflict (Political Economy of Institutions and Decisions)*. Cambridge (UK): Cambridge University Press.
- Knudsen, T. (2001) 'Nesting Lamarckism within Darwinian explanations: necessity in economics and possibility in biology?', in Laurent J., Nightingale, J. (eds) *Darwinism and evolutionary economics*, pp. 121-159. Cheltenham: Edward Elgar.
- Kuhn, T. (1962) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kukathas, C. (1989) *Hayek and Modern Liberalism*. Oxford: Clarendon Press.
- Ladyman, J., Lambert, J., and Wiesner, K. (2013) 'What is a complex system?', *European Journal for Philosophy of Science*, 3(1), pp. 33-67.



- Lagueux, M. (2010) *Rationality and Explanation in Economics*. New York: Routledge.
- Lange, O. and Taylor, F. M. (1938) *On the Economic Theory of Socialism*, New edition, Vol. 2, edited by Lippincott, B. E. University of Minnesota Press.
- Laplace, P. S. (1814) *Essai philosophique sur les probabilités*. Paris: Courcier.
- Lavoie, M. (2008) 'Neoclassical empirical evidence on employment and production laws as artefact', *2nd Seminario de microeconomía heterodoxa Facultad de Economía, UNAM, Mexico, 10-12 October 2007*.
- Liebowitz, S.J. and Margolis, S.E. (1990), 'The Fable of the Keys', *The Journal of Law and Economics*, Vol. 33(1), pp. 1-25.
- Leibowitz, N., Baum, B., Enden, G., and Karniel, A. (2010) 'The exponential learning equation as a function of successful trials results in sigmoid performance'. *Journal of Mathematical Psychology*, Vol. 54, pp. 338–340.
- Lewis, D. (2008), *Convention: A Philosophical Study*. Wiley [1969]. ISBN: 9780470692967.
- Lloyd, P., and Lee, C. (2018) 'A Review of the Recent Literature on the Institutional Economic Analysis of the Long-Run Performance of Nations', *Journal of Economic Surveys*, Vol. 32 (1), pp. 1–22.
- Loasby, B. J. (2000) 'Market Institutions and Economic Evolution', *Journal of Evolutionary Economics*, Vol. 10, pp. 297-309.
- Lorenz, E. (1963) 'Deterministic Nonperiodic Flow', *Journal of Atmospheric Studies*, Vol. 20, pp. 130-141.
- Lotzmann, U. (2010) 'Demonstrating the Theory 1: Traffic', in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 14, pp. 139-147.
- Lotzmann, U., and Möhring, M. (2010) 'Supporting the Theory: Formal Description of EMIL-S', in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 11, pp. 91-108.
- Lotzmann, U., Emde, R., and Troitzsch, K. G. (2010) 'Demonstrating the Theory 2: Wikipedia', in *EMergence In the Loop: Simulating the two way dynamics of norm innovation*, chap. 15, pp. 149-168.
- Luban, D. (2020) What is Spontaneous Order?, *American Political Science Review*, Vol. 114(1), pp. 68-80.
- Lucas, R. (1981) 'Tobin and Monetarism: A Review Article', *Journal of Economic Literature*, 19(2), pp. 558-567.

- Luhmann, N. (1995) *Social Systems*. Stanford: Stanford University Press.
- Luckmann, P. L., and Berger, T. *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*, Penguin Books.
- Machlup, N. (1967) 'Theories of the Firm: Marginalist, Behavioural, Managerial', *American Economic Review*, 57(1), pp. 1-33.
- Macy, M. (1998) 'Social Order in Artificial Worlds', *Journal of Artificial Societies and Social Simulation*, 1(1).
- Mayr, E. (1985) 'How biology differs from the physical sciences', in Depew, D.J., and Weber, B.H. (Eds.), *Evolution at a Crossroads: The New Biology and the New Philosophy of Science*. Cambridge, MA: MIT Press, pp. 43-63.
- Marciano, A. (2009) 'Why Hayek is a Darwinian (after all)? Hayek and Darwin on social evolution', *Journal of Economic Behavior & Organization*, Vol. 71, pp. 52-61.
- Marimon, R., McGrattan, E., and Sargent, T. J. (1990) 'Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents', *Journal of Economic Dynamics and Control*, Vol. 14, pp. 329-373.
- Marimon, R., and Miller, J. H. (1989), "Money as a Medium of Exchange in an Economy with Genetically Reproduced Decision Rules", Mimeo
- Maturana, H. R., and Varela, F. J. (1980) *Autopoiesis and Cognition : The Realization of the Living (Boston Studies in the Philosophy and History of Science)*. Dordrecht, Holland: D. Reidel Publishing Company.
- Menard, C., and Shirley, M. M. (2005) 'Introduction', in Menard, C. and Shirley, M. M. (eds.) *Handbook of New Institutional Economics*. Dordrecht, The Netherlands: Springer.
- Menger, C. (1892) 'On the origins of money', *The Economic Journal* Vol. 2, pp. 239-55.
- Menger, C. (1981) *Principles of Economics*, translated by J. Dingwall, and B. F. Hoselitz. New York: New York University Press, 1981.
- Menger, C. (1985) *Investigations into the Methods of the Social Sciences with Special Reference to Economics* (published in 1963 as *Problems of Economics and Sociology*), edited by Louis Schneider and translated by Francis J. Nock. New York: New York University Press.
- Meramveliotakis, G. and Milonakis, D. (2010) 'Surveying the Transaction Cost Foundations of New Institutional Economics: A Critical Inquiry', *Journal of Economic Issues*, 44(4), pp. 1045-1072.

- Milgrom, P. R., North, D. C., and Weingast, B. R. (1990), 'The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs', *Economics and Politics*, Vol. 2(1), pp. 1-23. DOI: <https://doi.org/10.1111/j.1468-0343.1990.tb00020.x>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0343.1990.tb00020.x>.
- Miller, M. S. (1996) 'Learning Curve: Review of John Holland, Hidden Order', *Reason Magazine*, December 1996, pp. 59-64.
- Miller, S. (2004) *Social Action. A Teleological Account*. Cambridge (UK): Cambridge University Press.
- Miller, J. H. and Page, S. E. (2007), *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press. ISBN: 9780691127026.
- Mirowski, P. (1990) 'Learning the meaning of the dollar: conservation principles and the social theory of value in economic theory', *Social Research*, Vol. 57, pp. 689-717.
- Mitchell, M. (2011) *Complexity: A Guided Tour*. Oxford University Press.
- Mitleton-Kelly, E. (2003) 'Ten Principles of Complexity & Enabling Infrastructures' in Mitleton-Kelly, E. (ed.) *Complex Systems and Evolutionary Perspectives on Organisations: The Application of Complexity Theory to Organisations*. Bingley (UK): Emerald Group Publishing Limited, pp. 23-50.
- Morgan, C. L. (1925) *Emergent Evolution*, 2nd edition. London: Williams and Norgate.
- Morgan, C. L. (1932) 'C. Lloyd Morgan' in C. Murchison (ed.) *A History of Psychology in Autobiography*, Vol. 2. New York: Russell and Russell, pp. 237-264.
- Morowitz, H. J. (1998) 'Emergence and equilibrium', *Complexity*, 4(6), pp. 12-13.
- Morowitz, H. J. (2002) *The emergence of everything: How the world became complex*. New York, USA: Oxford University Press.
- Nelson, R. H. (2001) *Economics as Religion: From Samuelson to Chicago and Beyond*. University Park (USA): Penn State University Press.
- Neumann, M. (2008) 'Homo Socionicus: a Case Study of Simulation Models of Norms', *Journal of Artificial Societies and Social Simulation*, 11(4).
- Neumann, M. (2010) 'Norm Internalisation in Human and Artificial Intelligence', *Journal of Artificial Societies and Social Simulation*, 13(1).

- New York Times (2008) *Edward N. Lorenz, a Meteorologist and a Father of Chaos Theory, Dies at 90*. Available at <https://www.nytimes.com/2008/04/17/us/17lorenz.html> (Accessed: 7 August 2019).
- Nicolis G., and Prigogine I. (1989) *Exploring complexity*. New York: W.H.Freeman & Co Ltd.
- Nisan, N., Roughgarden, T., Tardos, E., and Vijay, V. V. (2007) (eds): *Algorithmic Game Theory*. Cambridge University Press.
- North, D. C. (1981) *Structure and Change in Economic History*. New York: Norton.
- North, D. C. (1990) *Institutions, Institutional Change and Economic Performance*. Cambridge (UK): Cambridge University Press.
- North, D. C. (1992) 'Institutions, Ideology, and Economic Performance', *Cato Journal*, 11(3).
- North, D. C. (1993) *The New Institutional Economics and Development*. Available at <http://www2.econ.iastate.edu/tesfatsi/NewInstE.North.pdf> (Accessed 11 October 2019).
- North, D. C. (1994) 'Economic Performance through Time', *American Economic Review*, 84(4), pp. 359-367.
- North, D. C. (2005) *Understanding the Process of Economic Change*. Princeton: Princeton University Press.
- North, D. C. and Thomas, R. P. (1973) *The Rise of the Western World*. Cambridge: Cambridge University Press.
- North, D. C., Wallis, J. J., and Weingast, B. R. (2009) *Violence and Social Orders*. Cambridge: Cambridge University Press.
- Nowak, M., and May, R. (1992) 'Evolutionary Games and Spatial Chaos', *Nature*, Vol 359, pp. 826-829.
- Nozick, R. (1974) *Anarchy, State, and Utopia*, Oxford: Blackwell (Basic Books, New York).
- Okada, A. (2008) 'The Second-Order Dilemma of Public Goods and Capital Accumulation', *Public Choice*, Vol. 135(3/4), pp. 165-182.
- Okazaki, T., Okuno-Fujiwara, M. and Greif, A. (1998), 'Evolution of Economic Systems: The Case of Japan', in Y. Hayami and M. Aoki (eds) *The Institutional Foundations of East Asian Economic Development: Proceedings of the IEA Conference held in*

- Tokyo, Japan*, pp. 482-525. London: Palgrave Macmillan. ISBN: 978-1-349-26928-0, DOI: 10.1007/978-1-349-26928-0\_15.
- Olson, M. (1971) *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Olson, M. (1982) *The Rise and Decline of Nations: Economic Growth, Stagflation, and Social Rigidities*. Yale, USA: Yale University Press.
- Ostrom, E. (1991) *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Ostrom, E. (2005) 'The Complexity of Collective Action Theory', *Workshop in Political Theory and Policy Analysis, Indiana University*: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.563.7482&rep=rep1&type=pdf>
- Ostrom, E., and Ahn, K. (eds.) (2010) *Foundations of Social Capital*. Cheltenham, UK: Edward Elgar Publishing.
- Parsons, T. (1968) *The Structure of Social Action* (2 volumes). New York: Free Press.
- Peirce, C. S. (1903) 'The Nature of Meaning', in Peirce, C. S. (1998) *The Essential Peirce, Selected Philosophical Writings, Vol. 2*. Bloomington: Indiana University Press.
- Peirce, C. S. (1905) 'What Pragmatism Is', in Peirce, C. S. (1992) *The Essential Peirce, Selected Philosophical Writings, Vol. 1*. Bloomington: Indiana University Press.
- Perea, A. (2010) 'Backward Induction versus Forward Induction Reasoning', *Games*, 1(3), pp. 168-188.
- Polanyi, M. (1941) 'The Growth of Thought in Society', *Economica New Series*, Vol. 8, No. 32, pp. 428-456.
- Polanyi, M. (1951) *The Logic of Liberty*, London: Routledge & Kegan Paul.
- Popper, K. (1959) *The Logic of Scientific Discovery*. London (UK): Hutchinson & Co.
- Popper, K.R., and Eccles, J.C. (1977), *The Self and Its Brain*. Berlin: Springer.
- Posner, R. (1973) *Economic Analysis of Law, 2nd Edition*. Boston: Little, Brown.
- Posner, R. A. (1980) 'A Theory of Primitive Society, With Special Reference to Law', *Journal of Law and Economics*, Vol. 23(1), pp. 1-53.
- Prigogine, I. (1967) Dissipative structures in Claesson, S. (Ed) *Chemical Systems, Fast Reactions and Primary Processes in Chemical Kinetics*, pp. 371-382.

- Putnam, R. (1993) *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton: Princeton University Press.
- Radzicki, M. J. (1988) 'Institutional dynamics: An extension of the institutionalist approach to socioeconomic analysis', *Journal of Economic Issues*, Vol. 22, pp. 633–665.
- Radzicki, M. J. (2003) 'Mr. Hamilton, Mr. Forrester, and a foundation for evolutionary economics', *Journal of Economic Issues*, Vol. 37, pp. 133–173.
- Radzicki, M. J. (2009) 'Convergence of the social fabric matrix and complex modeling' in Natarajan, T., Elsner, W., & Fullwiler, S. T. (Eds.), *Institutional analysis and praxis: The social fabric matrix approach*, pp. 69–86. New York: Springer.
- Radzicki, M. J. (2010) 'Was Alfred Eichner a system dynamicist?' in Lavoie, M., Rochon, L.-P., & Seccareccia, M. (Eds.) *Money and macrodynamics: Alfred Eichner and post-Keynesian economics*, pp. 3–21. Armonk, NY: M.E. Sharpe.
- Railsback, S. F. and Grimm, V. (2011), *Agent-Based and Individual-Based Modeling: A Practical Introduction*. Princeton University Press, ISBN: 0691136742.
- Ritzer, G., and Stepnisky, J. (2013) *Sociological Theory (International Edition)*. New York: McGraw-Hill.
- Rizzello, S., and Turvani, M. (2000) 'Institutions meets mind: the way out of an impasse', *Constitutional Political Economy*, Vol. 11, pp. 165–180.
- Rizzello, S., and Turvani, M. (2002) 'Subjective Diversity and Social Learning: A Cognitive Perspective for Understanding Institutional Behavior', *Constitutional Political Economy*, Vol. 13, pp. 201–214.
- Robinson, R. (1963) *Definition*. Oxford (UK): Oxford University Press.
- Root, H. L. (2013) *Dynamics Among Nations*. MIT Press.
- Root, H. L. (2018) 'Keynes, Hayek, and the Roots of Complexity Theory in Economics', *Journal on Policy and Complex Systems* 4(1).
- Rosser, J. B. (1999) 'On the Complexities of Complex Economic Dynamics', *Journal of Economic Perspectives*, 13(4), pp. 169-192.
- Rosser, J. B. (2009), *Handbook of Research on Complexity*. Edward Elgar Publishing. ISBN: 9781845420895.
- Roth, A. E., and Erev, I. (1995) 'Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term', *Games and Economic Behavior*, Vol. 8, pp. 164-212.

- Rotherleder, N. J. (1996) 'Induction: Inference and Process':  
<https://www.eecs.umich.edu/techreports/cse/96/CSE-TR-288-96.pdf>
- Rumsfeld, D. (2002) *Defense.gov News Transcript: DoD News Briefing – Secretary Rumsfeld and Gen. Myers, United States Department of Defense*, February 12, 2002.
- Rutherford, M. (1995a) 'The Old and the New Institutionalism: Can Bridges Be Built?', *Journal of Economic Issues*, 29(2), June 1995, pp. 443-451.
- Rutherford, M. (1995b) *Institutions In Economics: The Old and the New Institutionalism*. New York: Cambridge University Press.
- Rutherford, M. (2001) 'Institutional Economics: Then and Now', *Journal of Economic Perspectives*, 15(3), pp. 173-194.
- Rutledge, R. B., Skandali, N., Dayan, P., and Dolan, R. J. (2014) 'A computational and neural model of momentary subjective well-being', *PNAS*, 111(33).
- Saam, N. J., and Harrer, A. (1999) 'Simulating Norms, Social Inequality, and Functional Change in Artificial Societies', *Journal of Artificial Societies and Social Simulation*, 2(1).
- Samuels, W. J. (1972) 'The scope of economics historically considered', *Land Economics*, 48(3), pp. 248-268.
- Samuelson, P. A. (1937) 'A Note on the Measurement of Utility', *Review of Economic Studies*, Vol. 4(2), February, pp. 155-161.
- Samuelson, P. A. (1969) 'Classical and Neoclassical Theory', in Clower, R. W. (ed) *Monetary Theory*. London: Penguin.
- Sandefur, T. (2009) *Some Problems with Spontaneous Order*, *The Independent Review*, Vol. 14 (1), pp. 5-25.
- Santa Fe Institute (2023) Complexity Explorer: Glossary. Available at <https://www.complexityexplorer.org/explore/glossary> (Accessed: 6 February 2023).
- Savarimuthu, T., Cranefield, S., Purvis, M. A., and Purvis, M. K. (2010) 'Obligation Norm Identification in Agent Societies', *Journal of Artificial Societies and Social Simulation*, 13(4).
- Sargent T. J. (1993) *Bounded Rationality in Macroeconomics*. Oxford (UK): Clarendon Press.

- Sawyer, R. K. (2001) 'Simulating emergence and downward causation in small groups', in Moss, S., and P. Davidsson (eds.) *Multi-agent-based simulation*, v. 1979, pp. 49–67. Heidelberg, Germany: Springer-Verlag.
- Sawyer, R. K. (2002) 'Emergence in sociology: Contemporary philosophy of mind and some implications for sociological theory', *American Journal of Sociology*, Vol. 108.
- Schelling, T. C. (1969) 'Models of Segregation', *American Economic Review*, Vol. 59, pp. 488-493.
- Schindler, J. (2012) 'Rethinking the Tragedy of the Commons: The Integration of Socio-Psychological Dispositions', *Journal of Artificial Societies and Social Simulation*, 15(1).
- Schofield, N. (1985) 'Anarchy, Altruism and Cooperation: A Review', *Social Choice and Welfare*, Vol. 2, pp. 207-219.
- Schotter, A. (1981) *The economic theory of social institutions*. New York: Cambridge University Press.
- Schultz, W. (2001) *The Moral Conditions of Economic Efficiency* (Cambridge Studies in Philosophy and Law). Cambridge: Cambridge University Press.
- Schumpeter, J. A. (1954) *History of Economic Analysis*. Great Britain: Routledge.
- Science Daily (2019) *Quantum Entanglement*. Available at [https://www.sciencedaily.com/terms/quantum\\_entanglement.htm](https://www.sciencedaily.com/terms/quantum_entanglement.htm) (Accessed 21 November 2019).
- Sciencing (2021) 'Coevolution: Definition, Types & Examples' Available at: <https://sciencing.com/coevolution-definition-types-examples-13719185.html> (Accessed: 28 January 2021).
- Searle, J. R. (1995) *The Construction of Social Reality*. London: Allen Lane.
- Searle, J. R. (2005) 'What Is an Institution?', *Journal of Institutional Economics*, 1(1), pp. 1-22.
- Seckler, D. (1975) *Thorstein Veblen and the Institutionalists*. Boulder: Colorado Associated University Press.
- Seltzer, N., and Smirnov, O. (2015) 'Degrees of Separation, Social Learning, and the Evolution of Cooperation in a Small-World Network', *Journal of Artificial Societies and Social Simulation*, 18(4).
- Sethi, R. (1999) 'Evolutionary stability and media of exchange', *Journal of Economic Behavior & Organization*, Vol. 40, pp. 233–254.



- Shackle, G. L. S. (1991) *Epistemics and Economics*. New Jersey (USA): Transaction Publications.
- Sheehan, J., and Wahrman, D. (2015) *Invisible hands: Self-organization and the eighteenth century*. Chicago: Chicago University Press.
- Simon, H. (1947) *Administrative Behavior: a Study of Decision-Making Processes in Administrative Organisations*. New York: Macmillan.
- Simon, H. (1957) *Models of man: social and rational; Mathematical essays on rational human behavior in a social setting*. New York: Wiley.
- Simon, H. (1976) 'From Substantive to Procedural Rationality', in Latsis, S. (ed) *Method and Appraisal in Economics*. Cambridge: Cambridge University Press, pp. 65-86.
- Simon, H. (1979) 'Rational Decision Making in Business Organizations', *American Economic Review*, 69(4), pp. 493-513.
- Simon, H. (1992) 'Decision-making and problem-solving', in Zey, M. (ed.) *Decision Making: Alternatives to Rational Choice Models*. Newbury Park: Sage Publications.
- Simon, J. (1987) 'Vygotsky and the Vygotskians', *American Journal of Education* Volume 95(4), pp. 609-613.
- Smith, A. (1776) *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: Strahan and Cadell.
- Smith, A. (2011) *The Theory of Moral Sentiments*. Urbana, Illinois: Gutenberg Publishers.
- Smith, J. M. (1982), *Evolution and the Theory of Games*. Cambridge University Press, DOI: 10.1017/CBO9780511806292.
- Smith, V. L. (1962) 'An experimental study of competitive market behavior', *Journal of Political Economy*, Vol. 70, pp. 111-137.
- Smith, J. M. and Price, G. R. (1973), 'The Logic of Animal Conflict', *Nature*, Vol. 246, pp. 15-18.
- Smith, V. L., Suchanek, G. L., & Williams, A. W. (1988) 'Bubbles, crashes, and endogenous expectations in experimental spot asset markets', *Econometrica*, Vol. 56, pp. 1119-1151.
- Smith, J. M., and Szathmáry, E. (1998) 'The Major Transitions in Evolution', Oxford University Press.

- Solow, R. M. (1985) 'Economic History and Economics', *American Economic Review*, 75(2), pp. 328-331.
- Southwood, N., and Eriksson, L. (2011) 'Norms and Conventions', in *Philosophical Explorations*, 14(2), pp. 195-217.
- Sperry, R. W. (1969) 'A modified concept of consciousness', *Psychological Review* Vol. 76, pp. 532-536.
- Sperry, R.W. (1991) 'In defense of mentalism and emergent interaction', *Journal of Mind and Behavior*, Vol. 12, pp. 221-246.
- Stacey, R. D. (2014) *Complexity and Group Processes: A Radically Social Understanding of Individuals*. Abingdon-on-Thames: Routledge
- Stanford University (2021) Stanford Encyclopedia of Philosophy. Available at <https://plato.stanford.edu/> (Accessed: 6 February 2023).
- Stanovich, K. (1986) 'Matthew Effects in Reading: Some Consequences of Individual Differences in the Acquisition of Literacy', *Reading Research Quarterly*, 21(4), pp. 360-407.
- . 'Money as a medium of exchange: An analysis with a genetic algorithm', Working Paper, Technical University, Vienna.
- Stigler, G. J., and Becker, G. S. (1977) 'De Gustibus Non Est Disputandum', *The American Economic Review*, Vol. 67(2), pp. 76-90.
- Stiglitz, J. E. (2001) 'Information and the Change in the Paradigm in Economics', *Nobel Prize Lecture*, 8 December, 2001.
- Stoelhorst, J. W. (2007) 'The Naturalist view of Universal Darwinism: an application to the evolutionary theory of the firm', in Hodgson, G. M. (Ed) *The Evolution of Economic Institutions: A Critical Reader*, pp. 233-251. Cheltenham: Edward Elgar.
- Sugden, R. (1989) 'Spontaneous Order', *Journal of Economic Perspectives*, Vol. 3(4), pp. 85-97.
- Sugden, R. (2005), *The Economics of Rights, Co-operation and Welfare*. London: Palgrave Macmillan. ISBN: 978-0-230-53679-1, DOI: 10.1057/9780230536791\_9.
- Takizawa, H. (2017), 'Masahiko Aoki's conception of institutions', *Evolutionary and Institutional Economics Review*, Vol. 14, pp. 523-540.

- Tesfatsion, L. (2006) 'Agent-Based Computational Economics: A Constructive Approach to Economic Theory', in Tesfatsion, L. and Judd, K. (eds.) *Handbook of Computational Economics: Vol. 2 Agent-based Computational Economics*. New York (USA): North-Holland Elsevier, pp. 831-880.
- Thatcher, M. (1995) *The Path to Power*. HarperCollins.
- Thebaud, O., and Locatelli, B. (2001) 'Modelling the emergence of resource-sharing conventions: an agent-based approach', *Journal of Artificial Societies and Social Simulation*, 4(2).
- Thorndike, E. L. (1898) 'Animal Intelligence: An Experimental Study of the Associative Processes in Animals', *Psychological Monographs*, Vol. 2(8).
- Toan, P., Winikoff, M., and Padgham, L. (2005) 'Learning Within the BDI Framework: An Empirical Analysis', *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems, 2005*.
- Tuckett, D. (2011) *Minding the Markets*. Springer.
- Tuomela, R. (2002) *The Philosophy of Social Practices: A Collective Acceptance View*. Cambridge (UK): Cambridge University Press.
- Turing, A.M. (1936) 'On Computable Numbers, with an Application to the Entscheidungs problem', *Proceedings of the London Mathematical Society*, Series 2, 42.
- Ullmann-Margalit, E. (1977) *The Emergence of Norms*. Oxford: Oxford University Press.
- Ullmann-Margalit, E. (1978) 'Invisible Hand Explanations', *Synthese*, Vol. 39(2), pp. 263-291.
- Urminsky, O., Bartels, D. M., Giuliano, P., Newman, G., Puntoni, S., and Rips, L. (2015) 'Choice and Self: How Synchronic and Diachronic Identity Shape Choices and Decision Making', *Marketing Letters*, 25(3), pp. 281-291.
- Vanberg, V. J. (1986) 'Spontaneous Market Order and Social Rules', *Economics and Philosophy*, Vol. 2, pp. 75-100.
- Vanberg, V. J. (1994a) *Rules and Choice in Economics: Essays in Constitutional Political Economy (Economics as Social Theory)*. Routledge.
- Vanberg, V. J. (1994b) 'Hayek's Legacy and the Future of Liberal Thought: Rational Liberalism versus Evolutionary Agnosticism', *The Cato Journal*, Vol. 14(2), pp. 179-199.

- Vanberg, V. J. (1994c) 'Cultural Evolution, Collective Learning, and Constitutional Design', In Reisman, D. (Ed.) *Economic Thought and Political Theory*, pp. 171—204. Boston: Kluwer Academic Publishers.
- Vanberg, V. J. (2001) 'Hayek's Legacy and the Future of Liberal Thought: Rational Liberalism versus Evolutionary Agnosticism' in Vanberg, V. (2001) *The Constitution of Markets: essays in political economy*, pp. 53-76. London: Routledge. Note: this is a longer version of [Vanberg \(1994b\)](#).
- Vanberg, V. J. (2011) 'Darwinian paradigm, cultural evolution and human purposes: On F. A. Hayek's evolutionary view of the market', *Papers on Economics and Evolution*, No. 1119, Max Planck Institute of Economics, Jena.
- Vanberg, V. J. (2018) *The Collected Works of F. A. Hayek*, Vol. 14. Oxford: Routledge.
- Vanberg, V. J., and Buchanan, J. (1988) 'Rational Choice and Moral Order' in *Analyse and Kritik*, Vol. 10, pp. 138-160.
- Veblen, T. B. (1898) 'Why is Economics Not an Evolutionary Science?', *The Place of Science in Modern Civilisation*. New York: Russell & Russell, pp. 56-81.
- Veblen, T. B. (1899) *The Theory of the Leisure Class: An Economic Study in the Evolution of Institutions*. New York: Macmillan.
- Veblen, T. B. (1900) 'The Preconceptions of Economic Science, III', in *The Place of Science in Modern Civilisation*, pp. 148-179. Reprint, New York: Russell & Russell.
- Veblen, T. B. (1914) *The Instinct of Workmanship, and the State of the Industrial Arts*. New York: Macmillan.
- Veblen, T. B. (1919) *The place of science in modern civilization and other essays*. New York: The Viking Press.
- Veenhoven, R. (1991) 'Is happiness relative?', *Social Indicators Research*, Vol. 24, pp. 1-34.
- Villatoroa, D., Andrighetto, G., Conteb, R., and Sabater-Mira, J. (2015) 'Self-Policing Through Norm Internalization: A Cognitive Solution to the Tragedy of the Digital Commons in Social Networks', *Journal of Artificial Societies and Social Simulation*, 18(2).
- Von Bertalanffy, L. (1968) *General System Theory: Foundations, Development, Applications*. New York: George Braziller.
- Vriend, N. J. (2002) 'Was Hayek an Ace?', *Southern Economic Journal*, Vol. 68(4), pp. 811-840.

- Waldrop, M. M. (1993) *Complexity: The Emerging Science at the Edge of Order and Chaos*. Pocket Books.
- Wallis, J. (2015) *Structure and change in economic history: The ideas of Douglass North*. Available at <https://voxeu.org/article/ideas-douglass-north> (Accessed: 24 October 2019).
- Watts, D. (2002) 'A Simple Model of Global Cascades on Random Networks', *Proceedings of the National Academy of Sciences of the United States of America*, 99(9), pp. 5766-5771.
- Weisberg, M. (2007) 'Who is a Modeller', *The British Journal of the Philosophy of Science*, Vol. 58(2), pp. 207-233.
- Whyte, J. (2019) The Invisible Hand of Friedrich Hayek: Submission and Spontaneous Order, *Political Theory*, Vol. 47(2), pp. 156-184.
- Wiesner, K., and Ladyman, J. (2019) *Measuring Complexity*, published online at <https://core.ac.uk/download/pdf/242723669.pdf> (accessed 8 March 2023).
- Williamson, O. E. (1975) *Markets and hierarchies, analysis and anti-trust implications: A study in the economics of internal organization*. New York Free Press.
- Williamson, O. E. (1980) 'The Organization of Work: A Comparative Institutional Assessment', *Journal of Economic Behavior and Organization*, 1(1), pp. 5-38.
- Williamson, O. E. (2000) 'The New Institutional Economics: Taking Stock, Looking Ahead', *Journal of Economic Literature*, 38(3), pp. 595-613
- Wolfram, S. (2002) *A New Kind of Science*. Champaign, Illinois (USA): Wolfram Media.
- Ya, H., Tan, F., Ding, M., Jia, Y. and Chen, Y. (2011) 'Sympathy and Punishment: Evolution of Cooperation in Public Goods Game', *Journal of Artificial Societies and Social Simulation*, 14(4).
- Yefimov, V. (2004) 'On pragmatist institutional economics', *MPRA Paper No. 49016*.
- Young, H. P. (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press. ISBN: 9780691086873.
- Younger, S. (2005) 'Reciprocity, Sanctions, and the Development of Mutual Obligation in Egalitarian Societies', *Journal of Artificial Societies and Social Simulation*, 8(2).
- Zellner, M., Wakins, C., Massey, D., Westphal, L., Brooks, J., and Ross, J. (2014) 'Advancing Collective Decision-Making Theory with Integrated Agent-Based Modeling and Ethnographic Data Analysis: An Example in Ecological Restoration', *Journal of Artificial Societies and Social Simulation*, 17(4).



# Appendix A

## Interaction Typologies

In this appendix we briefly look at alternative typologies presented by different researchers concerning human interactions. We present [Schultz's \(2001\)](#) framing as preferable and exhaustive<sup>1</sup>: this is used in the main text.

We begin with [Schotter \(1981\)](#) who developed [Ullmann-Margalit's \(1977\)](#) framing (Section [A.1](#)). Section [A.2](#) then discusses [Schultz's \(2001\)](#) typology.

### A.1 Schotter's Typology

[Schotter \(1981\)](#) includes four categories of interaction (the first three are taken from [Ullmann-Margalit, 1977](#)). The first three assume games “are played non-cooperatively or without communication among the agents.” ([Schotter, 1981](#), p. 28), which means the agents cannot make any binding agreements, including forms of utility transfer.

The four categories are:

1. Problems of coordination;
2. Prisoners' dilemma games;
3. Problems of inequality preservation; and
4. Cooperative games.

We look at these briefly in turn.

---

<sup>1</sup>Strictly speaking, it is exhaustive for interactions in which participants have overlapping preferences. This is discussed further below.

### Coordination Situations

The following  $2 \times 2$  game is a typical (pure) coordination-type interaction:

		Agent j	
		L	R
Agent i	U	1, 1	0, 0
	D	0, 0	1, 1

Table A.1: A Pure Coordination Game: “Matching Pennies”

If we rank the agents’ preferences we find (for both agents):  $UL = DR > UR = DL$ .

This is a simple coordination problem in the sense that the agents’ choices have to be coordinated in some way to ensure UL or DR. Both of these states are Nash equilibria, i.e., neither agent would want to change their strategy in these states.

### Prisoners’ Dilemma Games

In these interactions, assuming substantive rationality, the agents have strongly dominant strategies (D and R); however, the resulting outcome is pareto inferior.

		Agent j	
		L	R
Agent i	U	3, 3	0, 4
	D	4, 0	1, 1

Table A.2: The Prisoners’ Dilemma

If we were to rank their preferences, we find:

- Agent  $i$ :  $DL > UL > DR > UR$
- Agent  $j$ :  $UR > UL > DR > DL$

Note how  $UL > DR$  is common to both agents.

### Problems of inequality preservation

As [Schotter \(1981\)](#) notes, these are simply non-pure coordination games. Consider the game in Fig. [A.3](#) below.

		Agent j	
		L	R
Agent i	U	2, 1	0, 0
	D	0, 0	1, 2

Table A.3: A Non-Pure Coordination Game: “Battle of the Sexes”



Here the agents face a coordination situation but their preferences are not identical as they were in Table A.1 above. This time:

- Agent  $i$ :  $UL > DR > UR = DL$
- Agent  $j$ :  $DL > UL > UR = DL$

Schotter (1981) writes that for reasons of taxonomy it is helpful to classify them differently (p. 26) to pure coordination games.

### Cooperative Games

In this category of interaction, agents can communicate and bargain, or a social planner can attempt to influence the decisions of the agents. The preferences of the agents can take the form of the three interaction types listed above and many others.

## A.2 Schultz's Typology

Schotter's typology described above is a typical way of distinguishing between interaction types in the literature although many combine Schotter's first and third types as a single category of coordination games (pure and non-pure). Vanberg and Buchanan (1988), for example, distinguish between coordination and prisoners' dilemma situations.

Assuming we are only interested in non-cooperative games (as defined above), there are types of interaction other than coordination games and the prisoners' dilemma in which organic institutions might be helpful. This is the problem with Schotter's typology: it is non-exhaustive.

Consider the following game:

		Agent j	
		L	R
Agent i	U	2, 3	3, 4
	D	1, 2	4, 1

Table A.4: A Game with No Nash Equilibrium But Coinciding Preferences: "Around the Houses"

In this situation, the agents' preferences can be ranked as follows:

- Agent  $i$ :  $DR > UR > UL > DL$
- Agent  $j$ :  $UR > UL > DL > DR$

There are two interesting features here. First, the agents' preferences share some commonality. They are not identical but they both include:

UR > UL > DL.

The second interesting feature is that there is no Nash equilibrium.

The overlapping of preferences means there might be a role for an institution of some sort to bring about UR because it is preferred by both agents to two other potential outcomes. Clearly, however, if we assume substantive rationality then Agent  $i$  would prefer D over U if  $j$  chose R. This makes UR a challenging state to achieve: it is not a Nash equilibrium.

Schultz's (2001) framing is preferred in this thesis because it is exhaustive: it accommodates interactions like those in Table A.4 above (and it includes the Prisoners' Dilemma). In fact, it includes all interactions where agents might share preferences but where no Nash equilibrium exists.

Schultz (2001) notes how interactions can be described by two factors: (i) the agents' preferences; and (ii) their strategies.

The situations we are interested in are those where the agents share some commonality of preferences (this was the case in the games described in tables A.1 to A.4 above). Schultz (2001) describes situations where preferences strictly conflict as "moot" (p. 64). These types of interactions might be interesting for other reasons but here we limit ourselves here to where pareto superior outcomes can be achieved by institutions.

In respect of strategies, Schultz (2001) states that "[t]wo strategies coordinate ... if both were taken, both agents would achieve their desired social state." (p. 64). This appears to be identical to the concept of a Nash equilibrium: a combination of strategies where neither agent would choose to change.

From these two factors, Schultz (2001) develops two categories of interest, which he refers to as:

(1) coordination situations where "[a]t least one of each agent's preferences coordinates and each agent's best strategy coordinates with every other agent's best strategy." (p. 65); and

(2) collective action situations where "[a]t least one of each agent's preferences coordinates but each agent's best strategy conflicts with every other agent's best strategy." (ibid).

Put another way, in both situations agents' preferences share some commonality; but in the first, a Nash equilibrium exists at all the agents' (shared) most preferred state, whereas in the second there is no Nash equilibrium at this state.

We can use the four games illustrated above to illustrate Schultz's typology:

1. the pure coordination game fits in to the first category: UL and DR are preferred (equally) by both agents and both states are Nash equilibria;
2. the Prisoners' Dilemma game fits in to the second category: both agents prefer UL to DR but there is no Nash equilibrium at UL;
3. the non-pure coordination game fits in to Schultz's first category since  $DR > UR = DL$  and  $UL > UR = DL$  are shared by both agents, and both DR and UL are Nash equilibria; and
4. the "Around the Houses" game fits in to the second category: the agents' preferences share some commonality but there is no Nash equilibrium at the jointly preferred state (UR).

If we compare [Schultz's \(2001\)](#) typology with Schotter's (and that of many others), it appears the main difference is that researchers often use the Prisoners' Dilemma to represent Schultz's collective action situations. However, while Schultz's category contains the Prisoners' Dilemma, it is broader: it contains all situations when a pareto superior state cannot be achieved (assuming no enabling mechanism). This thesis uses [Schultz's \(2001\)](#) typology vis-à-vis types of human interaction because it seems, in a sense, more fundamental than others (including Schotter's).

The final point to note here is that [Schultz \(2001\)](#) categorises 'solutions' to coordination situations as *conventions*; and solutions to collective action situations as *normative constraints*.



# Appendix B

## Exploration of the Parameter Space: First Model

This appendix contains detailed discussions of the results when the parameter space of the first model was explored. A summary of this material was included in the main text (Chapter 8).

### B.1 Memory Decay

The default rate of decay of the agents' memories (denoted here as  $m_{dec}$ ) was 0.2, i.e., the weight of any grid square in memory decayed by 20% between rounds. To explore this,  $m_{dec}$  was adjusted between 0 and 1 in 0.05 increments. Three additional sets of simulations were run, when  $m_{dec} = 0.9501$ ,  $m_{dec} = 0.9751$ , and  $m_{dec} = 0.01$ , for reasons that are explained below.

#### B.1.1 Low Values of $m_{dec}$

There was one noteworthy observation when  $m_{dec}$  was below approximately 0.02, which is that in many simulations more than one market was sustained. An analysis of the data showed that very low values of  $m_{dec}$  slowed symmetry breaking.

For example, when  $m_{dec} = 0$  (which is when agents gave equal weight to every transaction they ever took part in), we observed a mean of 1.85 (and standard deviation of 0.48) markets (across 20 simulations) between rounds 900 and 1000.

The data showed significant variation between the 20 simulations analysed. For example, in one of the 20 simulations, two markets of approximately equal volume and attendance emerged and was sustained (see Fig. B.1 below). In six of the 20 simulations, one

dominant market emerged on its own, as in the default scenario. For the remaining thirteen simulations, two markets existed: one major and one minor.




Figure B.1: A time series of the total number of agents visiting specific target locations when the memory decay rate was zero and two markets of equal size emerged. These markets were maintained because symmetry breaking was weakened by the fact that all transactions were sustained in the agents' memories.

Very small increases in  $m_{dec}$  from 0 significantly increased the prevalence and speed of symmetry breaking. For example, when  $m_{dec} = 0.01$ , the mean number of markets over 20 simulations fell from 1.85 (0.48) to 1.20 (0.40) between rounds 900 and 1,000. In sixteen of 20 simulations, one dominant market emerged.

### B.1.2 Intermediate Values of $m_{dec}$

When  $(0.02 < m_{dec} < 0.95)$ , a single market always emerged by the end the 1,000th round, the agents specialised, and the total population stabilized at approximately 43 agents, as in the default simulations.

### B.1.3 High Values of $m_{dec}$

When the decay rate was set at 100% agents had no memories so the results of the second null scenario (Section 7.3)<sup>1</sup> were replicated.

When considering the impact of high values of the decay rate  $(0.95 \leq m_{dec} < 1)$ , it is worth recalling that locations were removed from memory when their weight declined

<sup>1</sup>This was when the agents had no memories but could specialise.

to below 0.05. This means that a single transaction at some location would not be remembered in the next round if  $m_{dec} > 0.95$ <sup>2</sup>.

Given that the proto-markets observed in the default scenario simulations were due to single transactions involving only two agents, we might expect that no markets would emerge when the  $m_{dec} > 0.95$ . We found, however, that this was not the case: in some simulations we occasionally observed 3 agents (sometimes more), who had walked around randomly and then transacted at the same location at least twice.

When  $0.95 < m_{dec} < 0.975$ , markets emerged in 55% of simulations: the agents' home locations were sufficiently close to each other that occasionally 3 (or more) agents would meet and record 2 or more transactions each. This meant all three of the agents would remember the location in the next round and go back to it. If all three agents did this and transacted more than twice each in every round, they would continue to return.

Fig. B.2 demonstrates this in one of the simulations. This time series chart agents' target locations when a market emerged. In this simulation, 3 agents met at location [25, 8] in Round 326 and transacted 10 times. All of these agents revisited this location in the next round when 7 transactions were recorded. One of these agents transacted 4 times and communicated this to another agent, which meant 4 agents visited the market in Round 328 (when 24 transactions were recorded). Sufficient volume was maintained at the market for these 4 agents to return and also for the location to be communicated to more agents, until most of the agents learnt of its location.

This example raises the question of why 3 agents would transact twice (or more) each. In fact, it was not unusual, as in the above example, for as many as 10 transactions to occur between 3 agents who happened to come across each other simultaneously. The reason for this is that transacting was bilateral and there was no mechanism built in to the model that optimized transactions between 3 agents. These two factors resulted in multiple transactions per agent.

Transactions between two agents were optimized in the sense that after any transaction, neither of the agents wanted to (or could) transact again. However, to achieve an equivalent optimization between 3 agents, we would have to calculate the Walrasian equilibrium price between the agents, given their holdings and personal resource arrays, and allow them to transact at that price only.

In these simulations, it was often the case that Agent  $i$  would transact with Agent  $j$ , and then with agent  $k$ , and then with Agent  $j$  again. Also, after transacting with  $i$ ,  $j$  would also transact with  $k$ , and so on. In general we saw several transactions between the 3 agents but in diminishing quantities. Approximately speaking, the three agents

<sup>2</sup>Two transactions at the same location would not be remembered if the decay rate was above 0.975, etc.




Figure B.2: A time series of the total number of agents visiting specific target locations when the memory decay rate was 0.9501. No market was sustained until Round 326 when 3 agents met at location  $[25, 8]$  and transacted (10 times). This market was sustained after it was reported to several agents who met and transacted, which ensured the market was maintained despite the agents' poor memories.

would iterate closer to a balance of resources that would have been achieved had the agents only traded at the equilibrium price. Sometimes this required 10 transactions.

It is important to mention here that a minimum value was imposed on the agents' transaction volume, of 0.01 units for either resource. This was in order to limit the run time of simulations: had agents been able to transact for infinitesimally small amounts of resources, the run time of the simulations would have been prohibitively long. This is important here because reducing this volume floor would have increased the number of transactions, even between 3 agents, which would have increased the transaction's weight in memory, *ceteris paribus*.

As it was, in the simulation shown in Fig. B.2 the number of transactions between the three agents was sufficient for them to continue returning to the same location and for other agents to be informed about this emerging market. Given a weight of 0.5 of other agents' transaction locations in memory, for an additional agent to visit the market it would have had to hear of at least four transactions for the location to remain in memory in the next round. This was observed. In fact, the number of transactions at the market per capita increased each time a new agent visited: approximately speaking, 3 agents created a mean of 3 transactions per agent and 4 agents created about 5-6 transactions



per agent<sup>3</sup>.

All of these factors meant markets emerged in 55% of the simulations when  $0.95 < m_{dec} < 0.975$ .

If we set  $0.975 < m_{dec} < 0.983$  it was only possible for locations to remain in memory if an agent had participated in 3 or more transactions. Given the discussion above, when  $m_{dec} = 0.9501$ , it should be clear that this was possible in principle. However, to be sustained, the initial transacting agents had to transact at least 3 times in every round and for another agent to join they had to hear about at least 6 transactions. These were very high thresholds for maintaining and increasing the participation at the market.

Ultimately, these thresholds were too high, which meant any markets that emerged were eventually forgotten. Fig. B.3 below shows a rare simulation, when a market emerged in Round 929 and was sustained for 40 rounds.




Figure B.3: A time series of the total number of agents visiting specific target locations when the memory decay rate was 0.9751. No markets were sustained until Round 929 when 3 agents met at location [32, 32] and transacted. This location was not, however, sustained because the agents had such poor memories that they eventually forgot about it.

When  $0.983 < m_{dec} < 0.990$ , four transactions were required for an agent to sustain a location in memory, which was highly improbable and no market was ever sustained.

<sup>3</sup>This exponential increase in transactions was part of the reason for capping transaction volumes at 0.01 units. It was also mitigated by limiting the agents to initiating only one transaction in any of the 50 time periods during the trading phase.

In summary, for simulations when  $m_{dec} > 0.975$  we saw, in essence, a replication of the second null scenario<sup>4</sup>, and when any markets that did emerge they were ephemeral.

## B.2 Population Density

The default torus was a grid of  $50 \times 50$  and the 25 agents at instantiation had home locations evenly distributed across this grid (Fig. 7.2). This meant that at the start of the simulation there were on average  $10^2$  grid squares per agent.

What happened if we changed the density of the agents? This was done by adjusting the number of grid squares per agent while keeping the number of agents constant at 25. The total number of moves in the trading phase was adjusted pro rata so agents could still access the whole torus: the total number of moves was kept equal to a grid dimension e.g., for a  $100 \times 100$  the agents were allowed 100 moves.

We found that if the population density was too low then no markets emerged. More specifically, if the population density of the grid was greater than approximately  $100^2$  grid squares per agent then the sparsity of the population meant agents never met. The agents would start the simulation moving randomly but the likelihood of any two agents meeting was trivially small so no transactions occurred and therefore no market emerged.

## B.3 Communication

In the default scenario each agent had a 1% chance of communicating with any other agent at the end of each round. Here we first look at eliminating communication altogether and then, second, at increasing the probability of communication from zero.

What happened if we eliminated communications between the agents? Specifically, was market emergence dependent on - or catalysed by - agents communicating about their transactions?

We found three things: (i) markets emerged more slowly than in the default scenario (this slower emergence meant it was necessary to increase the number of rounds in the simulation to 5,000 in order to explore this in more detail); (ii) multiple markets emerged and were sustained; and (iii) the turnover ratio was slightly lower than in the default scenario, at approximately 0.95 after 1,000 rounds, and most agents specialised.

The slower emergence of the markets was a result of agents not being able to learn of other agents' transaction locations: they either had to originate markets themselves by repeatedly going back to the same location of a previous transaction (which other

---

<sup>4</sup>When agents had no memories but could specialise.

agents had to discover) or by discovering markets that other agents had formed (via their random walk).

The speed of emergence is probably best demonstrated by looking at Fig. B.4 below, which shows the turnover ratio during the first 500 rounds of a typical simulation. This can be compared with Fig. 7.8, which shows the same metric over 500 rounds in a simulation that used the default parameters.




Figure B.4: Time series of the turnover ratio over 500 rounds when the agents did not communicate. The turnover ratio increased towards 1 more slowly than in Fig. 7.8: agents' convergence on markets was slowed down when they were prevented from communicating.

Multiple markets emerged in all of the simulations, with an average of 9.7 markets<sup>5</sup> sustained in the last 100 rounds of the simulations. This was because agents tended to remain faithful to 'local' markets they first encountered or originated: this prevented symmetry breaking.

The third point noted above is that the turnover ratio was below 1 on average, albeit is relatively high at about 0.95. An analysis of the data suggests that this slightly lower ratio was a result of new agents spending more time (than in the default scenario) finding a market.

What happen if we *increase* the probability of agent communication? Unsurprisingly, markets emerged much more quickly. In the extreme case of all agents communicating with all the other agents, every agent was informed about every transaction at the end

<sup>5</sup>In fact, only 2 agents visited some of these locations, which made them more like 'private understandings' between these agents than markets.

of the round it took place in. This meant that as soon as two agents bumped in to each other, this location became the market all agents visited for the rest of the simulation<sup>6</sup>.

## B.4 Trading on the Way to the Target Location

In the default scenario, when transactions were located in agents' memories they would only attempt to trade after they had reached their target destination on the grid. What if they are allowed to transact on the way to that target? This meant that if two agents were within sight of each other (on the same or an adjacent square), they could transact and then continue on their way to their target locations.

The results are broadly similar to the default scenario, i.e., markets emerged and agents specialised. The main difference was that one square typically emerged as a 'main' market (where 82.6% of transactions took place on average), with the remainder of transactions spread across other nearby squares. Also, the turnover ratio remained high despite transactions being more dispersed than in the default scenario (a mean of 0.96 here versus 1.00 in the default scenario). This enabled specialisation once again.

## B.5 Speed of Foraging Skill Change

Equation 7.3 above is the logistic equation used to update agents' foraging detection probabilities. The only parameter in this equation is  $t$ , which controls the rate of change of agents' skills: a higher value means agents' detection probabilities change more rapidly. Here we allow  $t$  to vary from its default value of 0.01.

The discussion of the emergence of specialisation in Section 7.5.2 is useful context for understanding changes in  $t$ . With this in mind, Fig. B.5 below shows the detection probabilities, the probability threshold (Equation 7.10), and the expected probability of transacting, for one agent in the first 500 rounds of a typical simulation using the default parameter set, i.e., when  $t = 0.01$ . It is equivalent to Fig. 7.12 but the chart below is shown over 500 rounds.

When analysing the simulation data, it was important to consider the value of  $t$  in the context of the availability of resources to the agents. This includes: the agents' starting resources (a mean of 50 units of each resource in the default simulations), their opportunity to forage for resources (the number of time slots during the foraging phase), the resources available at the fountains, and the number of agents competing for these resources.

---

<sup>6</sup>If there was more than one transaction in the first round in which any of the agents transacted then the agents would be fully informed about all of these locations. There would then be more than one market in subsequent rounds, until symmetry breaking meant one market dominated.




Figure B.5: A time series of agent specialisation. This chart depicts an agent’s skilled (blue line) and non-skilled (green line) resource detection probabilities, the probability threshold above which the agent chose to specialise (the red line), and the expected probability of transacting (the black line). The data is shown over the first 500 rounds of a typical simulation that used the default parameter set. The increase in the blue line to 1 depicts specialisation.

In the default scenario, agents typically saw their resource stocks decline by 5-10 units in approximately the first 50 rounds. Their stocks at instantiation were sufficiently high that none of the agents died in any of the simulations. After approximately 50 rounds, the agents’ detection skills typically increased, which raised the overall productivity of the agents (who also traded).

What happened when  $t$  was adjusted?

When  $t = 0$  the agents’ detection probabilities were static, i.e., they could not specialise. The results from the third null scenario were replicated (Section 7.4)<sup>7</sup>. The population fell to the environment’s carrying capacity of 15-16 agents where it stabilized and where the agents’ foraging yields were about 1 on average for each resource.

When  $t$  was relatively low but non-zero (e.g.,  $t = 0.0025$ ) it took longer for the agents to specialise. See Fig. B.6 below, which is equivalent to Fig. B.5 above but over 1,000 rounds. Note how the agent shown in this chart did not fully specialise until almost Round 1,000.

<sup>7</sup>This is when the agents had memories but could not specialise.

A low value of  $t$  meant that, typically, a few agents died because they did not specialise quickly enough. However, it was never the case that all the agents died: most of the agents lived long enough to specialise, survive, and bear children.




Figure B.6: A time series of agent specialisation when the speed of adjustment of detection probabilities was set very low, at 0.0025. See Fig. B.5 above for a description of the data, which here is shown over 1,000 rounds. The chart shows how specialisation was slowed down in these simulations.

A high value of  $t$  had a predictable effect: agents specialised much more quickly. For example, when  $t = 0.4$ , most of the agents were fully specialised by Round 100. None of the agents died and they bore children earlier.

## B.6 Randomizing Agent Home Locations

In the default scenario, agents were placed evenly on a grid (Fig. 7.2), which had the effect of maximizing the space between agents (and minimizing the probability they met while walking around randomly). What if agents were assigned random home locations on the grid?

An analysis of the data showed there was no discernible difference between these simulation results and those of the default scenario. An examination of the data indicated that markets had a slight bias toward emerging in locations closer to any clusters of homes. This should not be surprising: agents in close proximity were more likely to bump in to each other when walking around randomly in the early stages of the simulations.

Another interesting result was observed when we assigned home locations randomly and prevented the agents from using their memories to find each other, i.e., they only ever walked around the grid randomly. Fig. B.7 below shows the home locations of the agents surviving after the 10,000th round of a typical (extended) simulation. It shows a type of agglomeration effect whereby agents were more likely to survive if they were located close to other agents.




Figure B.7: Agent home locations after 10,000 rounds when the agents did not store transaction information and when their homes were allocated randomly. Agents walked around the grid randomly and tended to survive if their home locations were in close proximity to others. This is a type of agglomeration effect.

This agglomeration effect was sufficiently strong to ensure a high turnover ratio (an average of 0.84 by the end of the 10,000th round across all simulations). The agents whose homes were clustered ended up specialising and bearing children.

In effect, what happened here was that the clustering of homes took on the role of the market in the default scenario.

## B.7 Expectations Accuracy When Adjusting Foraging Strategies

When the agents adjusted their foraging strategies at the end of each round they had to estimate their expected total foraging yields under different scenarios. This was done by the model calculating accurate estimates and then adding an error term ( $\epsilon$ ), taken from a normal distribution with mean 0 and standard deviation ( $IQ_i$ ) of 0.1. We can

think of this standard deviation as representing the agents' ability to estimate foraging yields: the higher it was, the less accurate the agents' estimates were. Here we look at the results of varying this metric.

The data tell a simple story: (i) changing the standard deviation had little impact on the emergence of the markets (the turnover ratio remained high), which should not be surprising because this value was not directly relevant in this process; (ii) when the standard deviation was zero, the results of the default simulation were replicated, essentially, with the agents specialising slightly more rapidly (they made no errors); and (iii) the higher the standard deviation, the less likely it was agents specialised.

When the standard deviation was approximately  $0.1 < \epsilon < 0.7$ , agents specialised more slowly or not at all. To illustrate this, Fig. B.8 below shows the emergence of specialisation for a typical agent when the standard deviation was 0.4. This chart can be compared with Fig. B.5 above. In Fig. B.8, the agent reached skill perfection by approximately Round 500 whereas in the default scenario this happened by approximately Round 300.




Figure B.8: A time series of agent specialisation when the agents made inaccurate estimations of future foraging yields (here,  $IQ_i = 0.4$ ). See Fig. B.5 above for a description of the data, which is shown in this figure over 1,000 rounds. The chart shows how specialisation was slowed down relative to the default simulations but the agent nonetheless specialised.

For extremely high values of the standard deviation (approximately  $\epsilon > 0.7$ ), agents in effect adjusted their foraging strategies by the flip of a coin, i.e., foraging strategies were essentially random. Agents, therefore, did not specialise. Consequently, the results were



similar to the third null scenario (Section 7.4), when the agents had memory but were not allowed to specialise.

## B.8 The Weight of Other Agents' Transactions

Agents' memories of transactions were made up of (i) those in which they were a counterparty; and (ii) those they were informed about by other agents. In the default scenario, the former were given a weight of 1.0 in memory and the latter a weight ( $\beta$ ) of 0.5. What happened if this latter parameter was adjusted?

A weight of 0 had the same effect as when there was no communication between the agents (see Section B.3): symmetry breaking was eliminated and approximately 10 markets emerged.

An analysis of the data of simulations that used various positive values for this weight showed there is one point worth noting here: only a very small weight was required for a single market to emerge but this had to exceed 0.0625 for reasons mentioned below.

In the early stages of the simulations, the agents who had not originated a proto-market (approximately 15 agents in the default scenario) were informed about other agents' transactions. These would be the only transactions known to this group of agents. They would head toward the locations associated with these transactions even if the weight in memory of these locations was relatively small. However, a weight of above 0.0625 was required: any transaction communicated to the agent could only exist in memory in the next round if this were true because locations were removed from memory if the weight fell below 0.05 and weights decayed by 20% between rounds.

In colloquial terms, having a single location in memory with a very small weight is equivalent to a person hearing about a vague, unsubstantiated rumour of a transaction from a non-credible source. In the absence of any other information, the agent may as well check this location out. If they did not transact, they forgot about this location in the next round.

## B.9 Travel Distance

In the default scenario, agents could travel for up to 50 grid squares, which meant in principle they could traverse the whole torus (which had dimensions of  $50 \times 50$ )<sup>8</sup>. If an agent moved to a target square, it waited at this location for the remainder of the 50 moves and attempted to trade if they met other agents.

---

<sup>8</sup>Because it was a torus, agents could access any square within 25 moves.

What if this travel distance is allowed to vary? Here we limit the number of moves in the trading phase in which agents can move.

This was done by splitting the agents' time on the grid in to two: the number of time period in which they could travel (say, in moving to a target location) and the number of time periods they had for waiting (for other agents) and transacting. The latter was fixed at 25.

For example, in one of the simulations sets the agents were limited to 6 moves for travel and 25 moves for waiting and transacting.

Table B.1 below shows some key data for travel distances of between 6 and 24 grid squares, in increments of 2.

Included in the table is data labelled 'serviced location': this is the percentage of agent home locations falling within any market catchment area (defined as the grid squares within travel distance from any location where transactions had occurred in the previous 4 rounds).

We should note that the agents' homes were placed at 10 grid square intervals (Fig. 7.2), which means there were 9 grid squares between the agents: it was therefore impossible for agents to meet if the travel distance was 4 grid squares or less.

The three main observations from Table B.1 were that shorter travel distances were consistent with: (i) more squares with transactions; (ii) a lower turnover ratio; and (iii) a lower serviced ratio. We observed that in all these simulations, most of the agents were able to specialise and subsequently bear children.

Fig. B.9 and Fig. B.10 below help us to visualise what happened during the simulations. The first figure is a heatmap of transactions in the last 100 rounds of one simulation when the travel distance was 10 grid squares; and the second shows the 'catchment areas' of six emergent markets (the green boxes) in addition to the agents' home locations (the diamond shapes). The catchment areas are represented by grid squares that were within reach of the markets (darker greens were within the catchment areas of more than one market).

In Fig. B.10 we can observe that all but two agents fell within the six observed market catchment areas. In fact, on average 93.48% of agents' homes were within the catchment area of at least one market when the travel distance was 10 grid squares (in the last 100 rounds of all 20 simulations). The right hand column of Table B.1 shows this metric for a range of travel distances.

The emergence of multiple local markets raises a question about the optimality of market locations: even when travel distances were limited to 6 grid squares it was still possible (in theory) to achieve 100% market coverage and a turnover ratio of 1, e.g., if markets

Travel Distance	Mean Max Spec Value	Mean Max Detection Prob	Squares with Trans's	Turnover Ratio	Serviced Loc's (%)
6	4.17 (0.10)	0.81 (0.21)	14.80 (2.14)	0.87 (0.08)	77.90 (4.47)
8	4.37 (0.10)	0.87 (0.19)	12.20 (2.04)	0.85 (0.06)	88.18 (3.67)
10	4.39 (0.08)	0.90 (0.17)	7.40 (1.36)	0.92 (0.07)	93.48 (2.44)
12	4.48 (0.02)	0.92 (0.17)	5.20 (0.98)	0.92 (0.05)	93.49 (1.66)
14	4.53 (0.09)	0.94 (0.15)	4.00 (0.00)	0.96 (0.05)	99.18 (0.22)
16	4.53 (0.09)	0.93 (0.15)	5.00 (1.79)	0.94 (0.04)	97.46 (1.68)
18	4.50 (0.09)	0.93 (0.15)	3.80 (0.75)	0.93 (0.04)	97.86 (0.98)
20	4.52 (0.09)	0.94 (0.15)	2.60 (1.20)	0.95 (0.03)	95.70 (1.66)
22	4.55 (0.10)	0.95 (0.13)	2.00 (0.00)	0.95 (0.03)	99.65 (0.45)
24	4.57 (0.08)	0.94 (0.15)	1.30 (0.90)	0.99 (0.00)	99.28 (0.29)

Table B.1: Results of simulations when the agents' travelling distance on the grid was adjusted (last 100 rounds only). The data show that as the distance was increased, agents congregated at fewer markets, the turnover ratio increased, more agents were within the catchment area of at least one market, and the agents became more specialised.

Notes:

1. Mean max spec value is a scale from 3 to 5 where 3 means the agents were generalists and 5 means they were specialists.
2. Mean max detection probabilities measures the highest of each agent's two detection probabilities and then takes the mean of these for all agents.
3. Squares with transactions measures the total number of squares on the grid in which agents traded.
4. The turnover ratio is the actual net transactions divided by the market clearing volume of trade.
5. Serviced locations refers to the average percentage of agent homes locations that were within at least one market's catchment area, i.e., when the market was within the travel distance from an agent's home.

were placed evenly across the grid between agent homes. The fact this did not happen indicated there was a coordination failure.

## B.10 Selection of Target Location: Winner Takes All

In the default scenario, when agents had memories of previous transaction locations they selected their 'target' location using a 'Roulette Wheel' approach: agents selected a location in memory randomly but with probabilities of selection proportional to their weight in memory.

An alternative selection process is 'Winner Takes All', which is when the location with the largest weight is chosen.




Figure B.9: Heatmap of transactions in the last 100 rounds of a typical simulation when the agents' travel distance was limited to 10 grid squares. The chart shows that several 'local markets' emerged when the agents could not traverse the whole grid.




Figure B.10: Market catchment areas in the last 100 rounds of the same simulation depicted in Fig. B.9. Each green square shows locations within 10 grid squares of the market; and the 'diamonds' represent agent homes. We can see that all but two agents lived within at least one catchment area.

The results of these simulations showed no discernable difference with those that used the default parameter set. Though minor, the only interesting feature was the mean number of squares with transactions in the last 100 rounds of the simulations (1.1 versus 1.0 in the default scenario). This was because symmetry breaking was more difficult when the agents used 'Winner Takes All', which meant occasionally more than one market

location existed toward the end of the simulations.

## B.11 The Agents' Initial Resource Endowment

In the default scenario, agents were endowed with two resources in their personal resource arrays: both these values were drawn from a normal distribution with mean 50 and standard deviation of 5<sup>9</sup>. Here we examine whether varying either of these numbers influences the simulation results.

As mentioned above, the availability of resources to the agents in the default scenario meant that, typically, they saw the resources in their personal resource arrays decline by approximately 5-10 units in the first 50 rounds of the default scenario. After this, agents tended to specialise, which meant their foraging yields increased as did the resources in their personal resource arrays.

It should be clear from this that changes in the mean allocation of resources at instantiation would only make a significant difference if it was below approximately 15 units.

The results from simulations that set the initial mean resource allocations at 5, 10, and 15 units demonstrated that this was indeed the case. For example, when the agents' initial resource endowments were fixed at 5 units of each resource, we typically saw the population decline to approximately 12-14 agents by Round 100 (see Fig. B.11 below, which shows the agent population in a typical simulation). This was slightly below the carrying capacity of the environment, which begs the question of how this could happen.

The answer seems to lie in the combination of perturbations / noise during the foraging phase and the 'asymmetry of death': when an agent had few resources in its personal resource array and was particularly successful in foraging for a few rounds, its resources simply increased. However, if it was particularly unsuccessful, one or both of its resources declined to or below zero, so it would die. This asymmetry meant the agent population tended to undershoot the carrying capacity slightly before recovering on account of trading and specialisation.

---

<sup>9</sup>Note that the starting resources for children were adjusted in line with this, i.e., each child was instantiated with the same mean value as the original agents. When a child was born, half of this value was deducted from each of the parents' personal resource arrays. For example, if the original agents began with 20 units of each resource, the children did so too and 10 units were deducted from both resources in the parents' personal resource arrays.




Figure B.11: The total agent population, births and deaths when agents were born with 5 units of each resource over 1,000 rounds of a typical simulation. The total population initially declined to below the carrying capacity, which enabled 11 agents to survive. Agents eventually specialised and the increased productivity meant they bore children.

In these simulations, markets were formed before the population declined significantly and these were sustained by the agents who survived the initial population decline. The surviving agents traded with each other, specialised, and subsequently bore children<sup>10</sup>.

One other point worth noting about the simulations where agents started with 5 units of each resource was that the population plateaued at approximately 36 agents, which was less than in the default scenario where the equivalent number was 43. This was because children were born with initial resources of 5 each and in an environment that contained a large number of specialised agents who crowded them out during the foraging phases. When the population was 36 or more, none of the children born had the opportunity to specialise and survive.

For initial resources of between 5 and 15 units, we saw a similar pattern but with less of a decline in the population.

<sup>10</sup>In the extreme case of agents being endowed with no resources at all, the population fell to approximately 2-6 agents before recovering (unless only 1 agent survived in which case no children could be born). The reason why the agents did not all die immediately was because they would only be removed from the simulation at the end of the round if either of their resources declined to zero or less. This meant they could still forage and trade in the first round, allowing those agents who successfully collected more than one of each resource in that round to survive. The collapse in the population at the end of the first round then made it easier for these agents to continue surviving (the population was well below the carrying capacity) until markets were formed and they specialised.

Given the above analysis, the impact of varying the standard deviation of the initial resource endowment should be fairly obvious (assuming a mean of 50 units). Irrespective of the standard deviation, provided an agent started the simulation with resources in excess of approximately 3 units of each resource, they survived. Most of those with fewer resources died. For the population as a whole, the results of the simulations showed that there were always a sufficient number of agents who survived, formed markets and bore children for the population to always recover after an initial decline (rising to approximately 43 as in the default scenario).

## B.12 The Fountains' Initial Endowments

The environment for agents was relatively harsh by design: Section 7.2 above showed that not all the agents were able to collect sufficient resources to survive if they could not trade or specialise: we saw that the carrying capacity of the environment was about 15-16 agents.

In the default scenario, each fountain began each round with 50 resource units ( $L = 50$ ), or 2 units of each resource per agent at instantiation. Recall, also, that agents had five time slots in which to forage for resources during the foraging phase of each round.

When simulations were run for values of  $L \geq 1$  we found two main results: first, the population declined to 0 or 1 when  $L < 7$  (approximately); and, second, for values of  $L > 7$ , the resulting total population stabilized at a level that was approximately  $0.85 \times L$  (this was consistent with the default simulation results when  $L$  was 50 and the population stabilized at approximately 43 agents).

The decline of the population to 0 or 1 when  $L \leq 7$  is perhaps surprising if we consider that a value of  $L = 7$  should in principle allow 7 agents to survive. There were two problems: first, agents started the simulations with detection probabilities of 0.5, which meant their foraging yield was expected to be approximately 2.2 units of resources (in total) even if no other agents competed for resources<sup>11</sup>; and, second, 25 agents began each simulation so the competition for resources was fierce (this meant agents' personal resources declined rapidly in the early stages of the simulations). This combination of factors meant that when  $L \leq 7$ , either no agents, or just one agent, survived<sup>12</sup>.

When  $L = 8$ , an average of 2.5 agents (over 20 simulations) survived an initial decline in the population over the first 100 rounds. In most simulations, a market formed early on when most of the initial 25 agents were alive. In some, however, it took several

<sup>11</sup>Recall that one of each resource was deducted as a metabolism cost at the end of each round.

<sup>12</sup>If the agents started the simulations as specialists and there was a liquid market they knew about then a small group of agents could have survived. However, the agents started these simulations as generalists.

hundred rounds for 2-3 surviving agents to find each other. When the simulations were extended we saw that, occasionally, it took thousands of rounds for 2 agents to find each other but, if they did, they traded regularly and eventually specialised in different resources. Children were then born and the population increased. In a steady state, the environment sustained 6 specialised agents. Occasionally a seventh agent was born but it never survived.

For higher values of  $L$ , more agents survived the initial population decline and, as mentioned above, the agent population tended to stabilize at a value of approximately  $0.85 \times L$ .

Another question worth exploring here is whether abundant resources would lead agents not to specialise: would a market still emerge and would agents still specialise in this scenario? To test this,  $L$  was set at 200.

An analysis of the results indicated that the agents' behaviour was the same even if resources were not in short supply: agents sought out others to trade with, a market formed, and a high turnover ratio led the agents to specialise. When  $L = 200$  the population stabilized at approximately 170 agents.

## B.13 Geographically Locating Resource Fountains

In the default scenario we assume that the resource fountains are not geographically located. Here, we assume that the fountains are given random locations on the grid, and that the agents started trading at the last fountain they visited during the foraging phase (as determined by their foraging strategy arrays).

An analysis of the data showed there was no significant difference between these simulations and those of the default scenario vis-à-vis the turnover ratio, specialisation, and agent births. However, there were a number of interesting points worth noting.

In the early part of the simulations, markets invariably emerged at the two fountain locations. Agents typically returned to the location where they first traded even if they subsequently started the trading phase at the other fountain (this happened when the last slot in their foraging strategy changed).

Fig. B.12 shows the evolution of the agents' marginal rates of substitution (MRSs) during the first 40 moves of the trading phase in Round 20 of a typical simulation. Recall that an agent's MRS was the ratio of its holdings of one resource divided by its holding of the other. Fig. B.12 shows these ratios for each agent over the first 40 moves during the trading phase of Round 20 of a typical simulation.






Figure B.12: Agents' marginal rates of substitution during Round 20 of a typical simulation when resource fountains were located on the grid and agents started the trading phase from the last fountain they visited. The chart shows time series for each agent during the trading phase: their marginal rates of substitution converged on two different values, corresponding to transactions at each fountain location (agents remained at these locations so there was no arbitrage).

At this stage of the simulation the agents were generalists and there was a market located at both fountains.

The most striking feature of this chart is the different convergences of MRSs at the two markets. All the agents travelled to their target (if they were not already there) and only transacted at that location, so there was no opportunity for arbitrage (this would require agents to transact at more than one location).

As the simulations progressed, all of the agents came to trade on a single market square that was located at one of the two fountains: symmetry breaking occurred here too.

Fig. B.13 below shows the MRS data during Round 139 of the same simulation as that depicted in Fig. B.12 above. A single market had emerged by this stage; however, agents were not fully specialised.

An interesting feature of this chart is the two-stage convergence of MRSs at the market. Here, the market was located at the Fountain for Resource A: the agents at that location traded with each other during the first 10 moves or so and their MRSs converged on 0.95 (approximately). This can be seen in the upper left part of the chart.

At this stage the agents at Fountain B did not trade with each other because agents could not trade until they reached their target (here, the market at Fountain A).




Figure B.13: Agents' marginal rates of substitution during Round 139 of the same simulation depicted in Fig. B.12. Resource fountains were located on the grid and agents started the trading phase from the last fountain they visited. The chart shows time series for each agent during the trading phase: here, the agents were not fully specialised but their marginal rates of substitution converged on a single value because all agents targetted the same location on the grid. This market was located at one of the two fountain sites.

Fountain B was about 13 moves away from Fountain A in this simulation: after the agents from Fountain B had arrived at Fountain A, there was a second flurry of transactions, this time involving all the agents. The MRSs eventually converged on 0.89 (approximately).

As with the default scenario, the turnover ratio was sufficiently high in these simulations that agents fully specialised. Fig. B.14 below shows the agents' MRSs during Round 999 of the same simulation.

The interesting point to note here is that the two-stage MRS convergence phenomenon had disappeared by this round. This is because the agents at Fountain A only had Resource A in their baskets so they could not trade. They had to wait for the agents from Fountain B to arrive before any of them could transact. Similarly, the agents who started at Fountain B had to wait until they reached Fountain A to trade. When this happened, all of the transactions in the round occurred within about 10 time periods and the agents' MRSs converged on 1.31 (approximately).

## B.14 Agents with Limited Life Spans

In the default parameter set, agents would die if one of their resources stocks declined to or below zero but they never died of old age.




Figure B.14: Agents' marginal rates of substitution during Round 999 of the same simulation depicted in figures B.12 and B.13. Resource fountains were located on the grid and agents started the trading phase from the last fountain they visited. The chart shows time series for each agent during the trading phase: the agents were now fully specialised and their marginal rates of substitution converged because the single market location enabled arbitrage. Specialisation meant none of the agents could trade until they all met at the market square.

Here we will examine the impact of limiting life spans.

This was done by adjusting two parameters. First, a maximum age for the agents (e.g., 500 rounds) was introduced in to the model and all the agents died when they reached this age. Second, each of the agents at instantiation were given an age drawn from a uniform distribution with a minimum of 0 and a maximum of the age of death. Children were born with an age of zero.

Sets of 20 simulations were run with ages of death ranging from 100 to 500 in intervals of 50.

The data showed that the total population always declined to zero when the age of death was set at approximately 200 or less. The higher the age of death, the more likely it was that the population remained above zero. For values of the age of death of 400 and above, the agent population was always sustained. Table B.2 below shows the mean total population (and standard deviation) in the last 100 of 1,000 rounds in the simulations.

To help understand these data better it is worth recalling that agents were endowed with resources at instantiation that were drawn from a distribution with a mean of 50; and that agents bore children as soon as two of them had 125 units of both resources.

Age of Death	Mean Population (last 100 rounds)	Pop'n Survived (% of simulations)
100	0.0 (0.0)	0
150	0.0 (0.0)	0
200	0.0 (0.0)	0
250	1.80 (3.8)	20
300	11.8 (8.5)	70
350	14.7 (6.6)	90
400	21.4 (2.8)	100
450	21.4 (2.8)	100
500	21.8 (2.4)	100

Table B.2: Adjusting life spans. The table shows the mean population in the last 100 of 1,000 rounds and the percentage of simulations in which the agent population did not collapse. Unsurprisingly, as the agent's life span was increased, the mean population increased and the population collapsed fewer times.

When the simulation data was analysed, two phenomena were important vis-à-vis whether individual agents survived. The first was the total population relative to the carrying capacity of the environment. When the total population was above the carrying capacity, the surviving agents' resources declined on average (assuming they were generalists); and vice versa. Furthermore, this effect was more significant the greater the difference between the total population and the carrying capacity.

The second phenomenon was the degree of specialisation (and, therefore, productivity) among the population.

Both of these factors worked against the agents at the beginning of every simulation: the initial population of 25 agents exceeded the carrying capacity of 15-16 agents; and the agents started as generalists. The population therefore declined as the older initial agents died and were not replaced; however, consistent with the default simulations, the surviving agents became more productive as they specialised.

In order for the agent 'species' to survive, the agents at instantiation had to specialise (and benefit from any decline in the population to below the carrying capacity) quickly enough for at least 2 agents to have sufficient resources to bear more than one child. It was also necessary for these children to accumulate resources quickly enough for them to bear more than one child themselves.

For lower ages of death (200 or less) all the agents died before this could happen. By contrast, in the simulations where the ages of death were 400 or more, the agent 'species' always survived.

Fig. B.15 below shows the total population in a simulation when the age of death was set at 300.




Figure B.15: Total agent population over 1,000 rounds in a typical simulation when the agents' lifespan was limited to 300 rounds. In this simulation the total population fell to just 3 agents but these lived long enough to specialise and bear children. Their children also specialised and bore children, etc., which allowed the total population to recover.

The chart shows how some of the agents at instantiation survived long enough to bear 3 children before dying. These 3 agents benefited from a high per capita resource availability due to the population being significantly below the carrying capacity. In addition, they formed a market that enabled them to specialise<sup>13</sup>. These agents had 6 children between them before dying and these new agents then bore 15 agents before they died. The population then oscillated around a mean of approximately 22 agents.

## B.15 Skill Acquisition: A Linear Approach

Equation 7.3, which has a logistic form, was used to update the agents' foraging skills at the end of each round.

What happened when a linear approach was used? To test this we replaced the logistic equation with a linear one (see below Equation 7.3 for definitions of terms):

$$\Delta p_i^j = t.[w_i^j - d/x]$$

<sup>13</sup>The initial population of agents died before they could inform the 3 children of the location of their market. Two of the children formed a proto-market that the third agent was eventually informed about.

The main difference between this equation and Equation 7.3 is that learning did not slow down at high and low levels of skills (but skills were limited to a maximum of 1 and a minimum of 0.2).

This had no impact on the results. Agents still specialised and bore children after markets emerged. The only marginal change from the default simulations was that for the same value of  $t$ , the linear approach meant agents reached full specialisation slightly more quickly<sup>14</sup>.

---

<sup>14</sup>When the logistic equation was used, skills increased (decreased) more slowly at higher (lower) levels of skill.

# Appendix C

## Exploration of the Parameter Space: Second Model

This appendix contains a detailed analysis of the exploration of the parameter space of the second model. A summary of these results was presented in Chapter 11.

### C.1 Adjusting Agents' Starting Propensities

In the default simulations each agent was instantiated with propensities to steal and defend drawn from normal distribution with mean 0.5 and a standard deviation of 0.1 (children inherited propensities equal to the mean propensities of their parents). The question we address here is whether the default simulation results are sensitive to these starting values.

This question was tested by running 5 experiments. The first used a uniform distribution of between 0 and 1 to set the initial agents' propensities. The other 4 used various combinations of high (0.9) and low (0.1) values for the mean of the agents' initial propensities<sup>1</sup>:

- (i)  $p_i^S = 0.1$  and  $p_i^{FB} = 0.1$  for all  $i$ ;
- (ii)  $p_i^S = 0.1$  and  $p_i^{FB} = 0.9$  for all  $i$ ;
- (iii)  $p_i^S = 0.9$  and  $p_i^{FB} = 0.1$  for all  $i$ ; and
- (iv)  $p_i^S = 0.9$  and  $p_i^{FB} = 0.9$  for all  $i$ .

---

<sup>1</sup>Values were again drawn from a normal distribution with this mean and a standard deviation of 0.1

We found that in all five experiments the results of the default simulations were (essentially) replicated. For experiments (i) and (iii) above it was necessary to increase the agents' starting resources (by 50%) to compensate for the increased cost of fighting (the agents' propensities had further to 'travel' in these simulations) to see the same results (otherwise the population collapsed).

By way of illustration, Fig. C.1 below shows the 'cloud' of all living agents' propensities to steal in the first 150 rounds of a typical simulation (and a mean value - the blue line) when a uniform distribution was used to determine the agents' initial propensities.




Figure C.1: The cloud of agents' propensities to steal over the first 150 rounds of a typical simulation when the agents' propensities were determined by a uniform distribution at the start. Each dot represents a living agent's propensity in each round and the blue line is the mean of these propensities. The propensities evolved as they did in the default simulations, approximately speaking.

Perhaps the most interesting of these four experiments was when the agents' propensities started at  $p_i^S = 0.1$  and  $p_i^{FB} = 0.1$  for all  $i$ . Fig. C.2 below shows the cloud of the agents' propensities to steal in a typical simulation: the mean increased from 0.1 to 0.71 in Round 11 before declining back to and then below zero. This pattern reflected something noted before: when the agents' propensities to defend were below approximately 0.8 on average, their propensities to steal increased, and vice versa.

## C.2 Adjusting the Cost of Fighting

In the default scenarios, the cost of fighting ( $c$ ) was 0.3 units of both resources.

A range of fighting costs were simulated for this section, from 0 to 2.0 for each resource. Here, we focus on what happened when these costs were at both ends of this spectrum. Overall we found that fight costs had to be between approximately 0.1 and 0.6 for the results of the default simulations to be replicated.






Figure C.2: The cloud of agents' propensities to steal over the first 250 rounds of a typical simulation when the agents' propensities to steal and defend started relatively low (0.1 on average). Each dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. In this chart, the propensities to steal increased aggressively to approximately 0.5 - 1.0 before declining to below zero.

### C.2.1 Fight Cost of Zero

When the fight cost was zero, the agents essentially passed the foraged resources between each other via transactions, fights, and muggings, without any cost if they fought.

We observed that the mean propensities to steal of all the agents increased toward (and beyond) 1, albeit with considerable variation. Fig. C.3 below shows the 'cloud' of agents' propensities to steal over a typical simulation.

The agents' propensities to defend also increased but these became redundant as the propensities to steal exceeded 1 (defending resources was only relevant to agents whose initial choice was to trade).

The main reason for this trend was the removal of an important 'discipline' on agents (fight costs), which tended to reduce propensities to steal.

An unintended consequence of fixing fight costs at zero was that many of the agents became specialised (though not fully) in these simulations. This was particularly surprising given that so far in this thesis specialisation has only been observed after the agents respected each other's property and markets emerged.

An analysis of the data showed that this was caused in large part by the resource concentration effect we have now observed several times. The gini coefficient of resource holdings at the end of the interaction phase was typically about 0.88 after most of the agents' propensities to steal reached above 1 (Fig. C.4 below shows resource holdings before, mid-way through, and at the end of the interaction phase of Round 800 of the




Figure C.3: The cloud of agents' propensities to steal over the first 1,000 rounds of a typical simulation when the cost of fighting was 0. Each dot represents a living agent's propensity to steal in each round. The chart shows how, in general, the agents' propensities to steal increased to and above 1: with no cost associated with fighting, agents did not learn to respect others' property.

simulation represented in Fig. C.3 above). This persistent concentration of resources meant agents would often go many rounds without consuming any resources.

Recall from the first model that if the turnover ratio was low (as it was in these simulations) agents would choose to forage for the resource in which they were deficient. If agents did not consume any resources for, say, 10 rounds, then this deficiency would be the same for all of these 10 rounds.

In fact, given the volatility of consumption patterns, agents would often be deficient in the same resource for hundreds of rounds. The agent would then persistently choose to forage from the corresponding fountain in all of its time slots in all these rounds, which would increase its foraging skill for that resource. To quantify this, the mean maximum detection probability of the living agents at the end of Round 1,000 in the simulation shown in figures C.3 and C.4 was 0.82.

This productivity, combined with no fight costs, meant the agents bore some children. If we extended the simulations to 10,000 rounds we found the total population of agents oscillated around 25 agents<sup>2</sup>.

<sup>2</sup>Recall that the carrying capacity of the environment was approximately 16 agents when the agents did not specialise and approximately 43 agents when they did.




Figure C.4: Agents' resource holdings at the start, middle, and end of the interaction phase of Round 800 of the same simulation depicted in Fig. C.3 above. In these charts, the agents' resource holdings were put in ranked order at the beginning of the trading phase (the first time period), in the middle (the 25th time period) and at the end (the 50th time period). The charts show how the agents' resource holdings became more concentrated over the round.

## C.2.2 Fight Cost of 2.0

When fight costs were fixed at 2.0 (for each resource), there were no transactions and no fights. All of the agents' propensities to steal and defend therefore remained static.

This lack of interaction was due to agents knowing their own propensities to steal and defend (all instantiated above zero) and the consequential risk of incurring a high cost of fighting were they to fight. The expected pay-off from any interaction was therefore always negative.

The outcome of these simulations was the same as when agents could only forage as in the first null scenario of Chapter 7: if we extended the simulations for long enough, the agent population declined to 15-16 agents. Fig. C.5 below shows the total agent population over 10,000 rounds in a typical simulation - the total population at the end of the simulation was 15 agents.

## C.2.3 Intermediate Fight Costs

The last two sub-sections have demonstrated the two main patterns vis-à-vis fight costs: at zero fight costs, agents were not discouraged from attempting to steal; and when these costs were high (approximately above 1), agents did not interact for fear of incurring these costs.




Figure C.5: Time series of the total agent population in a typical (extended) simulation when the cost of fighting was high. The agents ignored each other: the lack of transactions and specialisation meant the total population declined to 15 agents, which was the carrying capacity of the environment (in the absence of specialisation).

We found that the default simulation results were replicated when fight costs were between approximately 0.1 and 0.6. Below this range the disciplining effect was too weak for property rights to emerge across the population, and above this cost, agents were reluctant to interact, so they did not learn.

### C.3 Adjusting the Influence of Other Agents' Experience

Recall from Section 9.4<sup>3</sup> that agents learned from their own experience in interactions and they learned from their counterparties. In the default simulations, a parameter ( $\beta = 0.5$ ) was used as a weight attributed to other agents' experience<sup>4</sup>.

This weight of  $\beta$  was adjusted between 0 and 1. The main result is that the results of the default simulations were replicated in all of these experiments:  $\beta$  had a marginal effect on the emergence of property rights and the defence of property.

An analysis of the data showed that the main impact of adjusting  $\beta$  concerned the speed of change in the agents' propensities to defend. If  $\beta = 0$ , agents could only learn when they themselves acquiesced or defended their resources; whereas if  $\beta > 0$  they

<sup>3</sup>Table 9.1 shows the impact of  $\beta$  on the changes to the agents' propensities.

<sup>4</sup>Recall that the rationale was for agents to use all relevant information when learning, which includes what happened to their counterpart.

also learned when other agents acquiesced or defended their resources. The net result of  $\beta > 0$  was to accelerate this learning: when  $\beta = 1$ , the mean propensity to defend of living agents exceeded 1 after 27 rounds on average whereas when  $\beta = 0$  it took 55 rounds.




Figure C.6: The cloud of agents' propensities to steal over the first 175 rounds of a typical simulation when the agents accounted for the experiences of other agents as much as their own when learning ( $\beta = 1$ ). Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. The chart shows how more Al Capone agents emerged in these simulations but they eventually died: these are the four blue dotted lines that can be seen drifting higher in the chart. The surviving agents' propensities to steal declined to below zero as in the default simulations.

## C.4 Adjusting the Errors Made in Expected Pay-Offs

Recall from Section 9.5 that agents evaluated the payoffs from all six scenarios (for both agents) in every interaction and that an error was added to each pay-off (drawn from a normal distribution with mean 0 and standard deviation of 0.05). We can think of this standard deviation as scaling the degree of information-related uncertainty concerning each interaction.

Below we look at the effect of adjusting this standard deviation.

### C.4.1 Standard Deviation of Zero

When the standard deviation was set to zero, the agents had perfect knowledge of all the pay-offs. We found that the results of these simulations almost exactly replicated

the default simulations. The only point worth noting is that agents made fewer mistakes when deciding which agents to interact with, e.g., in the default simulations we sometimes observed agents with no resources attempting to steal from agents who also had no resources. This type of mistake never happened when the standard deviation was zero; however, the impact of this on the final results was immaterial: respect for property still emerged across the population.

### C.4.2 High Standard Deviation

When a very large standard deviation was applied to pay-offs, the agents in effect interacted with other agents randomly<sup>5</sup>. To examine the implications of this, an experiment was run in which the standard deviation was set at 100.

The main result of these simulations was that the living agents' propensities to defend and to steal remained at approximately 0.5 on average, and the agent population collapsed to a handful of agents who tended not to interact with each other. Fig. C.7 shows the 'cloud' of the agents' propensities to defend over the first 300 rounds of a typical simulation<sup>6</sup>. The agents' propensities to steal followed approximately the same pattern.




Figure C.7: The cloud of the agents' propensities to defend over the first 300 rounds in a typical simulation when the agents made very large errors in estimating their pay-offs in interactions. These propensities remained at approximately 0.5 on average over the whole simulation: in effect the agents interacted randomly. These simulations showed that for property rights to emerge it was necessary for agents to discern between other agents.

<sup>5</sup>This was not quite true: in general, agents preferred to interact with agents who had low propensities to steal and defend.

<sup>6</sup>Here, the mean of the agents' initial resource endowments were increased to 1,000 units each, for purposes of illustration.

It is also worth noting that more interactions took place between the agents in these simulations because those with resources tended not to avoid interacting with other agents (there were approximately 4 times as many fights and interactions as in the default simulations).

An analysis of the data showed that an important factor in the increase in the agents' propensity to defend in the default simulations was their discerning of potential counterparts: this discernment was in effect missing in the simulations with large pay-off errors. Simply put, in the default simulations there were many interactions involving initiating agents with no or few resource and counterparts with multiple units of resources: whether the interaction was in Scenario 2F or 3F, initiating agents generally learned it was better to defend in such situations. This was because if they lost the fight, the cost was small but if they won, the gain was large.

By contrast, when the pay-off error term was large, there was no such discernment: the initiators of interactions held more resources on average and counterparts fewer. Here, agents learned it was not preferable to defend their resources because they lost more of their own resources.

To quantify this, if we look at the total resource holdings of agents in scenarios 2F and 3F in the default scenario, we saw that initiating agents typically held an average of approximately 0.5 resource units but the counterpart held a mean of approximately 5 units. In the simulations when the pay-off error term was 100, both the initiators and counterparts held a mean of approximately 2 units each.

Overall, the agents learned that defending their resources was preferable when they acquiesced but learned it was better to acquiesce when they defended their resources: the two effects approximately balanced out and the propensities to defend remained at 0.5 on average.

#### *Resource Concentration*

One of the surprising observations made when the data from these simulations was analysed was that resources became concentrated in the same way they did in the default simulations. At first blush, one might expect resource concentration to be *caused* by discernment (agents chasing resources).

An analysis of the data showed that resources became concentrated even when agents interacted randomly (which is effectively what happened in these simulations). It was only necessary for the agents to have target locations within close proximity to each other for resources to become concentrated.

To check this, a (separate) simple computation model was designed in which 25 agents, each with 2 units of both resources, interacted randomly with another agent in multiple

rounds (when agents did not discern between other agents). We assumed there was a 25% chance any one agent would interact with another agent and a 50% chance that one or the other agent would win the other's resources. The number of rounds it took for resources to become fully concentrated (i.e., when one agent held all 50 resource units and the gini coefficient was 1) was recorded.

When 100,000 simulations of this model were run, we found the mean number of rounds it took for full concentration was 92.1 rounds (with a standard deviation of 51.5 rounds). By comparison, if we allowed the agents to target the agent holding most resources, this mean fell to 13.0 (4.0) rounds. We can conclude from this that both proximity of the agents and discernment catalysed the resource concentration effect.

### C.4.3 Intermediate Sizes of the Standard Deviation

As we reduced the standard deviation from 100 to 0, the agents become more discerning about who they interacted with and, *ceteris paribus*, on balance they learned it was preferable to defend their resources than to acquiesce.

We found that for the default simulation results to be replicated, the pay-off error had to be between approximately 0 and 0.6. Above 0.6, the agents were not sufficiently discerning which meant the agents' propensities to defend failed to rise high and quickly enough for the agents' propensities to steal to subsequently decline.

## C.5 Adjusting the Standard Deviation of Children's Starting Propensities

In the default simulations children were born with propensities to steal and defend equal to the means of the corresponding propensities of their two parents. Here we take the same approach but add an error term to each child's starting propensities.

In the default simulations we saw that the surviving agents' propensities to steal were all well below zero by the time the first child was born<sup>7</sup>. Children were typically born when the mean propensity to steal was below approximately -2. Similarly, the surviving agents had propensities to defend close to or above 1 when the first child was born.

A range of values were used to test the impact of the standard deviation of the children's starting propensities. In each experiment, each child was born with a propensity to steal drawn from a normal distribution with a mean equal to the mean of its parents'

---

<sup>7</sup>This should not be surprising. There was typically a gap of approximately 200 rounds between the moment that all surviving agents had negative propensities to steal and the first child being born. During this period the agents' propensities to steal declined further below zero and the agents became fully specialised.



propensities to steal and the test standard deviation (the child's propensity to defend was found using the same approach).

We found that for smaller values of the standard deviation (approximately below 0.7), all the children were born with propensities to steal below zero (because their parents' propensities were typically below -2). This meant the agents continued to respect property rights and they only ever traded (and propensities to steal continued to decline). Under such circumstances, the propensity to defend with which any child was born was irrelevant and remained static: they were never involved in interactions when they or their counterpart acquiesced or defended their resources, and therefore no learning took place.

For standard deviations above approximately 0.7, some children were born with propensities to steal above zero (the higher the standard deviation, the more children were born with positive propensities). Fig. C.8 below shows the living agents' propensity to steal 'cloud' over 1,000 rounds in a typical simulation when the standard deviation was set at 5.




Figure C.8: The cloud of the agents' propensities to steal over the first 1,000 rounds of a typical simulation when 'black sheep' children could be born (these were new agents instantiated with positive propensities to steal). Each blue dot represents a living agent's propensity to steal in each round. In this simulation, six of such children were born but they all died: the population proved resilient.

There are three noteworthy points. First, six children were born with propensities to steal above 0 ('black sheep') and all of them died within approximately 70 rounds of

their birth (these data look like blue ‘wisps’ in the chart). An analysis of the data showed these agents attempted to steal from agents with propensities to steal below 0 but these agents invariably defended their resources. The black sheep agents endured considerable fighting costs whereas the costs to the agents who defended their resources were, broadly speaking, spread across the population.

Second, a number of children were born with propensities to steal below that of their parents (we can see four of these in Fig. C.8).

Third, property rights remained relatively resilient to the ‘black sheep’, i.e., most propensities to steal remained below zero.

We can see in Fig. C.8 there was one agent with a negative propensity that became positive as a result of the agent’s interactions with the black sheep. This agent had been born in Round 437 with a propensity to steal of -1.7 and a propensity to defend of -2.6, i.e., the agent was very doveish. Eventually its propensity to steal increased to above 0 and its propensity to defend rose to -2.1: it learned it was preferable to steal from the black sheep instead of attempting to trade and then doveishly acquiescing. This agent eventually died as a result of the fight costs it incurred (and the loss of resources from acquiescing).

In any case, if we extended the simulations we found the respect for property rights endured: ultimately the total population increased to approximately 43 agents, all of whom had propensities to steal below 0. This was true of any standard deviation.

## C.6 Adjusting the Agents’ Initial Resource Endowments

In the default simulations, the 25 agents instantiated at the start were given personal resource arrays with two values, each drawn from a normal distribution with a mean of 200 and standard deviation of 5. Here we describe the results of simulations for different values of this mean.

In the first sub-section below we look at the two types of simulation observed in these experiments. We will look at the split of these two types when the starting resources were below or equal to 200 units (the second sub-section). The third sub-section look at the results when the starting resources exceeded 200 units.

The results can be summarised by stating that if the agents’ starting resources exceeded approximately 180 units, the results from the default simulations were replicated. For starting resources of below 160 units, the agent population was at risk of collapsing before property rights emerged. However, if 2 or more agents survived in these simulations,

the population typically recovered and the agents in the resulting population eventually respected each other's property.

### C.6.1 The Two Types of Simulation

To contextualise the results presented below, it is worth noting that in the default simulations, the cost of fighting during the 'learning phase' meant that typically approximately 14 agents lost all their resources and died, and those agents who survived typically saw their resource arrays decline to approximately 50 units of each resource, on average.

Furthermore, in the simulations discussed below, the simulations fell in to two broad categories: (i) population collapse; and (ii) where the emergence of property rights prevented such a collapse.

In the **first type of simulation**, the total population would collapse to approximately 1-3 agents (and the mean propensity to steal of the surviving agents was positive). In these simulations, if only one agent remained then no new agents could be born so the population never recovered<sup>8</sup>.

If the population collapsed then fountain resources per capita were relatively high, i.e., resources were abundant (per capita) because the population was well below the (non-specialised) carrying capacity of the environment. If two or more agents survived then, eventually, they accumulated resources and had children and, if the simulations were run long enough, the agent population saw their propensities to defend rise above 1 and their propensities to steal fall below 0, i.e., they came to respect property.

The **second type of simulation** were those seen in all the default simulations: the population did decline but in general the surviving agents saw their propensities to defend rise above 1 (mostly) and their propensities to steal declined below 0. This happened sufficiently quickly for a population collapse to be averted.

### C.6.2 Starting Resources of Below and Equal to 200 Units

Table C.1 below shows for various starting resource values, how many of 20 simulations run were of the first type of simulation and how many the second type.

The table shows a simple and neat shift from the first type of simulation to the second as the starting resources increased from 20 to 200 units.

---

<sup>8</sup>Very occasionally, all the agents died.

Starting Resources	Type 1 Simulations	Type 2 Simulations
20	20	0
40	20	0
60	20	0
80	19	1
100	15	5
120	9	11
140	4	16
160	3	17
180	1	19
200	0	20

Table C.1: Adjusting the agents' starting resources and the resulting simulation types. The left hand column corresponds to the mean starting resources for each agent upon instantiation. The number of type 1 simulations (when the agent population collapsed) is shown in the middle column, and the number of type 2 simulations is shown in the right hand column (when the population survived). The data show that as the mean starting resources was increased, the agent population collapsed fewer time.

### C.6.3 Higher Values of Starting Resources

When agents start the simulations with resource stocks exceeding 200 units of each resource, property rights emerge but this happens more slowly than in the default simulations.

Recall from chapters 6 and 10 we observed six 'patterns' at work (see Section 6.4.1, p. 198). These patterns were seen in these simulations also; however, the 'selection' of passive-aggressive agents was delayed because the agents had more resources. As a result, the fifth pattern, which occurred when a single agent had a positive propensity to steal and the other agents' propensities were negative, was more prominent.

Fig. C.9 below is typical of what we observe in simulations when the agents started each round with 1,000 units of each resource. In the simulation shown, property rights did not emerge until Round 1,054. This is much slower than in the default simulations.

Let us examine the fifth pattern in more detail as it mostly explains what we observed in these simulations. Consider a situation in which one agent has a low but positive propensity to steal (say, 0.1) and all the other agents' propensities are below zero (all of them 'passive-aggressive'). In this situation, the 'lone wolf' agent can find that it benefits more from stealing than transacting and its propensity to steal increases as a result.

The reason for this is that the resource concentration effect has a disproportional and beneficial impact on the lone wolf. Suppose it successfully steals from another agent,




Figure C.9: The cloud of the agents' propensities to steal over the first 1,400 rounds of a typical simulation when agents were instantiated with a mean of 1,000 units of each resource. Each blue dot represents a living agent's propensity to steal in each round. It took approximately 1,200 rounds for all the agents' propensities to decline to below zero.

and then other agents attempt to passive-aggressively steal from it<sup>9</sup>. There is a small concentration of resources but here it is 'centred' around the lone wolf in a way that means there is approximately a 50% chance of it being left with the concentrated resources by the end of the simulation. The data shows that, over multiple rounds, the benefits of this tend to outweigh the additional fight costs: in net terms the agent learns it is advantageous to steal so its propensity to steal increases.

If this occurs, the lone wolf's interaction counterparts learn this lesson indirectly and their propensities rise on average. In this scenario we often observed several agents' propensities rise above zero. Moreover, this mechanism tends to enable a few agents to benefit in the same way (the benefits can outweigh the costs for approximately 3-4 agents). In these simulations we often observed an 'outbreak' of several agents' propensities to steal suddenly increasing after the population's propensities had compressed close to zero.

However, as more agents learn it is preferable to steal, this mechanism benefits agents less because the profit from the resource concentration effect is spread across multiple agents. Eventually, fight costs outweigh the benefits of stealing.

We might ask why the system does not settle into an equilibrium whereby propensities to steal stabilise at a positive but low value. This would be where the downward impact of fight costs and transactions equal the upward impetus of the resource concentration effect. The answer is due to the stochastic and complex nature of the whole system

<sup>9</sup>Note that agents often interact 20-30 times in each round.

which ensures the overall behaviour of the agents and the changes to their propensities is much ‘noisier’ than in equivalent equilibrium-based frameworks.

We might also ask how, even in light of these upward and downward ‘forces’, all the (surviving) agents’ propensities to steal eventually decline to below zero. There are two points worth noting. First, fight costs will eventually see the agents’ resource stocks decline toward zero (roughly evenly) and in any one round it is more likely that agents with positive propensities to steal will die first. This can give rise to all the agents with positive propensities to steal dying off, leaving only agents who respect property rights.

We see this in Fig. C.10 below, which shows the total stock of resources (Resource 1 + Resource 2) for the agents depicted in Fig. C.9 above. In this simulation, 20 agents died in various ‘outbreaks’. The last agent that died did so in Round 1,045. Property rights emerged almost immediately after this.

This tendency for agents with positive propensities to steal to die earlier explains why property rights emerged more quickly in the default simulations. Agents had fewer resources at the start of each round so these agents died earlier on, leaving agents with lower and negative propensities. The ‘selection’ of passive-aggressive agents occurs more quickly.

The second point to note is the combination of stochasticity and a natural asymmetry whereby property rights get locked in as soon as all the agents’ propensities decline below zero. This is a form of symmetry breaking: it is irreversible provided that children are not born with positive propensities to steal.

## C.7 Changing the Nature of Agent Learning

In Section 9.4 above we looked at how the agents’ propensities to steal and defend changed as a result of interaction. This involved transforming the reduced values of the agents’ gains / losses in three ways:

- The use of expected reduced gains / losses as a comparator for actual reduced gains / losses when adjusting the agents’ propensities to defend in scenarios 2 and 3 (but the use of absolute values for adjusting propensities to steal);
- A ‘cognitive coarseness’ parameter ( $\delta$ ) that was an exponential term that could range from 0 to 1 (the default value was 0.5); and
- A rate of change coefficient ( $r$ ) that mapped the resulting values on to propensity changes (the default value was 0.01).

See Table 9.1 on page 264 for the resulting equations used in the default simulations.




Figure C.10: A time series of the agents' total resources over 2,000 rounds (taken from the same simulation as that depicted in Fig. C.9 above). The y-axis shows for each round the sum total of resources held by each agent (Resource 1 + Resource 2). The chart shows how all but 5 of the agents died but when property rights emerged, these surviving agents thrived.

In the next three sub-sections we will look at the three factors listed above, in reverse order.

### C.7.1 Rate of Change Coefficient ( $r$ )

This metric can be viewed as determining the speed at which the agents' adjusted reduced gains / losses from an interaction impacted their propensities.

The experiments discussed below showed that the default simulation results were replicated when, approximately,  $r > 0.004$ .

#### C.7.1.1 Low Values of $r$ ( $\leq 0.004$ )

In extremis, if this value was set at 0, the agents' propensities never changed. When we ran 20 simulations using this value of  $r$ , the agent population always collapsed to 1-2 agents. It was not possible for respect for property to emerge because the agents' propensities never changed.

For positive but low values of  $r$  ( $0 < r \leq 0.004$ ), the results were approximately the same: the agents simply did not learn quickly enough for property rights to emerge. Most notably, their propensities to defend increased relatively slowly, which meant their propensities to steal spent longer above 0.5, resulting in higher fight costs than in the default simulations.

### C.7.1.2 Higher Values of $r$ ( $> 0.004$ )

For these values of  $r$ , the results of the default simulations were replicated. The agents went through the same learning process, whereby their propensities to defend rose above 1 and their propensities to steal declined below 0. The higher  $r$  was, the quicker this learning was. For example, when  $r = 0.2$  (20 times its default value), the living agents' mean propensity to defend increased above 1 within 10 rounds (versus 39 rounds in the default simulations), and the mean propensity to steal declined below zero within 13 rounds (versus 233 rounds in the default simulations).

## C.7.2 Cognitive Coarseness ( $\delta$ )

As mentioned in Section 9.4, a value of  $\delta = 0$  meant that the agents' reduced values were transformed to +1 in the case of a beneficial interaction (of any magnitude) and -1 in the case of a detrimental interaction<sup>10</sup>. We can think of this as 'coarse-graining' in the agents' perceptions of the interaction results (Gell-Mann and Hartle, 2007).

By contrast, when  $\delta = 1$  the agents' precise reduced values were used to adjust their propensities without any manipulation. We can think of this as a 'fine-grained' approach to agents' perceptions of the interaction results.

Simulations were run for values of  $\delta$  between 0 and 1 in increments of 0.1. Below, the data is presented and analysed for simulations at the two ends of this spectrum, when  $\delta = 0$  and  $\delta = 1$ .

The main conclusion we can draw from these simulations is that  $\delta$  had to be less than or equal to 0.9 for the default simulation results to be replicated, i.e., for property rights to emerge in all 20 simulations. This means that at least a small degree of cognitive coarseness was required for property rights emergence to be guaranteed.

### C.7.2.1 Coarse-Grained Cognition ( $\delta = 0$ )

The main effect of setting  $\delta = 0$  was that it magnified the impact of transactions (Quadrant 1) on the changes to agents' propensities to steal (putting more downward pressure on them). In the default simulations, transactions typically led to a very small decline in propensities to steal (approximately -0.001). When  $\delta = 0$ , this was typically ten times greater: it accelerated the emergence of respect for property rights because agents 'felt' they benefited more from transactions. In fact, property rights typically emerged before the agents' propensities to defend had reached 1.

To quantify this, the mean of the agents' propensity to steal declined to below zero in Round 50 on average (over 20 simulations), which compares with Round 233 in the

<sup>10</sup>For example,  $v_i^* = \text{sign}(v_i) \times |v_i|^\delta = +1$  when  $v_i > 0$  and vice versa, and  $v_{ii}^{**} = \text{sign}(v_i - E_i^{Q2}(v_i)) \times |v_i - E_i^{Q2}(v_i)|^\delta = +1$  when  $v_i - E_i^{Q2}(v_i) > 0$  and vice versa.



default simulations. Fig. C.11 below shows the ‘cloud’ of the agents propensities to steal over the first 80 rounds of a typical simulation when  $\delta = 0$ .




Figure C.11: The cloud of the agents’ propensities to steal over the first 80 rounds of a typical simulation when agents’ cognition was ‘coarse-grained’ ( $\delta = 0$ ). This meant agents attached  $\pm 1$  values to beneficial and detrimental interactions, respectively. Each blue dot represents a living agent’s propensity to steal in each round and the blue line is the mean of these propensities. All of the agents survived in this simulation and they came to respect property rights much more quickly than in the default simulations.

### C.7.2.2 Fine-Grained Cognition ( $\delta = 1$ )

By increasing the value of  $\delta$  from 0.5 to 1, agents became more sensitive to ‘windfall’ gains and losses (observed when the resource concentration effect was prominent) and less sensitive to gains from trade. In general, we would expect changes in the agents’ propensities to be more volatile and that it would be more difficult for property rights to emerge given the reduced dampening effect on propensities to steal from trading.

The results of 20 simulations showed that the agents’ propensities to defend increased as it did in the default simulations; however, there was much greater volatility seen in the agents’ propensities to steal. Fig. C.12 below shows the ‘cloud’ of the living agents’ propensities to steal over the first 160 rounds of a typical simulation when  $\delta = 1$ .

The cloud of propensities shown in Fig. C.12 displays much greater dispersion than in the default simulations. This resulted from the ‘windfall’ volatility noted above. In general this meant there were more fights, higher fight costs, and the agent population tended to decline more than in the default simulations.

However, the pressure on the agents’ propensities to steal was still net negative, and agents with higher propensities tended to die. In all, the mean total population of agents (over 20 simulations) bottomed out at approximately 7.9 agents when  $\delta = 1$  which




Figure C.12: The cloud of the agents' propensities to steal over the first 160 rounds of a typical simulation when agents' cognition was 'fine-grained' ( $\delta = 1$ ). Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. In this chart the agents' propensities declined more reluctantly than when their cognition was 'coarse-grained'. Nonetheless, property rights eventually emerged in this simulation.

compares with 11.3 in the default simulations. In addition, in 2 of the 20 simulations run for this parameter test, the total agent population collapsed to 1-2 agents.

Simulations run for lower values of  $\delta$  indicated that only a small reduction from 1 was needed for the default simulation results to be replicated in all 20 simulations in a set. For this to be true,  $\delta \leq 0.9$ .

### C.7.3 Relative and Absolute Approaches

As mentioned above, when it came to changes in propensities to defend (in scenarios 2 and 3), the agents' reduced values were compared with the expected gains / losses in the relevant scenario. By contrast when it came to changes in propensities to steal, an absolute value was used (this was discussed in Section 9.4).

What happened if we used absolute value of the agents' reduced gains / losses for changing propensities to defend? and what happened if we contrasted the agents' gain / loss with an expected gain / loss for the whole interaction when adjusting the agents' propensities to steal? We consider these questions in the next two segments.

### C.7.3.1 Using Absolute Reduced Values When Adjusting the Propensities to Defend

When absolute values of the agents' reduced gains / losses were used to adjust propensities to defend, these propensities tended to plateau at approximately 0.73 on average. Fig. C.13 below shows these propensities for all living agents in the first 100 rounds of a typical simulation. As a result of these propensities not moving higher, the agents' propensities to steal seemed to oscillate between 0.5 and 0.6<sup>11</sup>. Fig. C.14 below shows the living agents' propensities to steal in the same simulation as that shown in Fig. C.13.




Figure C.13: The cloud of the agents' propensities to defend over the first 100 rounds in a typical simulation when the agents used absolute gains / losses to adjust these propensities. Each red dot represents a living agent's propensity to defend in each round and the red line is the mean of these propensities. The agents' propensities increased as in the default simulations but they tended to remain below 1.

An analysis of the data showed that when expectations of pay-offs were incorporated in to the adjustment of the agents' propensities to defend, this had the effect of *increasing* the agents' propensities than if absolute values were used. This was most important when resources were concentrated. In fact, agents got stuck in this 'phase' of the simulations, when resources were concentrated, because defence of property did not emerge.

To understand this better, the results from an example interaction is shown below (Table C.2), which shows the outcomes of interaction when Agent  $i$  had no resources and Agent  $j$  had 10 units in total (assuming  $i$  attempted theft and  $j$  wanted to trade). This example was typical of interactions in the resource concentration phase.

Most notably, when an agent with no (or few) resources attempted to steal from an agent with lots of resources, the latter expected to lose significantly from the interaction

<sup>11</sup>Recall that agents' propensities to steal tended to decline when the agents' propensities to defend exceeded approximately 0.8 on average.




Figure C.14: The cloud of the agents' propensities to steal over the first 100 rounds of the same simulation depicted in Fig. C.13. Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. Here, because defence of property did not emerge, the agents' propensities to steal did not decline as they did in the default simulations.

Scenario	$v_j - E_j^{Q3}(v_j)$	$\rightarrow \Delta P_i^{FB}$	$v_j$	$\rightarrow \Delta P_i^{FB}$
3F - Counterpart wins	+7.2	+0.027	-0.6	-0.008
- Counterpart loses	-2.8	-0.017	-10.6	-0.033
3A - Counterpart acquiesces	-2.2	+0.015	-10.0	+0.032

Table C.2: An example of changes to a counterpart agent's propensity to defend using (i) relative outcomes ( $v_j - E_j^{Q3}(v_j)$ ); and (ii) absolute outcomes ( $v_j$ ), when the agent defended its resources (scenarios 3F, win or lose), and when it acquiesced (scenario 3A). The initiating agent ( $i$ ) had no resources and the counterpart ( $j$ ) had 10 units in total, and we assumed  $P_i^{FB} = 0.5$ . The most significant difference between the two was that  $j$ 's propensity to defend *increased* when it won the fight and the relative outcome was employed but it *decreased* when the absolute outcome was used.

( $E_j^{Q3}(v_j) = -7.8$  in the example). In Quadrant 3 this meant that if  $j$  defended its resources and won, it was (in effect) 'relieved' and  $P_i^{FB}$  increased despite the agents incurring fight costs.

In addition, if the agent fought and lost, its losses relative to expectations were much less than its gross losses, so  $P_i^{FB}$  fell less than it otherwise would have. On average, propensities to defend increased in such situations when agents defended their resources (when expectations were incorporated) whereas they fell if absolute reduced gains / losses were used.

The counter-balance to this was that when  $j$  acquiesced, relative losses were again less than gross losses, hence the counterpart increased in  $P_i^{FB}$  was less when expectations were incorporated. When we switched to using absolute values of  $v_j$ , this had the effect

of increasing the agents' propensities to defend. However, this effect undermined itself when  $P_i^{FB}$  increased (fewer agents acquiesced).

In the simulations when absolute values of  $v_j$  were used without expectations, the combination of a more aggressive increase in  $P^{FB}$  due to agents acquiescing and an average decline in propensities when agents defended their resources meant that the agents' propensities to defend plateaued at approximately 0.73, as noted above.

We can summarise the results presented here by stating that a relative measure of gains / losses (versus expected gains / losses) was necessary for property rights to emerge.

### C.7.3.2 Using Relative Reduced Values When Adjusting the Propensities to Steal

As mentioned previously, in the default simulations the agents' propensities to steal were adjusted using absolute gains / losses. Here we deduct the agents' expectations of their gains / losses for each interaction, and adjust their propensities to steal by this relative measure.

The agents' propensities to defend increased in a way similar to that seen in the default simulations but plateaued at approximately 0.8, similar to those seen in Fig. C.13 above. Fig. C.15 below shows the 'cloud' of the living agents' propensities to steal over the first 100 rounds of a typical simulation: in these simulations the agents generally learned it was preferable to steal.

An analysis of the data showed that when resources were concentrated, agents' expectations of their gains / losses from an interaction incorporated expected fight costs. As a result, in these simulations, changes to agents' propensities to steal became desensitized to these costs, which was most significant when both agents attempted to steal (Scenario 4). This had an impact that was equivalent to simulations fight costs being too low.

To quantify this, we observed that the mean aggregated impact of Scenario 4 interactions on the agents' propensities to steal in the default simulations was -26.6 per simulation, which compares with -5.9 in this parameter test. This reduced dampening effect meant that agents' propensities to steal remained high.

The lesson we can take from this parameter test is that for property rights to emerge, agents had to be sensitive to the detrimental impact of fight costs. The simulations described above reduced this sensitivity significantly because fight costs were incorporated in to the agents' expected gains / losses in each interaction.




Figure C.15: The cloud of the agents' propensities to steal over the first 100 rounds of a typical situation when these propensities were adjusted by outcomes relative to expectations. Each blue dot represents a living agent's propensity to steal in each round and the blue line is the mean of these propensities. Property rights did not emerge in this simulation. In fact, agents learned it was generally preferable to steal than to trade.

## C.8 Adjusting the Reputations Architecture

Recall from Section 9.4 that agents were given the ability to store information about other agents' interaction histories in order to estimate the propensities to steal and defend of potential counterparts. This included information from the agent's own experiences of interacting with other agents and what had been learned from others.

In the following sub-sections we adjust this reputations architecture in three ways. First, we assume agents employ substantive rationality if they have no historical information about a potential counterpart, i.e., if they are strangers; second, agents assume the very worst of strangers by prudently assuming their propensities to defend and steal are both 1; and, third, we adjust the agents' memory lengths for retaining any information about other agents.

### C.8.1 Strangers: Use of Substantive Rationality

In the default simulations, if an interaction counterpart was a stranger, the instigating agent assumed the stranger's propensities to steal and defend were both 0.5.

An alternative approach in this scenario was to employ substantive rationality (still using propensities when the counterpart was not a stranger). We made this assumption in a set of 20 simulations.

This approach served to inflate the *perceived* propensities to steal and defend of other agents during the first interactions (because decisions using substantive rationality almost always led agents to steal and defend). However, interactions in the first 100 rounds or so were typically dominated by the resource concentration effect, when the reputations of other agents were largely irrelevant<sup>12</sup>. Furthermore, the length of the agents' reputations array was 20 rounds, which meant any impact of a first interaction was soon forgotten.

Overall, therefore, the effect of this use of substantive rationality for strangers on the results was immaterial: it was mostly employed in the first interaction between two agents but it became irrelevant and was soon forgotten anyway.

### C.8.2 Strangers: Assuming the Worst

Here, if a potential counterpart was a stranger to an initiating agent, it prudently assumed the counterpart would always steal and defend. If the two agents did interact, their actual propensities were used to determine behaviour.

The overall effect on the results was marginal for the same reasons as stated in the preceding sub-section: this prudence assumption was typically used in the first interactions but it soon became irrelevant and was forgotten. The results of these simulations were, therefore, essentially the same as those that used the default parameter set.

### C.8.3 Adjusting the Memory Length of the Agents' Reputations Arrays

In the default simulations, agents could store information about other agents' interaction histories for the previous 20 rounds. In this section we look at the impact of varying this memory length.

In the first set of simulations below we look at results of simulations when the memory length was zero: here, the agents always assumed the propensities to steal and defend of potential counterparts was 0.5. In the second set we look at simulations when the agents had short memories (1-4 rounds).

We can summarise the results of these simulations by stating that the default simulation results were replicated when the agents' memory length was above 3 rounds.

---

<sup>12</sup>If an agent held a large number of resources, other agents would seek to interact with them irrespective of their perceptions of this 'wealthy' agent's propensities.

### **C.8.3.1 No Memories**

The data generated over 20 simulations showed that agents came to defend their resources and respect property rights in a way similar to the default simulations in seventeen simulations; and the population collapsed in the remaining three.

An analysis of the data indicated that the main consequence of agents assuming all their potential counterparties had propensities to steal and defend of 0.5 was that the agents consistently overestimated the likelihood that other agents would acquiesce in scenarios 2 and 3. In general this meant on average the agents' estimated gains from interactions were higher than they ought to have been.

As a result of this, more interactions took place in the earlier rounds than in the default simulations and agents with resources were more likely to interact.

We observed a slightly wider spread in the agents' propensities to defend and steal than in the default simulations. However, the social environment was not so much harsher that the agents population always collapsed: approximately 2 more agent died on average than in the default simulations. As a consequence, the agents came to respect property rights within 1,000 rounds in 17 of the 20 simulations.

### **C.8.3.2 Short Memory Lengths (1 - 4 Rounds)**

In the simulations run for these parameter tests we saw the same phenomenon as when agents' memory lengths were zero (a wider dispersion of propensities) but for longer memories, this was less significant. Simply put, the longer the agents could remember, the more information they stored about potential counterparts and the more accurate their propensity estimates became.

When the agents' memory lengths were four or more rounds, these estimates were sufficiently accurate that the results of the default simulations were replicated.

## **C.9 Changing the Agents' Target Location Weights**

In the default simulations, the agents stored information about the grid locations where they (and other agents) transacted and fought, which was used to select grid targets at the beginning of the interaction phase of each round. If the agent benefited from a transaction or from theft on a particular grid square, the weight attributed to that square was increased by 1. If they lost out, the weight was decreased by 1. For interactions



they heard about from other agents, the weights were increased by  $1/2$  or decreased by  $1/2$ , respectively<sup>13</sup>.

This approach was somewhat crude because these weights were not sensitive to the *degree* of benefit or loss from an interaction. One way of incorporating such sensitivity was to use the agents' reduced benefit / loss from each interaction to incrementally change the weight of each grid square. We took that approach for the simulations run for this section.

The main impact of this was to reduce the (positive) weights of transactions because these reduced values were typically very small (below 0.1). However, the overall impact of this was negligible and the results of the simulations were essentially identical to those using the default parameter set.

## C.10 Reducing Agent Clustering & Resource Concentration

It was notable in the results of the default simulations that agents tended to cluster in approximately the same part of the grid even when they repeatedly experienced losing the resources they had collected during the foraging phase. This raises questions of whether such congregation is reasonable and whether measures to reduce this would have an impact on the results.

In this section we employ two methods to reduce this clustering of agents. First, we enhance the impact of detrimental interactions by both decreasing the weight attributed to the locations where they took place (as in the default simulations) and also reducing the weights of all the adjacent ('king's moves') squares (by the same amount). The rationale is that agents would prefer to avoid getting even close to squares where they had experienced a detrimental interaction<sup>14</sup>.

Second, we allow agents to remove themselves from the grid if they had accumulated a certain amount of resources during the interaction phase. This was done by the agents heading back to their home locations if the reduced value of their basket arrays had increased by at least 2 resource units since the foraging phase.

These two methods are discussed in the next two sub-sections. The third sub-section employs both of them simultaneously.

---

<sup>13</sup>Recall that these weights were depreciated by 20% in between rounds.

<sup>14</sup>Recall that agents can see others on adjacent squares.

### C.10.1 Enhanced Avoidance of Previous Fight Locations

An analysis of the data showed that agents were indeed better at avoiding potential thieves than in the default simulations. They also spent more time moving around the grid randomly because it was more often the case that no grid squares existed in memory with net positive interactions (which was required for a target location to be chosen).

Fig. C.16 below shows a heatmap of fights during the first 100 rounds of a typical simulation, and can be contrasted with Fig. 10.14 in Section 10.2 (a typical simulation that used the default parameter set). It shows that fights were much more dispersed than in the default simulations (we saw the same patterns in the transactions heatmaps), which is consistent with agents moving randomly or targetting a broader range of grid squares during the interaction phase.




Figure C.16: A heatmap of fight locations in first 100 rounds of a typical simulation when agents used an enhanced target selection technique. This allowed them to better avoid locations they associated with their resources being stolen. Fights were much more dispersed in this simulation than in the same period in the default simulations (see Fig. 10.14 as a contrast).

Furthermore, if we look at the gini coefficient of agents' resource holdings at the end of the interaction phase, we found this was notably lower than in the default simulations. In those simulations, the resource concentration effect meant the gini coefficient typically peaked at approximately 0.8 (see Fig. 10.13). By contrast, in the simulations run for this parameter test, the coefficient typically peaked at approximately 0.6.

Despite this 'thinning out' of agent interactions and a reduced resource concentration effect, the agents interacted and learned, which meant their propensities changed. This process took longer than in the default simulations, e.g., it took an average of 148 rounds

for the agents' mean propensity to defend to increase above 1 (versus 37 rounds in the default simulations).

In the 20 simulations run for this parameter test, the agent population never collapsed - the surviving agents always came to respect each other's property quickly enough for such a collapse to be averted.

In summary, therefore, we can state that the default simulations results were largely replicated and property rights emerged, albeit more slowly than in the default simulations.

### C.10.2 Allowing Agents to Go Home

As mentioned above, in the simulations discussed here, agents were allowed to return home (which meant they were removed from the grid) if they had gained at least two more resources during the interaction phase<sup>15</sup> (this would require them to steal). Note that agents attempted to return home: other agents could steal from them on their return journey<sup>16</sup>.

The results of this parameter test were broadly similar to those described in the previous sub-section in that the resource concentration effect was reduced; however, here agents tended to congregate as they did in the default simulations.

Fig. C.17 below shows time series for the gini coefficient of the agents' resource holdings at the beginning and end of the interaction phase of each round (averaged over all 20 simulations). It shows that the coefficient peaked at 0.64, which compares with 0.8 in the defaults simulations (Fig. 10.13).

Despite the reduced resource concentration effect, the agents came to respect each other's property in all 20 simulations, albeit slightly slower than in the default simulations (and the agent population never collapsed).

### C.10.3 Enhanced Avoidance of Previous Fight Locations and Allowing Agents to Go Home

Here we combine the two methods used above. Note their effects had no discernable impact on each other: they were cumulative.

The overall impact was to further slow down the learning processes of the agents. Fig. C.18 below shows time series of the gini coefficients of agent resource holdings at the

---

<sup>15</sup>Strictly speaking, the reduced value of the resources in their baskets had to increase by more than 2.

<sup>16</sup>If this theft was successful, the agent remained on the grid and moved toward its original target location.




Figure C.17: Time series of the gini coefficient of agents' resources holdings (mean of 20 simulations) over the first 600 rounds when the agents removed themselves from the grid if they had gained more than 2 resource units during the interaction phase. In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. This chart shows that resources became more concentrated in approximately the first 250 rounds (when the dark blue line was higher than the light blue line) but less so than in the default simulations.

beginning and end of the interaction phases of each round (averaged over 20 simulations). This metric is a useful proxy for the speed of learning among the agents: it shows the end-of-interaction coefficient peaked at approximately 0.5, which compares with approximately 0.8 in the default simulations.

An analysis of the data showed there was essentially a race (in approximately the first 1,000 rounds) between the agents learning to defend their resources (and respect others' property) and the fact the total agent population was higher than the (non-specialised) carrying capacity of the environment. In eighteen of the 20 simulations, the agent population learned quickly enough that a population collapse was averted. However, in two of the simulations they did not and the total population fell to 5 agents with low but positive propensities to steal<sup>17</sup>.

<sup>17</sup>Again, the population would recover if the simulations were run for long enough because the population was less than the (non-specialised) carrying capacity of the environment. The resulting population would then eventually come to respect others' property.




Figure C.18: Time series of the gini coefficient of agents' resources holdings (mean of 20 simulations) over the first 600 rounds when the agents removed themselves from the grid if they had gained more than 2 resource units during the interaction phase and when they used an enhanced target selection technique. In each round the gini coefficient was recorded at the beginning of the interaction phase and then at the end. This chart shows that resources became more concentrated in approximately the first 400 rounds (when the dark blue line was higher than the light blue line) but less so than in the default simulations and in Fig. C.17 above.



# Appendix D

## Property Rights Model: Experimentation

In this appendix and Chapter 12 we focus on power and legal rules. The latter represents the main motivation for this appendix but power is interesting in and of itself.

Below we describe and examine the results from three sets of experiments that were developed using the second model set out in Chapter 9.

The first set of experiments below assumes agents never defend their resources in a conflict, i.e., their propensities to defend are fixed at zero. The second and third sets introduce power in to the model through its impact on the outcome of any conflict.

All three sets of experiments prepare the way for using legal rules to catalyse the emergence of property rights in the Chapter 12. In that chapter we will also consider a fourth experiment that assumes fight costs are too low to facilitate respect for property within the population. We will also consider the impact of corruption on the efficacy of legal rules in all four sets of experiments.

### **Representing Power in the Model**

In the default simulations, we assumed agents had an equal chance of winning any conflict over resources, i.e., they had the same ‘power’. In social systems, however, conflicts are often not symmetric because one party might have more power than the other. This asymmetry can take many forms, e.g., physical differences and fighting skills matter in altercations; in legal disputes, financial power might mean one side can afford better lawyers; and in political disputes, one person might have more influence than others.

For the second and third set of experiments presented in this appendix we change the even-power assumption and look at two different ways in which the outcome of fights

might be decided. In Section D.2 we introduce a new variable for each agent: fighting skill, which evolves depending on the number of fights an agent has participated in. The greater an agent's skill, the more likely it is to win the fight. Two types of experiment are developed: those in which no 'black sheep' children<sup>1</sup> are born; and those in which they are.

In the third set of experiments (Section D.3), we use the agents' personal resource arrays to represent power: the more resources an agent has, the more likely it is to win a fight. This set of experiments is also divided into two: in the first, all agents are born approximately equally vis-a-vis resources (as in the default simulations); and in the second, one of the agents is instantiated with 20 times the resources of each of the other agents.

In all of the experiments in sections D.2 and D.3, we assume the agents know the 'power' their counterpart has in any fight. This information is incorporated into their expectations of the likely gains (losses) from interactions.

## D.1 Yellow Agents

In this experiment all the agents' propensities to defend are fixed at zero, i.e., they all acquiesce in quadrants 2 and 3.

The results of the experiments were consistent across all 20 simulations in the set: the agents' propensities to steal increased to a mean of approximately 0.9 (Fig. D.1 below shows the mean propensity to steal in the first 200 rounds over 20 simulations for all of the living agents). The social environment was, therefore, hawkish and harsh enough that the population always collapsed.

This result should not be surprising given the observation in the default simulations that the agents' propensities to steal tended to increase when the prevailing propensity to defend was less than approximately 0.8.

Fig. D.2 below shows the gross contributions to the agents' propensities to steal by scenario in the same simulation. Both agents in an interaction learned it was preferable to steal when the instigating agent attempted theft and the other acquiesced (Quadrant 3A); and they learned it was better not to steal when both attempted theft (Quadrant 4). The former outweighed the latter so, on balance, the agents' propensities to steal increased.

This experiment presents us with a challenge vis-à-vis legal rules: property rights did not emerge endogenously so might some legal rule encourage the agents to respect each other's property? This is addressed in the first section of Chapter 12.

---

<sup>1</sup>Recall that 'black sheep' are children born with propensities to steal above 0.



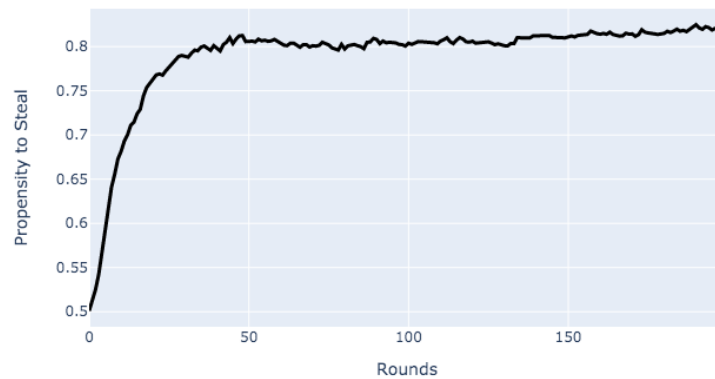


Figure D.1: Time series of living agents' propensities to steal (mean of 20 simulations) over the first 200 rounds when the agents never defended their resources. The chart shows that agents learned it was generally preferable to steal rather than to trade.




Figure D.2: A time series of contributions to the agents' propensities to steal by scenario over the first 200 rounds of a typical simulation within the 'Yellow Agents' experiment. The chart shows agents learned it preferable to steal when the instigating agent attempted to steal and the counterpart acquiesced (Scenario 3A); but they learned it was preferable to trade when both agents tried to steal from each other (Scenario 4). The net effect was positive so the agents' propensities to steal increased.

## D.2 Power from Fighting Skill

For these experiments a new state variable was created for each agent: fighting skill. These skill values were equal to the total number of fights the agent had been involved in over its life (deflated by 1% between rounds). In these experiments, the outcome of

a fight (i.e., which agent won) was determined by the agents' relative skill in fighting.

The following equation was used to calculate each agent's skill at the end of each round:

$$f_{t+1} = (1 - r^f) \cdot f_t + n_i \quad (\text{D.1})$$

where:

$f_{t+1}$  is the fighting skill of an agent in Round  $t + 1$ .

$r^f$  is the deflation variable.

$f_t$  is the fighting skill of the agent in Round  $t$ .

$n_i$  is the number of fights the agent was involved in during Round  $t$ .

We set  $f_0 = 10$ , i.e., we instantiated agents with a fighting skill of 10.

In terms of fight outcomes, if agents  $i$  and  $j$  have fighting skills of  $f_i$  and  $f_j$ , respectively, then the probability  $i$  wins is:

$$\pi_i = \frac{f_i}{f_i + f_j}$$

For this experiment, we increased the agents' initial resources to 400 units of each resource (this counter-balanced the fact the agents fight more in these simulations).

## D.2.1 Simulations without Black Sheep

In this set of simulations we endow the agents with fighting skills, as stated above, and we maintain the assumption that children are born with propensities to steal and defend that are equal to the means of their parents' propensities. This means the children's propensities are constrained to be within the population's minimum and maximum propensities (when born) and they are always born with negative propensities to steal.

We add variation to the children's starting propensities in the next sub-section, to allow 'black sheep' to be born. This means that some of the children's propensities start outside of the population's range and above 0.

Recall from Chapter 10 that we identified two broad groups of agents: 'Al Capone' agents, with relatively high propensities to steal; and 'passive-aggressive' agents, with relatively lower propensities. In general, Al Capone agents fought more in the default simulations: this is important for understand the results of the experiments discussed below.

In general, fighting skills encouraged propensities to steal higher. However, the cost of fighting sustained downward pressure on these propensities after defence of property emerged. The overall results were, therefore, dependent on the degree to which agents' fighting skills depreciated between rounds ( $r^f$  above). When  $r^f$  was higher (approximately  $r^f \geq 0.01$ ), propensities to steal stayed so high for so long that the agent population collapsed most of the time. By contrast, when  $r^f$  was lower (approximately  $r^f < 0.01$ ), the agent population collapsed fewer times.

Looking more closely at the data, property rights did not emerge and the population collapsed in 3 / 20 simulations when  $r^f = 0$ ; this happened in 13 / 20 simulations when  $r^f = 0.01$ ; and 17/20 simulations when  $r^f = 0.04$ .

Fig. D.3 below shows the agents' propensities to steal over the first 500 rounds of a typical simulation when  $r^f = 0.01$  and when the agent population did not collapse. Fig. D.4 shows the corresponding fighting skills in the same simulation. The seven Al Capone agents are noticeable in the second chart: they all died before Round 250 and, eventually, the surviving agents came to respect each other's property.




Figure D.3: Time series of living agents' propensities to steal over the first 500 rounds of a typical simulation when the agents had fighting skills, the change coefficient used ( $r^f$ ) was 0.01, and the agent population did not collapse. Each blue dot represents a living agent's propensity to steal in each round. In this particular simulation the surviving agents came to respect each other's property.

### D.2.2 Simulations with Black Sheep

In this set of simulations we made one adjustment to the model used in the above subsection: now the agents' children were born with propensities to steal and defend drawn from a normal distribution with a mean equal to the mean of their parents' propensities




Figure D.4: A time series of the agents' fighting skills over the first 500 rounds of the simulation depicted in Fig. D.3 above. Each line depicts the fighting skill of a single agent. The chart shows how all the agents' fighting skills increased and that the Al Capone agents with higher fighting skills (and propensities to steal) all died. The surviving agents' fighting skills declined toward zero as their propensities to steal also fell.

and a standard deviation of 3. This enabled the birth of 'black sheep' (children born with positive propensities to steal)<sup>2</sup>.

Note that in the simulations discussed below, we assume  $r^f = 0.01$ .

Recall from the parameter tests in Chapter 11 that agent populations that came to respect each other's property were resilient to these 'black sheep'. Do fighting skills make a difference to this result?

We found that they did: eventually a 'black sheep' would be born that became a single Al Capone agent. It then developed a much higher fighting skill than the other agents and then bullied them until all (or almost all) of them died, i.e., the population collapsed.

Fig. D.5 and D.6 below show the agents' propensities to steal and their fighting skills over 2,000 rounds of a typical simulation when property rights emerged among the initial agents. In this simulation, the sixth child (born in Round 1,174) was instantiated with a propensity to steal of approximately 1.5: in its first few interactions it learned it was preferable to steal and its fighting skill increased rapidly. Eventually, all the agents in this simulation, except the lone 'black sheep', died.

<sup>2</sup>See Section C.5 of Appendix C.




Figure D.5: Time series of living agents' propensities to steal over a typical simulation when the agents had fighting skills and 'black sheep' were born. Each blue dot represents a living agent's propensity to steal in each round. Here, a group of the initial agents came to respect each other's property; however, when the first 'black sheep' was born, it bullied the other agents until they all died.




Figure D.6: A time series of the agents' fighting skills over the simulation depicted in Fig. D.5 above. The fighting skills of the initial agents declined toward zero after property rights emerged. When the first 'black sheep' was born, its fighting skill increased more rapidly than the other agents because it bullied all of them, successfully, to the point of population collapse.

### D.3 Power from Personal Resources

In the section above we introduced a new state variable in to the model (fighting skill) that determined the outcome of fights. In this section we use the agents' own personal resource arrays to determine these outcomes.

Here, if two agents,  $i$  and  $j$ , interact, then the probability that Agent  $i$  wins a fight against Agent  $j$  is:

$$\pi_i = \frac{\sum_h \mathbf{r}_i}{\sum_h \mathbf{r}_i + \sum_h \mathbf{r}_j}$$

where:

$\mathbf{r}_i$  is Agent  $i$ 's personal resource array and  $\mathbf{r}_j$  is Agent  $j$ 's  
 $h$  denotes resources

For example, if  $i$  has 200 resource units in total and  $j$  has 50 units then the probability  $i$  wins a fight is  $200/250$ , or 0.8.

The first sub-section below considers experiments when all the agents were instantiated with two resources (A and B), each drawn from a normal distribution with mean 400<sup>3</sup> and standard deviation of 5. The second sub-section takes exactly the same approach but this time one of the 25 agents is given 8,000 units of each resource, i.e., approximately 20 times that of the other agents<sup>4</sup>.

Furthermore, in both sets of simulations, we assumed that the experiences of each agent's counterpart in any interaction had no effect on the agent's propensities to steal or defend, i.e.,  $\beta = 0$  (see Section C.3 of Appendix C). In these experiments we found that heterogeneity was more significant than in the default simulations (due to a single wealthy agent) so this assumption was more reasonable.

### D.3.1 Equal Starting Resources

At first blush we might expect the agents' propensities to change as they did in the default simulations since the agents started the round with almost identical resources. However, we found this was not the case because of the different strategies adopted by the agents, which led to resource (and therefore power) divergence.

An analysis of the data showed there was a general increase in the agents' propensities to defend and a decline in their propensities to steal as in the default simulations. This was for the same broad reasons: the agents learned it was preferable on the whole to defend their resources than to acquiesce; and to trade rather than steal (when the prevailing propensities to defend were high).

Fig. D.7 below shows the 'cloud' of the agents' propensities to steal over the first 250 rounds of a typical simulation.

<sup>3</sup>This was higher than in the default simulations because agents incurred more fight costs in these simulations and it helps us identify particular mechanisms.

<sup>4</sup>The wealthy agent is prevented from having children so that its 'power' is sustained.




Figure D.7: The cloud of the agents' propensities to steal over the first 250 rounds of a typical simulation when fight outcomes were determined by relative resource holdings. Each blue dot represents a living agent's propensity to steal in each round. Approximately speaking, these simulations replicated the results of the default simulations; however, the passive aggressive strategy was more successful in these simulations and, as a result, more of the initial 25 agents survived and specialised.

In general we found that Al Capone agents incurred greater fight costs than passive aggressive agents, which meant the former had fewer resources and, therefore, less power. This balance of power in favour of passive aggressive agents meant these agents 'preyed' on Al Capone agents who tried to avoid interacting with them. Al Capone agents tended to either die (as they did in the default simulations) or their strategy changed due to their experiences in interacting.

These observations raise the question of how passive aggressive agents can prey on Al Capone agent. This is the passive-aggressive theft noted in Chapter 11. Here, this was augmented by the passive-aggressive agents having more 'power' than the Al Capone agents.

In summary, when power was determined by resources it favoured passive aggressive agents who generally survived, bore children, and specialised. Note that this was the opposite of what we observed when fighting skills determined outcomes (this favoured Al Capone agents).

### D.3.2 One Wealthy Agent

Wealth concentration can occur for a variety of reasons not captured by the simulations presented in this thesis, e.g., successful entrepreneurship or exploiting a natural monopoly. Here we assume that one agent is wealthy for reasons outside of the model.

In these simulations, 24 agents are instantiated with a mean of 400 units of each resource, and one agent is instantiated with 8,000 units of each resource.

The outcome of these simulations was simple: the wealthy agent became an Al Capone agent who was able to bully the other agents. It quickly learned it was preferable to steal and its propensity to steal increased throughout the simulation.

Fig. D.8 below shows the ‘cloud’ of the agents’ propensities to steal over 100 rounds of a typical simulation. The wealthy agent’s propensity is shown as a red line: it increased immediately after the agent began interacting with others. In this simulation, its propensity increased to approximately 13.




Figure D.8: The cloud of the agents’ propensities to steal over the first 100 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 8,000 units of each resource) was included in the population. Each blue dot represents a (non-wealthy) living agent’s propensity to steal in each round. The red line represents the wealthy agent, which adopted an Al Capone strategy in this simulation. This meant it bullied the other agents and the population collapsed.

The bullying of the other agents meant the population collapsed to 2-4 agents who on the whole managed to avoid the wealthy agent.

### Reducing the Wealthy Agent’s Initial Resources

When we reduced the wealthy agent’s initial resources from 8,000 units, we observed a surprising outcome in some of the simulations: the wealthy agent came to respect property rights - its propensity to steal *declined* rather than increased. Fig. D.9 below shows the ‘cloud’ of the agents’ propensities to steal over 1,000 rounds of a typical simulation when the wealthy agent started with 4,000 of each resource (again, its metric is shown as a red line).






Figure D.9: The cloud of the agents' propensities to steal over the first 1,000 rounds of a typical simulation when fight outcomes were determined by relative resource holdings and one wealthy agent (who started with 4,000 units of each resources) was included in the population. Each blue dot represents a (non-wealthy) living agent's propensity to steal in each round. The red line represents the wealthy agent. Here, this wealthy agent adopted a passive aggressive strategy, which meant its propensity to steal became negative quite quickly. Ultimately, property rights emerged across the whole population.

When the wealthy agent came to respect other agents' property, the results of the simulations approximately replicated those of the default simulations: property rights emerged across the whole population.

This raises the question of why the wealthy agent was more likely to become an Al Capone agent when it started simulations with more resources, and vice versa. The answer is related to two different positive feedback effects. With more resources, the wealthy agent learned early on in the simulations it was preferable to steal (encouraged by propensities to defend being relatively low). It then entered a positive feedback loop in which it was more likely to steal and learn this was beneficial.

With fewer resources, it was more likely to enter a 'passive-aggressive feedback loop'. The wealthy agent in these simulations had less power and was more likely to lose resources to other agents, which put more downward pressure on its propensity to steal than otherwise. A relatively lower propensity to steal meant that when it attempted passive-aggressive theft, its propensity fell further, making a repeat more likely, etc.

Now that we have looked at three sets of experiments, in Chapter 12 we use some of these scenarios to conduct our 'liberal legislation' experiments.