Complexity Arbitrage

Using Complexity Science to Improve Outcomes in Asset Management

By Greg Fisher PhD, CFA, MA (Cantab), FRSA

3 August 2025 | v 3.0

Across the world, asset managers are being asked to navigate a far more turbulent landscape than in previous decades. Geopolitical fragmentation, climate instability, accelerated technological disruption, and systemic fragility now define the backdrop of day-to-day decision-making. Taken together, these interlocking disruptions are often referred to as the *polycrisis*: a condition marked not only by simultaneous shocks, but by their entanglement and mutual amplification.[1]

At its root, the polycrisis reflects a deeper intellectual and institutional challenge. Since the industrial and scientific revolutions, much of modern thought – including the foundations of economics and finance – has embraced a mechanistic world-view[2]. This paradigm has driven extraordinary advances in science and living standards. Yet the polycrisis exposes its limits: assumptions of linear causality, equilibrium, and control are poorly suited to today's interconnected, adaptive systems, leading to damaging unintended consequences and, at worst, systemic collapse.

Asset management is no exception. It must navigate these disruptions while embodying many of the same mechanistic habits that contribute to them. Dominant models prioritise optimisation over resilience, prediction over adaptation, and tractability over systemic realism. As stewards of capital, asset managers now face the challenge of operating in a world where these inherited assumptions are increasingly unreliable.

This paper sets out the case for using *complexity science* as a framework to complement and, in some areas, challenge established approaches in asset management. Rather than treating markets as machines, complexity science views them as complex systems[3], which draws attention to interaction, adaptation, feedback, and emergence[4].

The primary aim of this paper is to help asset managers improve valuation, strategy design, and risk management in a turbulent, interconnected world. To that end, we introduce the idea of *complexity arbitrage*: opportunities that arise because prevailing models misrepresent the true dynamics of financial systems. By adopting complexity-informed thinking, practitioners can better identify mispricings, design resilient strategies, and manage risks and uncertainties that conventional frameworks systematically underestimate.

Few firms have yet adopted this perspective in a deliberate way. Notable exceptions include Baillie Gifford and NZS Capital, which explicitly draw on complexity science in their investment philosophies. Their approaches illustrate that practical applications are already emerging, even if detailed methods remain confined to early adopters and specialist practitioners.

Practical guidance is included in this paper but it is treated lightly. Detailed applications are explored in companion papers. The goal here is more to establish conceptual foundations, showing where complexity-informed thinking opens opportunities for complexity arbitrage.

The following sections explore these ideas in depth: outlining the conceptual foundations of complexity science, identifying where conventional models fall short, and sketching the implications for a more robust and adaptive investment practice. Each of the first seven sections summarises a supporting paper, listed below.

- Section 1 **The Case for Complexity in Asset Management** introduces complexity science. It contrasts this with the mechanistic mindset, highlighting concepts such as emergence, feedback, and path dependence[5]. The section argues that complexity provides more realistic foundations for asset management in a turbulent and uncertain world. The counterpart companion paper is *An Introduction to Complexity Science for Asset Managers.*
- Section 2 **What Markets Really Look Like** argues that financial systems should be considered 'complex' given the defining features outlined earlier. It surveys empirical characteristics of asset price behaviour — including fat tails, clustered volatility, and return autocorrelation — that challenge classical assumptions. These *stylised facts* reveal markets to be adaptive, complex systems, highlighting the value of complexity-informed methods. This section mainly draws from the paper *Empirical Characteristics of Asset Price Returns: Stylised Facts for Investors.*
- Section 3 **Patterns, Not Just Models** encourages asset managers to focus on recognising evolving patterns rather than relying exclusively on formal models. Complexity science supports a more contextual, adaptive approach to modelling – one that embraces empirical realism, historical contingency, and the limits of prediction in financial systems. These arguments build on the analysis in *Formal Models and Complexity in Asset Management.*
- Section 4 **Prices as Emergent Phenomena** argues that prices are not static reflections of intrinsic value but emergent outcomes shaped by feedback loops[6], evolving narratives, and interactions among diverse market participants. This re-framing challenges equilibrium-based valuation and redirects attention to the generative processes behind market dynamics. The discussion here is grounded in the paper *Market Prices as Emergent Properties.*
- Section 5 **Narratives as Drivers of Market Dynamics** explores how shared stories influence market outcomes. Drawing on interdisciplinary research, it shows how narratives shape expectations, implicitly coordinate investment behaviour, and drive feedback mechanisms within price formation. This section synthesises material from the paper *Narratives in Financial Markets.*
- Section 6 **Facing Uncertainty** examines risk and uncertainty through the lens of complexity science. It distinguishes probabilistic risk from deeper forms of ambiguity and unknowability, critiques prediction-centric approaches, and advocates for humility and institutional resilience as more robust responses to real-world uncertainty. These ideas are explored in more detail in the supporting paper *Uncertainty and Risk in Finance.*
- Section 7 **Facing Time: Non-Ergodicity in Investment** explores how financial markets are shaped by non-ergodic[7] dynamics, where time-paths matter far more than is generally understood. It shows how compounding, irreversibility, and volatility drag[8] challenge conventional models, and how time-aware frameworks can improve risk assessment and long-term portfolio resilience. This discussion extends the arguments made in *Non-Ergodicity in Asset Management.*
- Section 8 **Toward a New Practice** concludes the paper by situating the discussion within its broader historical and intellectual context. The section revisits the concept of complexity arbitrage, outlines how complexity principles can reshape models, decision-making processes, and organisational behaviour, and highlights the potential system-wide benefits if these ideas were adopted at scale across the industry.

1 The Case for Complexity in Asset Management

This section summarises the main arguments of the accompanying paper *An Introduction to Complexity Science for Asset Managers*. It provides an overview of how complexity science can inform investment thinking and practice, and why it offers a valuable complement to traditional models. We first outline what is meant by a *complex system* and its defining characteristics, then place complexity science in its broader historical context. Next, we introduce its distinctive way of seeing the world through evolving patterns and describe a set of key concepts used to interpret complex phenomena. Finally, we contrast these ideas with traditional financial thinking, highlighting where mechanistic assumptions fall short and how complexity science can support better sense-making and decision-making in financial markets.

Readers interested in an accessible, impressive, and recent discussion of complexity science and its relevance for today's world should read Jean Boulton's book *The Dao of Complexity* (Boulton, 2024).

Defining Complex Systems

There is no single, universally accepted definition of a complex system, which reflects the field's development across multiple disciplines[9]. For instance, physicists may focus on statistical structure and phase transitions whereas biologists often emphasise evolutionary processes and feedback between organisms and their environments. Here, the framing is tailored to economics and finance, where the focus is on decentralised decision-making, networks of professional relationships, and the emergent nature of asset prices.

One helpful way to understand complex systems is that their collective behaviour cannot be understood by analysing their parts in isolation. Outcomes are shaped by dynamic interactions, feedback loops, and continuous adaptation among diverse, interdependent components. Such systems typically lack central control, evolve over time, and defy analysis through traditional linear or equilibrium-based models.

Four key features characterise complex systems:

- **Heterogeneity**: Actors typically differ in goals, preferences, information, and the cognitive rules that influence their behaviour.
- **Interdependence**: Each actor's behaviour affects others, sometimes in non-linear ways. Feedback loops – both reinforcing and balancing – influence the system's overall evolution.
- **Adaptation**: Actors adjust their strategies in response to feedback and changing conditions. These adaptations can reshape the system over time, leading to greater order or dis-order.
- **Emergence**: System-level patterns arise from local interactions among actors, without central control.

For asset managers, understanding these features is essential: market behaviour is not simply the aggregate of isolated decisions but the evolving product of interaction, feedback, and adaptation among heterogeneous participants. Note that we discuss the relevance of these four above points for financial markets in the next section.

Recognising these defining features sets the stage for understanding where complexity science came from and how its cross-disciplinary roots inform a richer, more realistic view of financial markets.

Complexity Science in its Historical Context

To understand how complexity science can inform asset management, it helps to see where the field came from. Complexity science did not emerge as a single discipline with a unified theory. Instead, it grew during the second half of the twentieth century from several interdisciplinary traditions, each reacting to the limitations of *mechanistic thinking* – the dominant paradigm shaping science, economics, and finance since the Enlightenment.

Figure 1 (the Map of the Complexity Sciences, 2021, by Castellani and Gerrits) visualises these origins and highlights influential researchers within each tradition. It shows five main intellectual streams that have converged to shape modern complexity science:

- Dynamical systems theory[10],
- Systems theory[11],
- Complex systems theory[12],
- Cybernetics[13],
- Artificial intelligence[14].

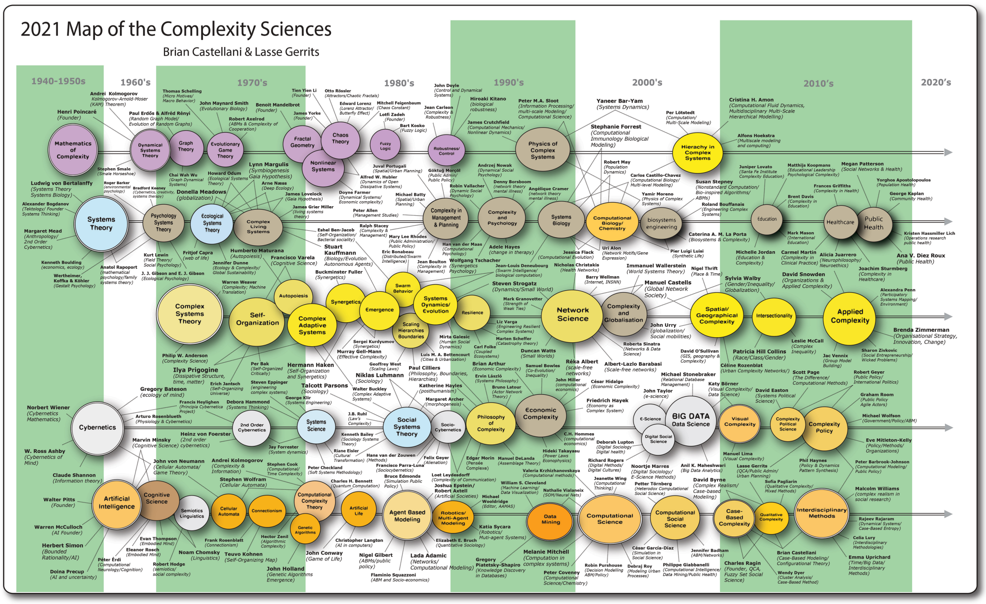

Figure 1: The Map of the Complexity Sciences (2021), developed by Brian Castellani and Lasse Gerrits. Reproduced under a Creative Commons Licence. The diagram organises the field into five major intellectual traditions: dynamical systems theory (purple), systems theory (blue), complex systems theory (yellow), cybernetics (grey), and artificial intelligence (orange). These schools represent the disciplinary origins of what has become a more integrated field of complexity research.

Although each stream developed its own language and methods, they share a common aim, broadly speaking: to understand how collective patterns of behaviour emerge from interacting parts, often without central control. These intellectual roots underpin today's complexity science – a field that spans natural and social systems alike, providing conceptual tools highly relevant for analysing financial markets and the evolving behaviour of asset prices.

This intellectual heritage forms the foundation for applying complexity science to financial markets. The next subsection builds on this, showing how complexity science interprets markets as evolving, patterned systems – a perspective with direct relevance for asset managers.

How Complexity Science Sees the World

It is helpful to understand the kind of worlds that complexity science is reacting to and seeking to describe, building on the four key features of complex systems outlined above. Conventional economics and finance draw heavily on the world-view of classical physics, particularly Newtonian mechanics and thermodynamics. This outlook treats the world as a machine governed by universal, time-invariant laws, where systems are assumed to be predictable, linear, and controllable. Within this "machine metaphor", uncertainty is treated as a problem of incomplete information or knowledge – something that can, in principle, be resolved through improved understanding.

Complexity science challenges this framing. It rejects the notion that economic and social systems operate according to fixed, universal laws but, at the same time, does not see them as chaotic or entirely random. Instead, it proposes a middle ground: a world understood through *patterns* – dynamic, evolving regularities that emerge through interaction, feedback, and adaptation. These patterns are historically contingent, context-dependent, and subject to change.

In economics and finance, such patterns take the form of (for example) dominant technologies, business practices, institutional norms, regulatory frameworks, and legal structures. These typically stabilise for a time before evolving, dissolving, or being replaced. They help shape expectations, behaviours, and outcomes, but they are products of history and social processes rather than fixed laws of nature.

Consider, for example:

- the widespread adoption of the microprocessor, which enabled the internet as well as new forms of trading infrastructure, algorithmic execution, and digital market platforms, fundamentally reshaping how financial markets operate;
- management practices such as just-in-time logistics, which transformed supply chain dynamics and corporate risk exposures, with implications for investment analysis and asset pricing;
- regulatory frameworks and monetary policy regimes – such as inflation targeting or post-crisis macroprudential rules – that (in principle at least) function as stabilising patterns until shifting political, economic, or social conditions prompt their reconfiguration;
- valuation norms, such as discounted cash flow models or risk premia conventions, which emerge from shared professional beliefs rather than immutable laws.

Crucially, these patterns do not arise through central design (except in the case of regulations and legislation) but through decentralised interaction, learning, and feedback. Complexity science highlights this evolutionary nature of economic and financial systems and encourages practitioners to focus on recognising and interpreting these shifting structures, rather than seeking fixed, universal laws of market behaviour.

Some Key Concepts

Having outlined the type of world that complexity science seeks to describe – one marked by openness, interaction, and evolution – we now turn to some of its core concepts. These do not form a single, unified theory but offer a flexible, interdisciplinary toolkit for recognising structure and regularity across complex systems. Many of these ideas have deep intellectual roots, predating complexity science itself. The field provides a coherent framework that connects them and makes them practically useful for interpreting phenomena in domains as diverse as ecosystems, cities, and financial markets.

These concepts can be understood as patterns in their own right – abstract regularities that recur in many settings. For asset managers and economists, they provide alternative lenses for interpreting market behaviour that traditional, linear or equilibrium-based models often struggle to capture.

- **Emergence**: System-level patterns, including behaviours, that arise from interactions among components or actors and cannot be understood by studying these parts in isolation.
- **Feedback**: Interactions that produce reinforcing or balancing loops, shaping system dynamics over time.
- **Adaptation**: Actors adjust strategies in response to experience, feedback, and environmental change.
- **Co-adaptation**: Multiple actors adapt to one another, creating iterative interdependencies.
- **Co-evolution**: Similar to co-adaptation, but occurring at the level of populations that evolve together over time (e.g., predators and prey).
- **Self-organisation**: Order arising spontaneously from local interactions without central control. While often associated with beneficial patterns, harmful or maladaptive patterns can also emerge, e.g., detrimental social norms.
- **Interdependence**: Actors within the system influence and are influenced by each other.
- **Non-linearity**: Small changes can lead to disproportionate and/or unexpected effects; outcomes are not necessarily additive.
- **Path dependence**: When a phenomenon evolves in a manner shaped by its own history.
- **Phase transitions**: Systems can shift rapidly from one regime to another, often unpredictably.
- **Robustness and fragility**: Systems may resist some shocks but remain vulnerable to others due to structural features.
- **Ergodicity**: A system is said to be *ergodic* if, over time, it visits all the possible states – or configurations – available to it in its state space. This implies that the *time average* of an observable property for a single trajectory converges to the *ensemble average*.
- **Mental models**: Internal cognitive representations that actors use to interpret the world and guide decisions.
- **Reflexivity**: Actors' beliefs and models influence the system they are trying to predict or control. Their collective behaviour can lead to a re-patterning of the system and subsequent adjustments in mental models, ad infinitum.
- **Uncertainty**: Some aspects of complex systems are fundamentally unknowable.

Why Traditional Thinking Falls Short

Traditional financial thinking has long been shaped by elegant but overly simplified mathematical models rooted in a Newtonian world-view and thermodynamic analogies. These include Modern Portfolio Theory (MPT)[15], the Capital Asset Pricing Model (CAPM)[16], the Efficient Markets Hypothesis (EMH)[17], and option pricing frameworks such as the Black–Scholes model[18]. While these frameworks have provided structure and tractability, they rest on strong assumptions: fully rational investors, normally distributed returns, ergodic dynamics, stable equilibria, and linear risk–return relationships. These assumptions downplay the complexity, interdependence, and unpredictability of real-world markets.

They also treat uncertainty as stemming solely from incomplete information, implying that once enough is known, asset managers can control for precise outcomes. Financial markets, however, are not closed mechanical systems: they are open, adaptive, and historically contingent. As a result, traditional approaches struggle to explain phenomena such as bubbles, contagion, regime shifts, and feedback loops. They are particularly weak in environments of deep uncertainty, where probabilities are unknown and participant behaviour reshapes the system itself.

The limitations of traditional approaches highlight the need for alternative ways of interpreting and navigating financial markets. Rather than proposing a single predictive framework to supplant traditional models, complexity science offers a set of tools and perspectives for making sense of evolving market dynamics. This shift from seeking precise control to developing adaptive judgement underpins the discussion that follows on decision making in complex contexts.

Sense-Making and Decision Making in Complex Contexts

Building on this shift from control to adaptive judgement, complexity science provides a cognitive toolkit for navigating uncertain and dynamic environments. It offers concepts, heuristics, and frameworks that help practitioners interpret unfolding situations and refine their understanding of market dynamics. The emphasis is not on deterministic solutions, but on inquiry and adaptation – asking questions such as: What patterns are emerging? How confident are we in our interpretation? What actions might clarify the system's structure?

For asset managers, this means moving beyond reliance on static models or predefined playbooks. Decision-making under complexity requires continuously updating mental models, testing investment theses against incoming information, and staying alert to shifts in market regimes driven by feedback loops and evolving narratives. Complexity thinking supports a more iterative and exploratory investment process, where strategies are adjusted as understanding deepens.

A related and essential capability is pattern recognition, which draws on tacit and learnt knowledge, as well as lived experience. In markets characterised by non-linearity and reflexivity, the ability to detect subtle contextual changes early – before they are visible in traditional metrics – can provide a significant edge. This agility helps avoid oversimplified or delayed reactions and supports more resilient portfolio positioning in a world where certainty is rarely attainable.

2 What Markets Really Look Like

Should financial markets be viewed as complex systems? This is a foundational question for any attempt to apply complexity science in asset management. One way to approach it is to ask whether markets exhibit the four defining characteristics of complex systems introduced in the previous section. We consider these immediately below. A complementary way is to examine whether the empirical properties of asset price returns align with the behaviours typically associated with complex systems. This is the focus of the second part of this section.

- **Heterogeneity**: Financial markets bring together a wide variety of actors – institutional investors, retail traders, central banks, algorithmic funds, market makers, and others – each operating with distinct mandates, information sets, and behavioural rules.
- **Interdependence**: Global finance forms a highly interconnected network, where shocks in one node – such as liquidity withdrawals or policy announcements – propagate rapidly through others. The 2008 financial crisis is a clear illustration of systemic effects driven by these interdependencies.
- **Adaptation**: Market participants adjust their strategies continually in response to new information, changing conditions, and the actions of others, sometimes reshaping market dynamics in the process.
- **Emergent patterns**: As explored further in the supporting paper *Market Prices as Emergent Properties*, market-level phenomena such as volatility regimes, bubbles, and crashes arise from decentralised interaction among participants and cannot be attributed to any single actor or directive.

Taken together, these features suggest that financial markets are better understood as evolving systems shaped by interaction, feedback, and historical contingency, rather than as equilibrium mechanisms processing exogenous information in a linear fashion. In this light, complexity science offers not only a theoretical lens but also a practical and empirically grounded framework for interpreting market behaviour – a perspective that opens up opportunities for what we earlier referred to as *complexity arbitrage.*

Having established that financial markets exhibit the core features of complex systems, we now turn to the empirical question: do patterns in asset price behaviour reinforce this view? The following sub-section reviews recurring statistical regularities – the so-called "stylised facts" – observed across asset classes, time horizons, and geographies. These facts are significant for asset managers in their own right, but they also strengthen the case for viewing markets through the lens of complexity.

Empirical Characteristics of Asset Price Returns

As mentioned in the previous section, traditional frameworks in finance – including MPT, the CAPM, the EMH, and Black–Scholes option pricing – assume that asset returns are normally distributed, serially uncorrelated, and driven solely by exogenous news. Risk is modelled as variance; actors are treated as representative and rational; and markets are expected to move smoothly toward equilibrium. These assumptions provide mathematical tractability and elegant solutions – but they are not supported by empirical evidence.

Decades of research paint a different picture. Across equities, bonds, commodities, currencies, and derivatives – and over horizons ranging from intraday to multi-decade – a consistent set of statistical features appears:

1. **Fat tails**: The distribution of returns is non-normal, exhibiting heavy tails and excess kurtosis[19]. Extreme price moves occur far more frequently than Gaussian models predict, leading to systematic underestimation of tail risk.

2. **Absence of simple return autocorrelations**[20]: Raw returns are typically uncorrelated at short lags in liquid markets, likely reflecting adaptive strategies that eliminate short-term predictability. This pattern is sometimes cited as evidence of market efficiency.

3. **Volatility clustering and long memory**: Large price moves are often followed by further large moves (regardless of direction), while calm periods also cluster together. Volatility is strongly autocorrelated and displays long-range dependence – a sign of feedback, regime-switching, and memory in the system.

4. **Power laws and scale invariance**: The distribution of returns, drawdowns, and other market quantities frequently follow power law[21] relationships, suggesting that market patterns are self-similar across scales[22] – behaviour typical of systems far from (or without any meaningful) equilibrium.

5. **Asymmetries in behaviour**: Market reactions to positive and negative news differ. Volatility tends to rise faster in downturns than it falls in rallies (the "leverage effect"), and correlations often increase during stress periods, highlighting non-linear dependencies and amplification effects.

These empirical features are persistent and widespread, not anomalies or artefacts of specific instruments or time periods. They point to a system marked by heterogeneity, interaction, and adaptation.

Viewed through a complexity lens, each stylised fact reflects system-level outcomes of decentralised interaction. Fat tails emerge from amplification mechanisms, where local shocks propagate through networks of actors. Volatility clustering results from herding, learning, and institutional routines that adapt slowly over time. Power law scaling signals self-organised[23] dynamics without a characteristic scale. Asymmetries arise from reflexive expectations, leverage constraints, and non-linear decision rules under uncertainty.

For asset managers, these realities matter directly. Models built on Gaussian assumptions underprice risk, underestimate the frequency of large losses, and ignore non-linear exposures. Frameworks that neglect memory and adaptation miss regime shifts, volatility states, and correlated positioning. Complexity-aware approaches, by contrast, anchor analysis in the actual properties of markets rather than equilibrium and optimisation. This perspective reframes market patterns not as noise but as signals of structure, history, and interdependence, which helps sharpen conventional investment analysis.

3 Patterns, Not Just Models

Formal models underpin many activities in asset management. They guide asset pricing, portfolio construction, risk management, and strategic allocation. Over recent decades, models based on MPT, equilibrium pricing frameworks, and stochastic calculus have become deeply embedded in both practice and regulation. They offer an impression of precision and predictive control but many of these models rest on assumptions that misalign with the empirical and ontological realities of financial markets.

As discussed earlier, most conventional models are grounded in a *mechanistic ontology*, whereas markets are better understood as complex systems. This mismatch creates two risks for asset managers. First, strategies built on overly simplified models can falter in environments marked by non-linearity, regime shifts, and feedback amplification. Second, the apparent rigour of these models – reinforced by their mathematical sophistication and academic pedigree – can instil overconfidence, discouraging critical reflection on assumptions or alternative perspectives. As a result, poor modelling is not merely a technical flaw; it is a strategic vulnerability.

These issues are not hypothetical. Decades of market experience reveal recurring limitations in conventional approaches – the following list contains a number of overlapping points:

- **Instrumentalism (Black-box Modelling)**: Some models are designed to generate plausible predictions while ignoring whether their internal structure reflects real-world processes. Following Milton Friedman's (flawed) interpretation of *instrumentalism*, such models aim for prediction over explanation. They often fail in out-of-sample contexts, particularly in reflexive markets where modelling itself alters behaviour.
- **Over-reliance on fit**: Statistical fit to historical data can produce models that perform well in-sample yet fail when regimes shift. Over-fitted models capture noise rather than structure, producing fragile forecasts and false confidence – failures repeatedly exposed during market crises.
- **Missing important structure (reduction)**: Simplifying assumptions may omit critical interactions, feedbacks, or institutional constraints. Models that reduce complexity instead of abstracting from it[24] risk missing the drivers of system behaviour, particularly under stress.
- **Mis-specified assumptions (wrong ontology)**: Models based on rational agents, equilibrium convergence, or independent shocks may be mathematically serviceable but conceptually flawed if they misrepresent how markets actually operate.
- **Ignoring tail risk and instability**: Gaussian assumptions and constant-volatility frameworks tend to understate the frequency and impact of extreme events. Before and during the global financial crisis, misplaced faith in such formulaic approaches amplified systemic fragility.
- **Goodhart's law and reflexivity**: Widely used models can change the system they attempt to describe, undermining their own predictive power. As market participants act on model outputs, feedback loops reshape dynamics, requiring constant adaptation.
- **Excess confidence and misuse (model risk)**: Beyond conceptual flaws, practical errors in coding or calibration, combined with poor model governance, can lead to significant financial losses. Even robust models can harm when applied without humility or oversight.

A complexity-informed approach points to a different orientation. Rather than searching for optimal solutions under strong assumptions, it emphasises the search for *robust patterns*: regularities that persist across time, context, and evolving structures. Models become tools to support judgement, learning, and collective decision-making, not engines of precise prediction.

This shift reshapes modelling practices:

- **Start from realism**: Build models that reflect heterogeneity, networked interactions, and adaptive behaviour. For example, distinguish between investor types, learning rules, institutional constraints, or liquidity effects rather than assuming representative agents with limited interaction.
- **Aim for patterns, not just fit**: Capture underlying structure and dynamic tendencies – such as momentum, crowding, or volatility clustering – rather than optimising statistical fit to past data.
- **Use models as aids, not answers**: In complex systems, even well-calibrated models cannot provide definitive forecasts. They can, however, highlight plausible future paths, stress-test scenarios, and point to where the system is fragile or adaptive.
- **Experiment and diversify**: Employ multiple approaches – agent-based simulations, regime-switching models, heuristic stress tests, and narrative mapping – to avoid model monocultures and reduce fragility.
- **Hold models lightly**: Apply models with humility, recognising their limits and provisional nature in uncertain, reflexive environments.

This is not an argument to discard formal models, but to use them more wisely. In complex markets, the most valuable models are not those that promise certainty, but those that enhance understanding, sharpen judgement, and guide adaptive decision-making under uncertainty.

4 Prices as Emergent Phenomena

Traditional asset pricing theory typically assumes that market prices reflect an equilibrium – a state in which they efficiently incorporate all available information. Whether framed in terms of marginal utility (as in classical economics), arbitrage-free pricing (as in modern finance), or simple supply and demand dynamics, prices are usually seen as the outcome of optimisation within a stable, well-defined model. This perspective implies that the "correct" price exists independently of the market process, discoverable and deterministic, even if it fluctuates in response to new information.

In practice, prices behave differently, especially in financial markets. As noted in previous sections, financial systems exhibit the hallmarks of complex systems: heterogeneity, interdependence, adaptation, and emergence. From this perspective, prices arise from decentralised interactions among market participants with differing beliefs, information, and constraints. These interactions generate evolving patterns of trading and valuation. In short, market prices are better understood as emergent properties[25] of a complex system.

This section draws on the supporting paper *Market Prices as Emergent Properties*, which reviews relevant literature and develops a more realistic conceptual framing of price formation. Here, we focus on what it means to treat prices as emergent and the implications of this view for asset management.

Why does this matter? A core task for asset managers is to assess the reasonable value of assets and compare these estimates with prevailing market prices; discrepancies between the two drive buy and sell decisions. Complexity science informs both valuation and price interpretation. This section focuses on the latter – clarifying what market prices truly represent when understood as outcomes of decentralised, interactive processes.

A Shift in Ontology: From Discovery to Construction

To understand the difference between conventional and complexity-informed views of price, it is helpful to contrast their ontologies – their assumptions about what prices *are*.

In conventional asset pricing theory, the market's role is to *discover* the correct price, assumed to be unique, stable, and firmly anchored in fundamentals. Participants are generally modelled as homogeneous, rational actors solving well-defined optimisation problems. Even in more sophisticated approaches, such as rational expectations equilibrium or no-arbitrage pricing, the assumption remains that market participants eventually converge on the "true" valuation, with deviations being temporary.

The complexity perspective challenges this view. It suggests that prices are *constructed* through the ongoing interaction of heterogeneous, adaptive participants. There is no single correct price *ex ante*; rather, prices emerge as evolving settlements shaped by narratives, learning, imitation, liquidity constraints, and institutional structures. Price formation, in this view, is not a solution to a fixed optimisation problem but the by-product of a dynamic, collective process.

Emergence and Emergent Properties

To further unpack this ontological shift, it is useful to clarify what is meant by 'emergent properties'.

Markets often display patterns and regularities that are not deliberately designed or centrally imposed but instead arise organically from the interactions of many participants. This phenomenon, known as *emergence*, is described by the Santa Fe Institute as the "process by which a system of interacting subunits acquires qualitatively new properties that cannot be understood as the simple addition of their individual contributions" (The Santa Fe Institute, 2023). In other words, the system as a whole exhibits behaviours and patterns that differ from, and cannot be reduced to, the properties of its individual components.

This view challenges what Ullmann-Margalit (1978) terms the *artificer bias* – the belief that social outcomes must stem from deliberate planning or top-down design. Instead, markets frequently exhibit *spontaneous order*, a concept explored by Enlightenment thinkers and – more recently – Friedrich Hayek, where structures and coordination emerge without central control.[26]

Two processes are especially relevant for understanding emergence:

- **Upward effects**: Micro-level actions – such as trading decisions, beliefs, and strategies – collectively generate macro-level patterns like price movements, volatility clustering, and institutional norms. Historically labelled 'upward causation', this terminology can be misleading, as emergent properties are not 'caused' in the deterministic sense.[27]
- **Downward effects**: Once macro-level patterns are established, they feed back to influence individual behaviours – for example, prevailing prices, conventions, or norms shaping future trading activity. Traditionally referred to as 'downward causation', but again, this should not be understood as mechanical causality.

Recognising emergence shifts our understanding of market dynamics. Rather than treating prices as intrinsic reflections of fundamental value within a mechanical equilibrium, complexity science invites us to explore how collective behaviours, feedback loops, and adaptive interactions continually construct market prices.

Agent-Based Models of Financial Markets

This subsection offers a concise overview of the literature relevant to the idea of prices as emergent properties that is based on Agent-Based Models (ABMs). For a more detailed treatment of this research landscape, readers are referred to Axtell and Farmer (2025).

ABMs provide a bottom-up framework for studying how financial markets evolve as complex, emergent systems. Unlike traditional equilibrium-based models, ABMs simulate markets as interactions among heterogeneous, adaptive actors, each following behavioural rules or strategies. Prices arise endogenously from these interactions rather than from an imposed equilibrium condition, allowing researchers to investigate how real-world market patterns develop over time.

Early research included the Santa Fe Artificial Stock Market (SFI-ASM) developed in the 1990s (see Palmer et al., 1994; and Arthur et al., 1996). This pioneering model showed that diverse traders, using evolving strategies based on past performance, could generate equilibrium-like states, bubbles, crashes, and volatility clustering without requiring perfect rationality or external shocks. It demonstrated that complex market phenomena can emerge purely from decentralised decision-making.

Subsequent 'second-generation' ABMs focused on reproducing the stylised facts summarised in Section 2, such as heavy-tailed return distributions, volatility clustering, and the absence of autocorrelation in returns. Models developed by Thomas Lux and colleagues, for example, showed that simple architectures with fundamentalist and chartist agents, switching probabilistically between strategies, could mimic real market statistics and fractal-like price patterns. A third generation introduced more realistic order-book dynamics, capturing effects such as boom–bust cycles and persistent oscillations in price formation.

Typical of ABMs is heterogeneity and adaptation: agents differ in beliefs, information, and risk preferences, and update their strategies through learning. These interactions create self-reinforcing feedback loops, regime shifts, and an ever-changing 'ecology of strategies' where no single approach dominates indefinitely. Simplified models, such as minority games, further demonstrated that complex price fluctuations can arise from minimal behavioural rules.

By the 2000s, ABMs had become an established methodology in economics and finance, with applications extending to market microstructure, derivatives, and cryptocurrencies. They offer a powerful lens to examine how local interactions generate global price dynamics, challenging equilibrium-based theories and providing tools to explore how policy, market design, or trading behaviour can drive systemic phenomena.

Information and Decentralized Price Discovery

Here we briefly explore the literature related to how prices in financial markets emerge as mechanisms for aggregating dispersed information held by individual participants. Rooted in Hayek's hypothesis[28], it describes prices as real-time statistics that synthesize private beliefs and signals into a collective valuation.

It is important to note, however, that asset markets differ fundamentally from goods and services markets due to their forward-looking nature, reliance on expectations, and susceptibility to shifting narratives and reflexivity. This temporal and interpretative dimension makes asset prices highly path-dependent and subject to uncertainty beyond calculable risk.

Empirical support for decentralized price discovery comes from laboratory experiments. Smith (1962) showed that competitive prices emerge rapidly even when traders know only their own supply or demand conditions. Plott and Sunder (1988) extended this to asset markets, demonstrating that dispersed signals, when aggregated through trading, allow prices to converge toward the true underlying value. These findings underpin modern prediction markets, where prices serve as collective forecasts.

Gode and Sunder (1993) provided further evidence with their "zero-intelligence trader" experiments, showing that market efficiency can emerge even when participants act randomly under budget constraints. The structure of continuous double auctions itself drives convergence toward efficient allocations, indicating that sophisticated cognition is not strictly necessary for price discovery.

Theoretically, this aligns with the EMH, which views prices as information-reflective. Yet the Grossman and Stiglitz (1980) paradox highlights that perfect efficiency is unattainable: if information were fully embedded in prices, no one would invest in discovering it. Instead, prices reflect the outcome of costly, decentralized efforts to gather and act on information.

Overall, the above literature points to price discovery that is best understood as an emergent, collective computation shaped by interaction, competition, and institutional mechanisms rather than by a central planner or perfectly rational actors. Prices dynamically evolve to encode dispersed information, but their accuracy and stability depend on the diversity, independence, and incentives of market participants.

Implications for Asset Management

Viewing market prices as emergent properties reshapes several aspects of asset management practice. This shift in perspective influences how firms think about investment strategy, risk management, organisational design, and incentives, as outlined below:

- **Interpreting prices as processes, not fixed points**: Prices represent evolving settlements shaped by interaction and adaptation, not unique or deterministic fair values. Portfolio decisions should therefore be framed as contingent on – or with an expectation of – shifting market dynamics.
- **Attention to narratives and implicit coordination**: Price formation is strongly influenced by shared stories, expectations, and conventions that can change abruptly. Monitoring narrative shifts complements traditional analysis of fundamentals (discussed further in the next section).
- **Role of feedback loops**: Trading activity and positioning feed back into price dynamics, sometimes amplifying volatility or creating instability. Understanding these loops is essential for anticipating crowding risks or reflexive sell-offs.
- **Limits of equilibrium valuation**: Valuation models assuming a single discoverable price can mislead. Scenario-based approaches and adaptive valuation frameworks, leading to multiple price points, offer more resilience to changing market conditions.
- **Microstructure awareness**: Prices evolve within specific institutional settings, shaped by order-book dynamics, liquidity constraints, and execution flows. These factors can significantly influence realised prices.
- **Diversity and resilience**: Markets dominated by similar strategies or beliefs become fragile. Encouraging diversity of perspectives within teams and portfolios can help avoid systemic vulnerabilities and improve long-term decision quality. This idea is linked to the concept of requisite variety in complexity science.

5 Narratives as Drivers of Market Dynamics

If prices emerge from interaction and feedback, then the ideas, stories, and framing devices used by market participants become central to collective market behaviour. Among these, narratives play a particularly important role: they shape how actors interpret information, form expectations, and make decisions. Unlike other cognitive patterns such as metaphors or analogies, narratives have a temporal structure that links past events to present signals and future possibilities. They influence not only what market participants believe, but also what they notice, how they react, and how their individual decisions combine into system-wide outcomes. In this sense, narratives function both as cognitive instruments and as powerful forces shaping financial systems.

This section draws on the supporting paper *Narratives in Financial Markets*, which reviews the growing body of research on this topic. It highlights key mechanisms through which narratives influence market behaviour, surveys empirical findings, and sets out why narrative dynamics matter for asset managers. The discussion proceeds in four parts: how narratives help market participants make sense of uncertainty, the main mechanisms through which they act, evidence from empirical studies, and practical implications for investment practice.

Narrative as Sense-Making in Uncertain Systems

Financial markets are inherently uncertain environments. Actors must make decisions without complete information, stable causal relationships, or well-defined probability distributions. Under such conditions, narratives become essential sense-making devices: they offer explanations, provide emotional and cognitive anchoring, and enable more informed decision making.

This contrasts with standard models of decision-making under risk, where preferences are fixed, beliefs are probabilistic, and the environment is fully specified. In a complexity framing, uncertainty is ontological as well as epistemic: actors cannot know all possible future states or assign meaningful probabilities to many outcomes. Narratives, in this context, are how people navigate indeterminacy – simplifying, interpreting, and projecting events to enable action despite pervasive uncertainty.

Key Mechanisms

Researchers have identified several ways in which narratives influence investor behaviour and shape price dynamics. These mechanisms often operate simultaneously, reinforcing one another and contributing to emergent market outcomes:

- **Framing effects**: Narratives guide how investors interpret incoming information. A well-crafted story can draw attention to particular aspects of news while downplaying others, leading the same data to be perceived in bullish or bearish terms. Experimental studies[29] show that investors often give greater weight to narratively framed information than to abstract or purely statistical data.
- **Emotional engagement**: Narratives frequently resonate on an emotional level, triggering responses such as fear or exuberance more strongly than numerical facts alone. As David Tuckett

argues, narratives provide the emotional scaffolding needed to justify action in ambiguous environments (Tuckett, 2017). Strong emotional content can override risk aversion, fuelling episodes of optimism, overconfidence, or panic selling.

- **Contagious spread**: Narratives are socially contagious. Like memes, compelling stories diffuse rapidly through media, social networks, and professional communities. Shiller (2019) documents historical episodes – from tulip mania to Bitcoin – where narrative contagion played a key role in driving market dynamics. Social media has accelerated this process, enabling near-instant transmission of narratives (as seen during the GameStop episode). Importantly, even those who do not believe a narrative may act as if it matters, knowing that others will be influenced by it.
- **Social validation and feedback**: In uncertain environments, investors often look to others for cues. When a narrative becomes widely accepted, it gains authority – not necessarily because it is true, but because many believe it. This herding behaviour can be rational: collective belief can shape market outcomes. For example, widespread confidence in stable credit markets prior to the 2008 crisis reduced perceived risk, fuelling lending and asset inflation. As narratives gain traction, they form shared mental models that guide expectations. Confirmation bias can then entrench the dominant story, filtering out contradictory evidence. In extreme cases, narratives become self-fulfilling – belief drives behaviour that reinforces the narrative, until a competing narrative overturns it.

Taken together, these mechanisms show that narratives are far from "noise" in financial markets. They are powerful forces shaping interpretation, emotion, social coordination, and feedback – all central to how prices form and evolve in complex systems.

Empirical Observations

A growing body of empirical research supports the idea that narratives play a significant role in shaping market behaviour. The supporting paper reviews various methods used to track and analyse narrative dynamics, including natural language processing, media analysis, surveys, interviews, and ethnographic studies.

Key findings from this literature include:

- The rise and fall of dominant macroeconomic narratives (e.g., "secular stagnation," "reflation," "soft landing") often coincide with notable shifts in asset prices, volatility, and capital flows.
- Sudden narrative regime changes – such as the transition from "low inflation forever" to "monetary tightening is necessary" – can trigger sharp repricing across multiple asset classes.
- Narrative alignment across media channels, policymakers, and investors tends to amplify market moves, increasing the risk of herding, overshoot, and abrupt reversals.

These patterns cannot be fully explained by fundamentals or news flows alone. As discussed in the next subsection, narratives arise from within financial markets themselves, emerging from interactions among participants and evolving dynamically over time.

A Complexity Perspective

From a complexity science viewpoint, narratives can be understood as distributed cognitive patterns that both shape and are shaped by market dynamics. They are not exogenous inputs or top-down directives but arise endogenously from decentralised communication among heterogeneous actors – investors, analysts, journalists, policymakers, and others. Narratives spread, evolve, and gain traction through these interactions, influencing how market participants collectively process information and form expectations.

This positions narratives as another form of emergence in financial systems. Like volatility clustering or market bubbles, narratives originate at the micro level – in the interpretations, forecasts, and framings of individual actors – and coalesce into macro-level patterns that exert systemic influence.[30] Once established within a financial environment, narratives feed back into participants' behaviour. Narrative dynamics are therefore recursive: they emerge from interaction and in turn reshape the structure of that interaction. This echoes the upward and downward effects of emergent properties discussed earlier.

Narratives also offer a bridge between qualitative and quantitative approaches. They provide interpretive depth – helping to explain how actors make sense of uncertainty, coordinate action (intentionally or otherwise), and respond to emotionally charged developments – while increasingly being amenable to empirical study. Advances in natural language processing (NLP) and data science now allow narrative dynamics to be examined at scale. Techniques such as topic modelling, sentiment analysis, and semantic clustering can identify dominant storylines, track their evolution, and relate them to market outcomes. As noted in the supporting paper, this opens the door to new forms of pattern recognition: mapping "narrative landscapes" and detecting "narrative regimes" that co-evolve with volatility regimes, liquidity conditions, and investor positioning.

Seen through this lens, narratives are not secondary phenomena in markets. They are core components of the adaptive landscape in which financial actors operate – shaping expectations, influencing positioning, and contributing directly to the feedback structures that drive market outcomes.

Implications for Asset Managers

Narrative dynamics have direct and significant implications for asset management practice. Key considerations include:

- **Narrative understanding as a complementary lens**: Narratives influence sentiment, regime dynamics[31], and price formation. Incorporating narrative analysis adds an interpretive layer that strengthens traditional fundamental, technical, and quantitative approaches.
- **Identification of narrative regimes**: Detecting whether markets are dominated by coherent or fragmented narratives helps explain volatility, directional biases, and pricing signals.
- **Monitoring narrative evolution**: Tracking changes in narrative content, sentiment, and prominence over time can provide early indicators of potential price movements or regime transitions.

- **Cross-market narrative diffusion**: Narratives often spread across asset classes and geographies, generating correlations and co-movements that may not appear in standard risk models. These dynamics can operate at the micro level (e.g., single securities), meso level (e.g., sectors), or macro level (e.g., global themes).
- **Narrative risk**: Abrupt shifts in dominant narratives can trigger sudden repricing, herding behaviour, or liquidity crises, creating material risks for portfolios and counterparties.
- **Hidden concentration risks**: Portfolios that appear diversified may in fact be exposed to a single underlying narrative driver, leading to correlated losses under stress.
- **Reputational and regulatory exposure**: Narratives around themes such as ESG or geopolitics can influence investor sentiment, trigger regulatory scrutiny, or redirect capital flows independently of fundamentals. Asset managers can benefit from anticipating these dynamics before they materially affect portfolios or firm operations.
- **Timing asymmetry**: Narratives may lead or lag market prices, creating complex dynamics between perception, expectation, and realised outcomes.
- **Organisational intelligence**: Narrative analysis fosters cross-team understanding, reduces cognitive silos, and enables more coherent decision-making under uncertainty.
- **Dedicated analytical capabilities**: Depending on resources available, firms may benefit from establishing narrative analytics teams, developing bespoke dashboards, and drawing on interdisciplinary expertise (e.g., psychology, linguistics, data science) to track and interpret narrative landscapes effectively.

Seen this way, narrative competence is an increasingly valuable capability for asset managers operating in complex market environments – yet it remains underdeveloped across much of the industry.

6 Facing Uncertainty

Asset managers operate in a world of uncertainty – but not all uncertainty is alike. A core contribution of complexity science is to clarify the nature of uncertainty and how it differs from calculable risk. This distinction, which has deep roots in economic thought, is not just academic: it has direct implications for modelling, strategy, and judgement in investment practice.

Risk, Uncertainty, and Ambiguity

A longstanding distinction in decision theory – and one highly relevant to asset management – is that between *risk* and *uncertainty*. Frank Knight (1921) and John Maynard Keynes (1937) both argued that not all unknowns can be treated alike. Risk refers to situations where outcomes and their associated probabilities are known. By contrast, Knightian uncertainty describes situations where neither outcomes nor their probabilities can be meaningfully specified – they are not merely hard to estimate, but fundamentally indeterminate. Keynes similarly observed that many economic decisions are made under conditions where "we simply do not know" the relevant probabilities (Keynes, 1937, p. 214).

A third category, *ambiguity*, is also worth highlighting. It refers to situations where the possible outcomes are known, but the probabilities attached to them are unclear or contested. This is common in finance, where one might know an asset's potential price range – even if it stretches from -100% to $+\infty$ – yet have no reliable basis for assigning probabilities to those outcomes.

These distinctions matter because many conventional financial models treat all unknowns as measurable risks, assuming that probabilities can always be specified. A complexity-informed approach recognises that asset managers frequently operate under ambiguity or genuine uncertainty, where probabilistic reasoning breaks down and alternative approaches are required. Table 1 summarises the differences between risk, ambiguity, and uncertainty.

	Outcomes	Probabilities
Risk	Known	Known
Ambiguity	Known	Unknown
Uncertainty	Unknown	Unknown

Table 1: Comparison of Risk, Ambiguity, and Uncertainty

Sources of Uncertainty

Understanding uncertainty involves more than analysing outcomes and probabilities; it also requires attention to its underlying sources. Drawing on literature from institutional theory and complexity science, five distinct sources can be identified. These are not exhaustive, and several may be present in a single context.

Limited information reflects the reality that economic actors rarely have access to all the data needed for decision-making. Gaps relating to counterparties, product quality, or enforcement mechanisms create the need to gather, interpret, and verify information – often at a cost.

Limited knowledge refers to incomplete understanding of the patterns that link phenomena. While data points describe what has happened, knowledge concerns why events occur. All actors rely on mental models to interpret the world, but these are inevitably imperfect. Complexity science can help refine such models, while recognising that complete accuracy is rarely achievable.

Limited human cognition builds on the notion of bounded rationality: our ability to store and process information is finite. Actors cannot optimise across all possible strategies and instead depend on cognitive patterns, including narratives (as discussed earlier), to navigate decision-making.

Beyond these three 'classical' constraints, complexity science highlights two deeper sources of uncertainty:

Mutual contingency arises when an actor's decisions depend on its expectations about others, and vice versa. This interdependence creates an infinite regress problem: even with full information and computational power, strategic uncertainty remains unless stabilised or resolved by norms, conventions, or laws.

Future novelty refers to the emergence of genuinely new phenomena – properties, behaviours, or events that were not foreseeable in advance. This reflects not only epistemic limits (what we do not yet know) but also ontological emergence (what cannot be known even in principle). Such novelty places hard limits on prediction and underscores the need for adaptive rather than purely probabilistic strategies.

Taken together, these five sources illustrate why uncertainty is not merely a matter of missing data. It is a fundamental condition of decision-making in complex systems, arising from the limits of knowledge, cognition, interdependence, and the evolving nature of the world itself.

Behavioural Perspectives and Ambiguity Aversion

Behavioural research shows that investors do not respond to risk, uncertainty, and ambiguity in the same way. A foundational insight comes from the Ellsberg Paradox (Ellsberg, 1961), which demonstrated that people systematically prefer known risks over ambiguous ones – a phenomenon referred to as *ambiguity aversion*. This behaviour is distinct from risk aversion: it reflects discomfort not with volatility itself, but with the lack of reliable information about underlying probabilities.

Building on this insight, economists have developed models to formalise decision-making under ambiguity. One influential framework is the multiple-priors model by Gilboa and Schmeidler (1989), which assumes that investors consider a set of possible probability distributions rather than relying on a single known one.

Empirical evidence further suggests that ambiguity is priced into financial markets. Investors demand an *ambiguity premium* – a higher expected return for bearing ambiguous rather than well-defined risks. Studies such as Bossaerts et al. (2010) and Brenner and Izhakian (2018) confirm that assets exposed to greater ambiguity trade at wider discounts relative to their expected payoffs.

Taken together, these findings reinforce the view that ambiguity and uncertainty are not simply more complicated versions of risk. They are qualitatively different conditions that meaningfully shapes investor behaviour and asset pricing in ways that standard risk-based models fail to capture.

Impact on Asset Pricing and Portfolio Management

Ambiguity and Knightian uncertainty carry significant implications for how assets are priced and portfolios are constructed. Standard asset pricing models assume that investors can assign reliable probabilities to future outcomes but when ambiguity or deep uncertainty is present, these assumptions break down. Theoretical work such as Epstein and Wang (1994) shows that under Knightian uncertainty, asset prices are no longer anchored to unique values but instead fall within a 'plausible range'.

Consistent with this, Haldane (2009) notes that when investors lack a shared probability model, valuations become more conservative and less precise. This often results in systematically lower prices for a given set of expected cash flows – effectively, higher risk premia. The discounting reflects not volatility itself, but a deeper aversion to the unknown. This dynamic was evident during the 2008 financial crisis, when asset managers struggled to assign values to complex securities in the absence of trustworthy models, leading to market freezes and distressed selling.

Portfolio decision-making is also reshaped by under ambiguity / uncertainty. Traditional approaches such as mean-variance optimisation[32] rely on the assumption that future return distributions are known and can be reliably estimated. When this assumption does not hold, ambiguity-averse investors shift their focus toward robustness rather than optimisation. They often accept lower expected returns in exchange for portfolios that are less vulnerable to errors in model assumptions or data inputs. In practice, this can mean favouring familiar or well-understood assets (a form of home bias), applying wider margins of safety to valuations, or avoiding investments where the underlying probabilities are unclear altogether.

These behavioural patterns are difficult to explain using risk-only models. They become intelligible once ambiguity and uncertainty are recognised as first-order concerns. In this way, accounting for uncertainty not only improves descriptive realism but also helps rationalise observed deviations from classical portfolio theory.

Challenges for Risk Management and Decision-Making

Deep uncertainty and ambiguity pose fundamental challenges for conventional risk management practices. Standard tools – such as value-at-risk (VaR)[33], volatility forecasts, and credit models – typically assume that historical data can be used to infer future probabilities. When uncertainty is Knightian in nature, or when ambiguity dominates, this assumption is invalidated. In such cases, probabilistic models may provide a false sense of precision, masking the reality that key risks are unmeasurable or structurally unknowable. This can be dangerous and, in extreme cases, catastrophic for portfolios.

We can distinguish between problems at different conceptual levels. First, there is model risk: the possibility that the chosen risk model is misspecified or rests on flawed assumptions. Second, even with a well-specified model, the relevant probability distribution(s) may be unknown or unknowable. Third – and most radically – there may be no stable distribution at all. This is particularly likely in complex, open systems characterised by novelty, reflexivity, and mutual contingency.

Behavioural research shows that in contexts of deep uncertainty, investors often fall back on heuristics or narratives, as discussed in the previous section. In financial markets, this can manifest as herding (copying others when individual models lack credibility), a preference for liquidity (to preserve flexibility), or a bias toward actions that maintain optionality. While these behaviours often make sense under uncertainty, they diverge from the predictions of classical finance theory – underscoring the need for an approach that acknowledges the reality of unknown unknowns.

Frameworks for Managing Uncertainty in Practice

Faced with such deep uncertainty and ambiguity, some asset managers and researchers have developed practical frameworks that go beyond conventional risk models. These approaches aim to embed robustness into decision-making and portfolio design, recognising that uncertainty cannot always be reduced to measurable probabilities. Four broad strategies stand out:

- **Robust optimisation and min-max strategies** avoid reliance on a single probability distribution. Instead, they assess performance across a range of plausible models, often targeting worst-case outcomes. The objective is to minimise maximum regret or loss, thereby protecting against severe model error. While potentially less efficient under a 'correct' model – assuming one exists – these approaches provide resilience under ambiguity and can help guard portfolios against adverse tail risks.
- **Info-gap decision theory** offers a non-probabilistic framework for severe uncertainty. Rather than asking what is most likely, it asks how far reality can deviate from expectations before a strategy fails. This generates a robustness profile for each decision. For example, a bank might use info-gap analysis to determine which loan portfolio remains solvent across the widest range of adverse conditions. Unlike worst-case strategies, info-gap approaches seek robustness across a broad middle ground of unknowns.
- **Scenario planning and stress testing** are long-standing tools for preparing for diverse future states. Instead of relying on a central forecast, decision-makers examine optimistic, pessimistic, and extreme scenarios – such as recessions, market crashes, or pandemics – and assess portfolio performance under each. These methods encourage thinking beyond statistical baselines and help uncover structural vulnerabilities. Regulatory stress tests and proprietary tools, such as those used by BlackRock, operationalise this approach at scale.
- **Model risk management and adaptive strategies** emphasise recognising model limitations. Some financial institutions maintain model inventories, conduct regular validation exercises, and allocate buffers for model risk. Adaptive techniques, including Bayesian updating and model averaging, reduce reliance on any single model. Even simple heuristics – such as equal-weighting across assets – have been shown to perform well under uncertainty, particularly when estimation error is high and confidence in model forecasts is low.

All four frameworks share a common goal: to construct decisions and portfolios that remain viable across many possible futures, not just the one perceived to be most probable. This often means sacrificing some expected (arithmetic mean) return in exchange for reduced exposure to unknown risks. In both academic and industry contexts, there is evidence of a shift from optimisation towards robustness. As Debelle (2010) observed, designing financial systems with lower leverage and greater capital buffers is one straightforward way to enhance resilience. More broadly, the complexity-aware investor seeks not to eliminate uncertainty, but to respect it – by designing systems capable of withstanding surprise.

7 Facing Time: Non-Ergodicity in Investment

A central but often implicit assumption in many standard financial frameworks is that markets are *ergodic*. This assumption, touched on earlier and defined more precisely in the next subsection, suggests that long-term outcomes for a single investor mirror the mean outcomes observed across many investors at a point in time. In practice, financial markets rarely behave this way. Treating them as ergodic can lead to flawed risk assessments, misaligned strategies, and poor long-term portfolio outcomes. The accompanying paper *Non-Ergodicity in Asset Management* develops this argument in detail, showing how recognising the non-ergodic nature of markets reshapes our understanding of risk, return, and investment strategy over time.

Ergodic and Non-Ergodic Systems

In an *ergodic* system, the long-run performance of a single investment path is assumed to match the average result observed across many possible paths at a given moment. In other words, the time average of one trajectory approximates the 'ensemble average'[34] of all possible states. This simplifying assumption underpins much of orthodox financial theory: expected utility maximisation, mean–variance optimisation, and standard asset pricing models all rely on ensemble averages as a guide for decision-making.

Financial markets, however, do not conform to this structure. Investment returns evolve along a single, path-dependent trajectory where shocks, sequencing effects, and compounding dynamics create outcomes that cannot be replicated or averaged away. Most investment processes are multiplicative: gains build on past gains, and losses erode the base from which future returns compound. For example, a gain of $+50\%$ followed by a loss of -50% (or vice versa) results in a net wealth decline of -25%, not zero. In such non-ergodic settings, long-term investor outcomes often diverge sharply from ensemble averages, meaning strategies based on ergodic assumptions can systematically underestimate the risks to capital over time.

Kauffman's Evolutionary Framing

Stuart Kauffman offers a philosophically grounded view of non-ergodicity with implications that extend from biology to economics. Central to his account is the idea of path dependence: a system's past decisions, actions, and external shocks shape its present state and constrain its future trajectory. These paths cannot be rewound or re-run under different conditions. Financial markets share this property, as past events, market structures, and compounding effects continually influence what outcomes are now possible.

Kauffman introduces the concept of the *adjacent possible* – the set of potential future states reachable from a system's current configuration. Change is incremental and historically contingent: not every future is accessible, and new possibilities emerge as systems evolve. This framing aligns with the notion of open-ended evolution, where economies explore an expanding, indeterminate landscape of opportunities shaped by innovation, adaptation, and interaction.

His work also highlights 'ontological emergence', where new properties or behaviours arise from interaction that cannot be reduced to component parts. For example, the 'wetness' of water is not inherent in any individual H_2O molecule but emerges among a group of molecules, i.e., at a higher level of organisation. This notion extends to social systems, where new emergent properties like social norms are observed over time.

Finally, Kauffman emphasises self-organisation and group selection. Cooperative structures – whether in biology, sports teams, or firms – can outperform individuals operating in isolation, suggesting that evolutionary processes favour coordination as well as competition. These ideas challenge the reductionist assumptions embedded in mainstream economics and support a complexity-aware, non-ergodic perspective on financial markets.

Theoretical Foundations

The concept of ergodicity originates in late 19th- and early 20th-century statistical physics, where it was developed to explain how physical systems evolve over time. Ludwig Boltzmann introduced the idea in the 1870s, and George Birkhoff later formalised it in the 1930s. The ergodic hypothesis proposed that, under certain conditions, time averages for the properties of a single system would equal ensemble averages across all accessible microstates. This assumption enabled powerful simplifications in thermodynamics and statistical mechanics. By the mid-20th century, however, researchers studying biological and complex systems increasingly recognised that many natural phenomena are non-ergodic – shaped by history, exhibiting irreversibility, and generating novel structures. These insights paved the way for later applications in economics and finance.

Post-Keynesian economists, such as Paul Davidson and Sheila Dow[35], argued in the 1990s that economies are fundamentally non-ergodic: they evolve from an irrevocable past toward an uncertain future, and historical data cannot reliably forecast what comes next. Economic processes are path-dependent and do not converge to stable, long-run equilibria. This challenges the ergodic assumptions embedded in much of mainstream theory.

Building on this perspective, Ole Peters and collaborators launched a research programme called *Ergodicity Economics* around 2011. They showed that many puzzles in economics stem from conflating ensemble averages with time averages. Expected utility[36] theory often assumes ergodicity incorrectly. Replacing ensemble expectations with time-based growth measures resolves long-standing paradoxes such as the St. Petersburg paradox[37] and leads to different conclusions about optimal decision-making. A time-average approach, often linked to logarithmic utility, better reflects the compounding and path-dependent nature of wealth accumulation.

In multiplicative systems – including financial markets – arithmetic mean returns (ensemble average) and geometric mean[38] returns (expected time average) diverge when volatility is present. The geometric mean is lower and more relevant to long-term investors. Peters' work on leverage demonstrates that strategies optimised for ensemble averages, such as those implying infinite leverage, can be ruinous over time. Widely used metrics like the Sharpe ratio overlook this effect and may give misleading signals. Time-based optimisation methods, such as the Kelly criterion[39] (discussed below), identify leverage levels that maximise long-run growth.

Finally, empirical evidence reinforces the non-ergodic nature of real-world finance. Wealth distributions in the U.S. are strongly shaped by initial conditions and multiplicative effects, while experiments show that even trained decision-makers often fail to navigate non-ergodic environments effectively.

These findings underline the importance of putting time and sequencing at the centre of investment thinking, including portfolio design.

Ergodic Assumptions in Traditional Asset Management Models

Many foundational models in asset management rest on the implicit assumption of ergodicity. While this simplifies analysis, it misrepresents the dynamics of wealth accumulation in real-world, non-ergodic markets.

MPT[40], for example, optimises a portfolio's expected return over a single period, abstracting away from how wealth compounds over time. As a result, it overlooks key non-ergodic effects such as volatility drag, sequencing risk, and the impact of irreversible losses. Although Markowitz (1952) allows for subjective return expectations, the common practice of extrapolating historical returns to estimate future performance reinforces an ergodic view.

The CAPM[41] extends MPT's single-period logic to asset pricing. It assumes that returns are independent and identically distributed (i.i.d.) and that beta[42] fully captures risk. From a non-ergodic perspective, this is misleading: assets with identical expected returns but different volatilities produce very different long-term outcomes once compounding is considered.

Expected Utility Theory, widely used to model decision-making under risk, similarly assumes ergodicity. It presumes that the expected utility of a one-period gamble can guide behaviour over repeated plays. In non-ergodic settings, however, sequential outcomes often diverge dramatically from ensemble expectations.

Finally, the notion of time diversification – the belief that risk diminishes with a longer holding period – relies on ergodic reasoning. In reality, in volatile, multiplicative systems, the range of potential wealth outcomes typically widens as time extends, increasing exposure to extreme paths and potential ruin.

In sum, traditional models implicitly treat risk as something that averages out, whereas in practice it compounds and accumulates along a single, irreversible path. This oversight can lead to flawed investment strategies and poor long-term results.

Rethinking Risk, Return, and Long-Term Wealth in a Non-Ergodic World

The accompanying paper expands substantially on these topics, beginning with the observation that financial returns are typically generated by multiplicative processes – a class of non-ergodic systems. These systems have four defining features: path dependence, proportionality, quantifiability in monetary terms, and market-based interchangeability.

Recall the simple thought experiment above, which is worth repeating, that illustrates the asymmetry inherent in multiplicative dynamics. Suppose an asset gains +50% and subsequently loses −50% (or vice versa). The gain multiplies wealth by 1.5, but the loss divides it by 2, producing a net result of −25%. This tells us that losses (expressed as a percentage) have a disproportionately large impact when returns compound over time. As a result, outcome distributions become skewed, often following a log-normal shape where the median return is lower than the mean.

Examples such as Peters' coin-toss game and simulated investment returns demonstrate how volatility drag – the wedge between arithmetic and geometric mean returns – erodes long-term wealth. If volatility increases and the arithmetic mean return remains unchanged, the corresponding geometric mean declines. A useful approximation captures this effect:

$$\bar{x}_g \approx \bar{x}_a - \frac{\sigma^2}{2} \tag{1}$$

where $\bar{x}_g$ is the geometric mean of variable x, $\bar{x}_a$ is the arithmetic mean, and σ^2 is the variance of the underlying distribution from which returns are drawn. This rule-of-thumb arises from Geometric Brownian Motion in continuous time and relies on simplifying assumptions so it should be applied with caution; but, nonetheless, it remains a powerful heuristic to think about the drag imposed by volatility.

Figure 2 illustrates this phenomenon. It plots the compounded returns of portfolios that all have an arithmetic mean return of 8% but differ in standard deviation (0–50% in 10% increments). An initial investment of €100 over 10 years shows a clear pattern: higher volatility systematically reduces the geometric mean return.

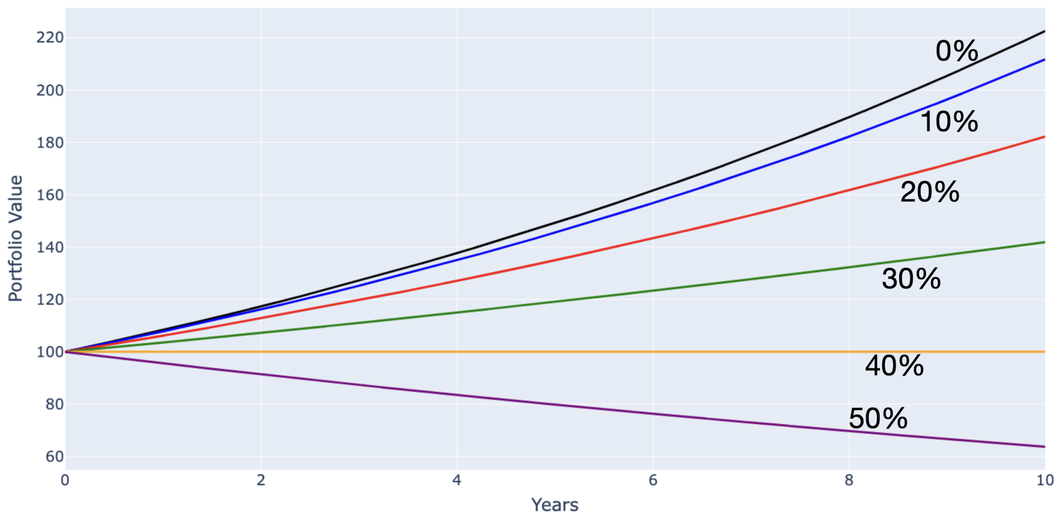

Figure 2: Time series of geometric mean returns for a portfolio with an underlying distribution that has a mean of 8% and varying standard deviations (0–50% in 10% steps). The graph highlights how greater volatility lowers long-term compounded returns.

This insight reframes risk aversion. Investors must be compensated not only for the dispersion of returns – assuming they are known – but also for the structural drag volatility imposes on long-term wealth growth.

The accompanying paper then examines path dependence in portfolio management, particularly sequence-of-returns risk. For investors making withdrawals (e.g., retirees) or managing active portfolios, early losses can permanently impair future wealth, even if long-run average returns appear attractive. This highlights a crucial principle: for long-term investing, survival – "staying in the game" – should take precedence over maximising expected returns.

Finally, the Kelly criterion is introduced as one approach to optimising the fraction of wealth allocated to risky assets in order to maximise long-term geometric mean growth:

$$f^* = \frac{\mu - r_f}{\sigma^2}$$

where f^* is the Kelly fraction, μ is the expected return of the asset, r_f is the risk-free rate (sometimes omitted), and σ is the standard deviation of returns.

If $f^* < 1$, only a fraction of wealth is allocated to risky assets, with the remainder held in cash. If $f^* > 1$, leverage would be required to reach the optimal exposure. Simulations show that this allocation balances volatility drag against long-run growth potential. The conclusion is clear: in long-term investing, conservative approaches (the tortoise) often outperform high-volatility, high-return strategies (the hare).

Practical Implications for Asset Managers

Recognising non-ergodicity has far-reaching consequences for how portfolios are built, risks are assessed, and success is measured. As noted above, traditional mean–variance optimisation, rooted in ensemble averages, can be misleading because it ignores the time-path nature of real-world wealth accumulation. A focus on maximising geometric mean returns – through log-utility or Kelly-based approaches – better reflects the compounding dynamics of wealth. This often favours portfolios with more stable returns, even if their arithmetic average is lower. Diversification takes on a renewed purpose: not just to smooth volatility, but to reduce the risk of ruin by limiting exposure to adverse return sequences. Systematic rebalancing can also turn volatility into a source of incremental growth, an approach sometimes referred to as "volatility harvesting".

Risk management, viewed through a non-ergodic lens, shifts emphasis from optimising expected risk-adjusted returns to safeguarding the investor's trajectory. Position sizing becomes more conservative, while tail-risk hedges – such as protective options or allocations to safe-haven assets – are valued for their ability to prevent catastrophic losses. Liquidity buffers gain additional importance, helping to avoid forced asset sales during market stress. Tools like scenario analysis and stress testing, already common in practice, gain a stronger rationale: they help evaluate vulnerability to adverse time-path outcomes rather than just probabilistic risks.

For client communication, re-framing investment outcomes in time-average terms can help set more realistic expectations. Explaining that arithmetic mean returns do not guarantee realised results – particularly when sequencing risk is significant – supports more resilient decision-making. Highlighting concepts such as geometric returns (CAGR) and the asymmetry of losses can reinforce the value of steady, conservative strategies over aggressive, high-volatility approaches.

Manager incentives often remain aligned with short-term ensemble performance rather than long-term wealth outcomes. Adjusting evaluation horizons, introducing clawbacks, and linking rewards to drawdown-aware or log-wealth growth metrics could improve alignment with investor interests.

Finally, examples of firms already incorporating non-ergodic thinking into practice include Baillie Gifford and NZS Capital, both of which apply long-term, complexity-informed approaches to portfolio management. While such thinking remains far from mainstream, it demonstrates how embracing a time-path perspective can reshape investment strategy and risk management. A deeper exploration of these practices is provided in the companion paper *Non-Ergodicity in Asset Management.*

8 Toward a New Practice

The preceding sections have explored how complexity science can reframe investment thinking and practice. To close, it is useful to place this discussion in its wider historical and intellectual context, highlighting why change in asset management is both necessary and possible.

The dominant mindset shaping modern finance has deep roots in European Enlightenment thought of the 17th to 19th centuries. Earlier philosophical traditions, notably those of ancient Greece, laid important groundwork, but the Enlightenment marked a decisive turn toward rationalism, reductionism, and mechanistic models of cause and effect. These ideas powered major scientific and technological advances and continue to underpin much of economic theory and financial practice today.

Yet this paradigm – based on assumptions of individual optimisation and self-correcting systems – has proved fragile when applied to modern, interconnected economies. Instead of delivering stable, beneficial order, it has contributed to what is now termed the polycrisis, mentioned at the beginning of this paper: a web of interlocking global challenges, many amplified by financial dynamics rather than contained by them.

Asset management is deeply embedded in this history. The industry has played an important enabling role in applying scientific advances and deploying new technologies since the Enlightenment. However, alongside governments and other economic actors, it has reinforced this mechanistic orientation. Also, its focus on short-term profit, reliance on ergodic assumptions (discussed in Section 7), and role in generating endogenous volatility have contributed to cycles of market instability and repeated public interventions. These interventions, in turn, constrain future state responses, compounding systemic fragility.

It is important to clarify the relationship between financial crises and the individuals working within asset management. Public discourse often frames systemic crises as the direct result of bad actors, implying that financial turmoil is caused by unethical or reckless individuals. This is a mechanistic, reductionist view – what I refer to as a 'fractal explanation' of systemic phenomena. In reality, systemic crises can emerge even when intelligent, ethical professionals act responsibly within their roles. As argued in Section 2, endogenous volatility often arises precisely from such micro-level behaviour, amplified through feedback loops in market-based systems. In this sense, responsibility is neither singular nor absent but distributed across the system as a whole.

Against this backdrop, complexity science offers more than an alternative intellectual framework. It provides a practical orientation for reshaping asset management so that it contributes not only to investment success, but also to the resilience of the wider financial and economic systems on which it depends.

Two broad assertions follow from this paper. First, the concept of *complexity arbitrage* highlights how applying complexity science can improve outcomes for asset managers. These improvements extend beyond higher returns: they include more effective management of risk and uncertainty, closer alignment between client needs and portfolio construction, and stronger stewardship within investment firms themselves.

Second, adopting complexity principles more widely has the potential to lessen the industry's contribution to systemic fragility. If embraced at scale, such principles could help dampen endogenous volatility, reduce dependence on mechanistic risk models that underestimate tail events, and encourage investment practices better aligned with long-term societal value creation. While asset management alone cannot resolve the polycrisis, its influence over capital allocation gives it considerable leverage in shaping broader economic and financial outcomes.

This second claim is necessarily cautious. Systemic change cannot be guaranteed by improvements at the level of individual firms. Nonetheless, wider adoption of complexity-informed practices would probably increase the likelihood that asset management supports more resilient market dynamics over time. Financial crises and other forms of instability may still occur, but their frequency and severity could be reduced. Broader resilience will probably require collective measures – new institutions, stronger safeguards, and innovations in market design – yet complexity science offers valuable orientation for moving in that direction.

From this analysis, several themes stand out where complexity science can reshape asset management practice:

Re-framing Reality. Improving decision-making in asset management first requires clarity about the nature of the reality we are trying to navigate. Much of modern finance – and everyday thinking – implicitly assumes a mechanistic ontology, while complexity science challenges this view, offering an alternative framing, as outlined in Section 1.

This shift matters because it changes what we notice, what we measure, and how we act. It underpins key topics explored in this paper: the emergent nature of prices and narratives, the distinction between measurable risk and deeper uncertainty, and the non-ergodic, path-dependent character of returns.

A Cohering Lens. Expertise in asset management is fragmented across disciplines, each with its own language, assumptions, and methods. While diversity is valuable for resilience, it often leads to siloed analysis and incoherent decision-making.

Complexity science provides a shared conceptual frame. Its core ideas – feedback loops, adaptation, emergence, and path dependence – cut across specialisations, offering common ground for dialogue. This does not (and should not) eliminate variety, but structures it, making it easier to integrate diverse perspectives. Developing this shared language takes effort and openness, yet it can strengthen collaboration, improve judgement, and support more robust investment strategies.

Understanding Financial Regimes. Markets do not behave as stable, single-state systems. They tend instead to settle into temporary regimes – patterns of prices, volatility, and behaviour sustained by shared beliefs and reinforcing feedback loops. Narratives help coordinate expectations, making these regimes self-sustaining for a time.

Regimes are neither permanent nor predictable in their transitions. They may shift gradually as evidence undermines prevailing views or collapse abruptly when confidence falters, sometimes triggering crises. Recognising markets as regime-based, rather than perpetually in stable equilibrium, shifts the focus from forecasting a single path to preparing for multiple, evolving patterns of market dynamics.

Rethinking Models. Traditional investment models offer a sense of precision and control, yet they rest on assumptions that often fail in complex, adaptive markets. Their elegance can conceal fragility: patterns shift, feedback loops amplify shocks, and outcomes depend heavily on history and path dependence.

A complexity perspective reframes models as tools for exploration, not prediction. Models should capture empirical patterns, test hypotheses, and be used in plural rather than singular form, recognising that no single model reflects the full dynamics of markets. Above all, models need to be 'held lightly' – continuously updated, challenged, and sometimes discarded – as part of an adaptive approach to decision-making under uncertainty.

Bridging the Application Gap. Complexity science, particularly in its social science applications, often struggles to translate theory into practice. Academic incentives reward conceptual advances, while specialisation can distance research from the realities of decision-making. For asset managers, what matters is whether these insights help navigate markets more effectively.

The orientation of this paper has therefore been pragmatic: to sketch how a complexity-informed perspective can change what asset managers *do*, not only how they *think*. The practical ideas summarised in this paper are starting points for experimentation, not fixed recipes. The real task is to build organisations that learn continuously, adapting tools and practices as market patterns evolve.

Looking Forward

This paper is a first step toward embedding complexity science more fully into asset management practice. As mentioned previously, some firms – notably Baillie Gifford and NZS Capital – are already experimenting with complexity-informed approaches, showing that these ideas can be put into practice in distinctive ways. Such initiatives are still rare, highlighting considerable untapped potential for the broader industry.

Progress will require collaboration between theorists and practitioners to design new tools, test them in live market settings, and integrate them into investment processes. Pilot projects, shared research platforms, and experimental approaches within firms can help turn conceptual insights into practical capabilities. Complexity science provides orientation rather than prescription; its value will only be realised through collective experimentation and adaptation.

My contribution is to help bridge theory and practice, bringing together perspectives from both worlds. Yet this challenge is larger than any single author or paper. If asset management is to evolve toward approaches that better reflect complex, evolving markets – and that contribute to systemic resilience rather than fragility – it will take shared effort, openness to change, and a willingness to rethink long-standing assumptions.

The stakes reach beyond investment performance alone. Asset management directs vast amounts of capital, shaping market outcomes and wider economic and social systems. A complexity-informed approach can make this influence more adaptive, constructive, and resilient. Whether the industry embraces this opportunity – experimenting, learning, and building new practices for a world that is likely to grow more complex – remains an open question.

Endnotes

[1]Recent manifestations of the polycrisis include: the war in Ukraine, which has disrupted energy and grain markets while reshaping geopolitical alliances; escalating tensions between Israel and Iran, contributing to instability across the Middle East and volatility in oil markets; climate-induced agricultural disruptions in the Horn of Africa and the American Midwest, affecting food security and inflation; long-tail effects of the COVID-19 pandemic on supply chains, labour markets, and public debt; regional banking failures and institutional fragility in the US and Europe Arnold (2023), revealing vulnerabilities in financial architecture; and the rapid deployment of generative AI, altering employment, competition, and governance dynamics across sectors.

[2]A way of thinking that models systems as predictable, linear, and decomposable into independent parts.

[3]Systems composed of many interacting elements, whose collective behaviour cannot be understood simply by analysing the parts in isolation.

[4]The arising of macro-level patterns or outcomes from micro-level interactions, often in ways that are not obvious or intended.

[5]The principle that current and future states of a system are shaped by its unique past trajectory and cannot simply reset or repeat.

[6]Circular cause-and-effect processes where an outcome of a system influences future behaviour of that same system, reinforcing or dampening changes.

[7]A system is said to be *ergodic* if, over time, it visits all the possible states – or configurations – available to it in its state space. Non-ergodic systems do not. In ergodic systems, the *time average* of an observable property for a single trajectory (e.g., one individual, one asset) converges to the *ensemble average* – that is, the average across many parallel instances of the system at a single point in time. This is not true of non-ergodic systems.

[8]The reduction in long-term investment returns caused by the asymmetric effect of losses and gains when returns compound.

[9]See Arthur (2013) for a discussion of how complexity science can be understood as a 'movement' across disciplines, and for an excellent discussion of how complexity science can be deployed in economics.

[10]A mathematical approach to modelling how systems evolve over time, particularly those exhibiting sensitivity to initial conditions, non-linearity, and chaotic behaviour.

[11]A framework developed in biology and engineering to study how systems maintain their integrity through interdependence, regulation, and feedback among parts and wholes.

[12]An approach that focuses on the interaction and adaptation of many decentralised components, often leading to emergent and non-linear outcomes.

[13]The study of communication, control, and feedback in both living and artificial systems, often emphasising self-regulation.

[14]Originally the attempt to model human reasoning and learning through algorithms, now encompassing a broader range of adaptive, decision-making systems

[15]An investment framework introduced by Markowitz (1952), which proposes that investors can construct an optimal portfolio by balancing expected return and variance (risk) of returns through diversification, assuming known probabilities and normally distributed returns.

[16]A traditional financial model that estimates the expected return of an asset based on its risk relative to the market as a whole.

[17]This hypothesis holds that asset prices fully incorporate available information, so persistent outperformance on a risk-adjusted basis is unlikely. It is often presented in three forms: weak (prices reflect past price data), semi-strong (prices reflect all public information), and strong (prices reflect all information, including private).

[18]The Black–Scholes model, introduced by Fischer Black and Myron Scholes (Black and Scholes, 1973), provides a closed-form solution for pricing European options. It assumes frictionless markets, constant volatility and risk-free rates, and asset prices following a geometric Brownian motion. This model underpins much of modern derivatives pricing but has well-known limitations in capturing real-world market dynamics.

[19]A statistical measure of the "tailedness" of a distribution, indicating how often extreme values occur relative to the normal distribution.

[20]Statistical correlations between current values and past values of a time series, often used to detect patterns or predictability.

[21]Mathematical relationships where the frequency of an event scales as a power of its size, often implying the absence of a characteristic scale.

[22]A property where patterns look similar at different levels of magnification, indicating that no single time scale or size dominates system behaviour.

[23]Arising without central control, through local interactions that generate global structure or pattern.

[24]Reduction is when significant explanatory patterns are missing from a model; abstraction is when a model includes all such relevant patterns.

[25]System-level behaviours or patterns that arise from many individual interactions and cannot be explained by analysing those parts alone.

[26]My PhD research overlapped significantly with this subject: socially constructive forms of order can emerge spontaneously in complex social systems but this is not inevitable. Conditions have to enable such emergence, and damaging forms of order can also develop depending on the wider environment.

[27]This terminology follows Geoffrey Hodgson (2011).

[28]The idea, proposed by economist Friedrich Hayek (1945), that market prices act as a decentralized information-processing mechanism, aggregating local and private knowledge held by individual actors into a collective signal that guides economic decisions.

[29]See, for example, Tversky and Kahneman (1981), Kirchler et al. (2004), and Nair et al. (2022).

[30]Narratives can be understood as emergent in two distinct senses. First, at the micro level, the narratives embedded in our mental models are often products of historical and cultural evolution. Over time, societies develop recurring story patterns – for example, archetypal “hero and villain” narratives – that shape how individuals interpret events and navigate uncertainty. Second, at the macro level, financial markets reflect a constantly shifting interplay of these micro-level narratives. At any given time, some narratives gain prominence while others recede, influencing market sentiment and price formation to varying degrees. We can think of narratives as carrying different “weights” in the market, their influence waxing and waning as collective attention and conviction shift over time.

[31]Patterns of market behaviour, such as periods of high or low volatility, that persist for a time before shifting into different patterns.

[32]A standard portfolio construction approach that aims to maximise expected return for a given level of risk, assuming returns are normally distributed.

[33]A risk management metric estimating the maximum potential loss of a portfolio over a given period with a specified confidence level.

[34]An average outcome calculated across many possible parallel versions of a system, assuming all outcomes occur simultaneously.

[35]For example, Davidson (1991, 1994, 1995, 1996).

[36]A decision-making concept in economics that calculates the best choice by weighing all possible outcomes according to their probabilities and associated ‘utility’ (value).

[37]A classical problem in probability theory where a gamble with infinite expected value fails to correspond to how individuals actually behave, highlighting the limitations of expected value as a decision criterion.

[38]The geometric mean is a way of finding an average that reflects compounding effects or proportional changes. It is commonly used in finance to describe long-term investment growth because it accounts for the fact that gains and losses build on one another. Unlike the arithmetic mean, it gives a more realistic picture of overall growth when outcomes fluctuate. It can also summarise the central tendency of a distribution of positive values observed at a single point in time.

[39]A mathematical formula used to determine the optimal fraction of wealth to invest in a risky asset to maximise long-term growth.

[40]An investment framework introduced by Markowitz (1952), which proposes that investors can construct an optimal portfolio by balancing expected return and variance (risk) of returns through diversification, assuming known probabilities and normally distributed returns.

[41]A traditional financial model that estimates the expected return of an asset based on its risk relative to the market as a whole.

[42]β is a measure of an asset's sensitivity to movements in the overall market. In CAPM, it represents the expected change in the asset's return for a one-unit change in the return of the market portfolio. A β of 1 implies the asset moves with the market, $\beta > 1$ indicates higher volatility than the market, and $\beta < 1$ indicates lower volatility.

References

Arnold, M. (May 2023), *Recent banking turmoil exposed flaws in liquidity rules, say regulators*, Accessed 28 July 2025, Financial Times, URL: https://www.ft.com/content/abf7b3f5-ce68-48c1-be78-7744aa0e7421.

Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R. and Taylor, P. (1996), *Asset pricing under endogenous expectations in an artificial stock market*, tech. rep., URL: www.santafe.edu/research/results/working-papers/asset-pricing-under-endogenous-expectations-in-an- (visited on 04/05/2023).

Arthur, W.B. (2013), *Complexity Economics: A Different Framework for Economic Thought*, Working Paper 2013-04-012, Sante Fe Institute, URL: https://sfi-edu.s3.amazonaws.com/sfi-edu/production/uploads/sfi-com/dev/uploads/filer/a1/3e/a13e8ad4-cd39-4422-8cc3-86c543699f6d/13-04-012.pdf.

Axtell, R.L. and Farmer, J.D. (2025), 'Agent-based modeling in economics and finance: Past, present, and future', *Journal of Economic Literature* 63 (1), pp. 197–287.

Black, F. and Scholes, M. (1973), 'The Pricing of Options and Corporate Liabilities', *Journal of Political Economy* 81 (3), pp. 637–654, DOI: 10.1086/260062, URL: https://doi.org/10.1086/260062.

Bossaerts, P., Ghirardato, P., Guarnaschelli, S. and Zame, W.R. (2010), 'Ambiguity in Asset Markets: Theory and Experiment', *Review of Financial Studies* 23 (4), pp. 1325–1359, DOI: 10.1093/rfs/hhp100, URL: https://doi.org/10.1093/rfs/hhp100.

Boulton, J. (2024), 'The Dao of complexity: making sense and making waves in turbulent times', *The Dao of Complexity*, De Gruyter.

Brenner, M. and Izhakian, Y. (2018), 'Asset Pricing and Ambiguity: Empirical Evidence', *Journal of Financial Economics* 130 (3), pp. 503–531, DOI: 10.1016/j.jfineco.2018.07.009, URL: https://doi.org/10.1016/j.jfineco.2018.07.009.

Castellani, B. and Gerrits, L. (2021), *2021 Map of the complexity sciences*, URL: www.art-sciencefactory.com/MAP2021Sharing.pdf.

Davidson, P. (1991), 'Is Probability Theory Relevant for Uncertainty? A Post Keynesian Perspective', *Journal of Economic Perspectives* 5 (1), pp. 129–143, DOI: 10.1257/jep.5.1.129.

Davidson, P. (1994), *Post Keynesian Macroeconomic Theory: A Foundation for Successful Economic Policies for the Twenty-First Century*, Aldershot, Hants: Edward Elgar.

Davidson, P. (1995), 'Uncertainty in Economics', *Keynes, Knowledge and Uncertainty*, ed. by S. Dow and J. Hillard, Aldershot, Hants: Edward Elgar, pp. 107–116.

Davidson, P. (1996), 'Reality and economic theory', *Journal of Post Keynesian Economics* 18 (4), pp. 479–508.

Debelle, G. (2010), 'On Risk and Uncertainty', *BIS Review* (111), Speech at the Risk Australia Conference, Sydney, 6 July 2010, URL: www.bis.org/review/r100709c.pdf.

Ellsberg, D. (1961), 'Risk, Ambiguity, and the Savage Axioms', *The Quarterly Journal of Economics* 75 (4), pp. 643–669, DOI: 10.2307/1884324, URL: https://doi.org/10.2307/1884324.

Epstein, L.G. and Wang, T. (1994), 'Intertemporal Asset Pricing under Knightian Uncertainty', *Econometrica* 62 (3), pp. 283–322, DOI: 10.2307/2951650.

Gilboa, I. and Schmeidler, D. (1989), 'Maxmin Expected Utility with Non-Unique Prior', *Journal of Mathematical Economics* 18 (2), pp. 141–153, DOI: 10.1016/0304-4068(89)90018-9, URL: https://doi.org/10.1016/0304-4068(89)90018-9.

Gode, D.K. and Sunder, S. (1993), 'Allocative Efficiency of Markets with Zero-Intelligence Traders', *Journal of Political Economy* 101 (1), pp. 119–137.

Grossman, S.J. and Stiglitz, J.E. (1980), 'On the Impossibility of Informationally Efficient Markets', *The American Economic Review* 70 (3), pp. 393–408, URL: www.jstor.org/stable/1805228.

Haldane, A.G. (2009), *Why Banks Failed the Stress Test*, Speech delivered at the Marcus-Evans Conference on Stress-Testing, London, URL: www.bis.org/review/r090218d.pdf.

Hodgson, G.M. (2011), 'Downward Causation - Some Second Thoughts', URL: www.geoffreymhodgson.uk/downward-causation.

Keynes, J.M. (1937), 'The General Theory of Employment', *Quarterly Journal of Economics* 51 (2), pp. 209–223, DOI: 10.2307/1882087.

Kirchler, E., Hoelzl, M. and Kamleitner, K. (2004), 'Framing Effects, Selective Information, and Market Behavior: An Experimental Analysis', *Journal of Economic Psychology* 25 (4), pp. 441–452, DOI: 10.1016/S0167-4870(03)00041-6.

Markowitz, H. (1952), 'Portfolio Slection', *The Journal of Finance* 7 (1), pp. 77–91, DOI: https://doi.org/10.1111/j.1540-6261.1952.tb01525.x, eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1952.tb01525.x.

Nair, S., Haridas, H. and Nair, K. (2022), 'Can Framing Affect Investment Decisions? An Experimental Study', *Journal of Behavioral Finance and Economics* 12 (2), Preprint available on ResearchGate, pp. 45–56, URL: www.researchgate.net/publication/365710950_Can_Framing_Affect_Investment_Decisions_An_Experimental_Study.

Palmer, R., Brian Arthur, W., Holland, J.H., LeBaron, B. and Tayler, P. (1994), 'Artificial economic life: a simple model of a stockmarket', *Physica D: Nonlinear Phenomena* 75 (1), pp. 264–274, ISSN: 0167-2789, DOI: https://doi.org/10.1016/0167-2789(94)90287-9, URL: www.sciencedirect.com/science/article/pii/0167278994902879.

Plott, C.R. and Sunder, S. (1988), 'Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets', *Econometrica* 56 (5), pp. 1085–1118, DOI: 10.2307/1911359.

Shiller, R.J. (2019), *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*, Princeton University Press, ISBN: 9780691210261, URL: http://www.jstor.org/stable/j.ctv10vm1xf (visited on 09/11/2023).

Smith, V.L. (1962), 'An Experimental Study of Competitive Market Behavior', *Journal of Political Economy* 70 (2), pp. 111–137.

The Santa Fe Institute (2023), *Complexity Explorer: Glossary*, URL: www.complexityexplorer.org/explore/glossary (visited on 21/04/2023).

Tuckett, D. (2017), 'The Role of Conviction and Narrative in Decision-Making under Radical Uncertainty', *Theory & Psychology* 27 (5), pp. 501–523, DOI: 10.1177/0959354317713158, URL: https://journals.sagepub.com/doi/full/10.1177/0959354317713158.

Tversky, A. and Kahneman, D. (1981), 'The Framing of Decisions and the Psychology of Choice', *Science* 211 (4481), pp. 453–458, DOI: 10.1126/science.7455683, URL: www.science.org/doi/10.1126/science.7455683.

Ullmann-Margalit, E. (1978), 'Invisible-hand explanations', *Synthese* 39, pp. 263–291.