WILEY

Review

Reviewed Work(s): The Theory of Industrial Organization. by Jean Tirole

Review by: Marius Schwartz

BOOK REVIEW

THE THEORY OF INDUSTRIAL ORGANIZATION
by Tirole, Jean. Cambridge, MA: MIT Press, 1988,
pp. xii, 479, ISBN 0-262-20071-6.

Introduction: What, How and for Whom?

In economics, unlike some other social sciences, aca-
demic brownie points are awarded primarily for publish-
ing journal articles rather than books. Correspondingly,
the incentive to write books—other than introductory or
intermediate texts—is quite weak. Practicing what they
preach, economists respond by focusing on the journals.
It is therefore rather unusual for a scholar at the peak of
his creative energy to devote the kind of time and effort
needed to write a work that surveys an entire field. This is
an exception, and we are the fortunate beneficiaries.
Tirole has written a remarkable book that will serve both
as an invaluable reference source and a wonderful teach-
ing aid in a wide variety of courses.

The book's title, The Theory of Industrial Organization,
accurately describes the content. The focus is on theory;
institutional and empirical material is introduced only to
motivate and provide some perspective on the theoretical
analysis. Also, the label 'The Theory' is fully justified: the
coverage of theory is thorough and comprehensive, treat-
ing virtually every major topic in industrial organization.
Most importantly, the treatment is not in an 'X says this
and Y says that' manner, but through a series of self-
contained, elegant models that capture the essence of
each argument. Tirole has an uncanny knack for stripp-
ing down a problem to its bare essentials, no more and no
less. It is no mean feat to sustain this level of exposition
over almost five hundred pages, but he pulls it off.

Given the quality of the product, it is hard to quarrel
with the author's decision to limit himself to the theory
and largely leave out, in addition to empirical work,
subjects such as antitrust and regulation. To attempt a
comprehensive coverage of such issues as well would
have meant a book of unwieldy length. Rather than be
subjected to half-hearted job, the reader is better off using
Tirole's book for the theory of industrial organization
and supplementing it with empirical and institutional
sources according to interest. One possibility is the forth-
coming text by F. M. Scherer and David Ross (Industrial
Market Structure and Economic Performance, 3rd edition,
Boston: Houghton Mifflin, 1989).

I should stress that the theory presented is not 'art for
art's sake'. Far from it. The underlying concern is always
with the economic phenomena. Models are suggested by
the problem, not vice versa, and they are presented in a
way that highlights the basic intuition as clutter-free of
technical details as possible while retaining the necessary
rigor. Much of what passes for theoretical work in
economics today can be rather esoteric. But an applied
economist impatient with such work will find this book a
different kettle of fish. It attempts to simplify rather than
complicate, clarify rather than obscure. What makes all

this possible is Tirole's terrific and rare blend of technical
skills, economic intuition and expositional talent. These
are evident throughout the book.

While the book does not cover all aspects of industrial
organization, in some ways its scope extends beyond this
field. Some of the main theoretical developments in
economics in the last decade or so originated in industrial
organization contexts, and the ideas and techniques are
increasingly applied elsewhere. Examples include the
study of imperfect competition and strategic trade policy
in international trade, and gaming aspects in macro-
economics (credibility of government policy to private
agents, policy co-ordination among sovereign states).
Tirole's book provides a good foundation for studying
such problems and thus should be interesting to many
economists outside of industrial organization proper.

Before turning to the content, let met make a few
remarks about the writing style. It is concise, and takes a
bit of getting used to. Sophisticated ideas are often ex-
pressed succinctly and without repetition (though the
reader is reminded when a given model is encountered in
a new context), so it can take several readings to make
sure one has not glossed over something. Also, occasion-
ally one has to fill in some missing steps. The standard of
exposition, however, is consistent; once on the author's
wavelength, the reader will find the exposition uniformly
excellent. Tirole does not waste words or the reader's
time, the arguments are tight, and there is no hand waving
at difficult points. Tirole is an economist's economist. It is
amazing how much state-of-the art material he has
organized coherently and explained lucidly.

This feat is all the more impressive considering that
Tirole is not exclusively an industrial organization
economist. True, he has made significant contributions in
industrial organization (mainly in the area of dynamic
oligopoly), but his main work has been in the area of
incentives. This different background, while requiring
him to master an additional field, also enables him to
treat old problems from fresh perspectives and provide a
unified treatment of seemingly distinct issues. Many of
the arguments and techniques used are familiar to people
working in the area of incentives but less so within
industrial organization proper, and the cross-fertilization
is highly productive.

A simple example will help to make this concrete.
Consider the result that a monopolist whose marginal
cost curve becomes lower will not reduce output. Using a
revealed preference argument, Tirole offers the following
elegant proof that dispenses with typical restrictions such
as the shape or differentiability of marginal cost. Since
the initial output q^* was profit-maximizing given the old
cost function, it was not worth reducing output below q^*;
with the lower marginal cost function, the cost savings
from reducing output are less while the revenue loss
would be the same—hence the new profit-maximizing
output cannot be below q^* (and will be above if marginal
revenue is continuous). This property is used later in the

book to effortlessly establish further nice results. For instance, final price must be higher under successive monopoly than if the two stage are integrated, because the integrated structure would price the input to the downstream division at marginal cost while the unintegrated monopolist charges a price above marginal cost; the higher input price implies (assuming the input is not inferior) that the downstream unit's marginal cost will be higher under successive monopoly, hence (assuming continuous marginal revenue) so will final price. Finally, in an oligopoly context, reducing marginal cost through strategic investments in capital will shift a firm's reaction function in the direction of higher output (Cournot) or lower price (Bertrand). This follows because a reaction function gives a firm's optimal choice taking the rival's choice as given; for any such choice, the investing firm's perceived demand is fixed and we can thus invoke the result developed for a monopolist.

There is a lot of flexibility in how to use the book, both in the depth and range of material covered. Regarding depth, each chapter is divided into the main text and a supplementary section that treats selected points in more detail or introduces more advanced material. There are also exercises in each chapter, indexed by one, two or three stars according to their level of difficulty (three stars usually involve sacrificing the weekend). Hints are provided for the exercises, but they are just that rather than complete solutions. Consequently, one can learn a lot from attempting the exercises even if one succumbs and peeks at the hints. The exercises are definitely a highlight of the book. They test and reinforce the understanding of the material, and develop important results of independent interest from both industrial organization and other branches of microeconomics.

If one foregoes some of the supplementary sections and many of the more technical footnotes (and stays away from those nasty three stars), then the material is quite accessible. The only prerequisites are intermediate microeconomics and multivariable calculus. What is vital, however, is theoretical sophistication—appreciating the difference between necessary and sufficient conditions, understanding proofs by contradiction, etc. The book itself goes a long way towards developing such sophistication, but some background in abstract reasoning is necessary to feel comfortable with its style. (It is not uncommon to reach the end of a proof and have the students staring at you eagerly awaiting the punchline.)

The range of material covered is enormous. We are told in the preface that the book has been used for courses in graduate industrial organization and advanced topics at MIT, and (in yet another indictment of the US educational system) by French and Swiss undergraduates. Given the amount of material and the book's versatility, I will make it a main objective of this review to provide the prospective reader with a coherent roadmap. In order to avoid an unwieldy bibliography, I will omit references to original papers; the book itself is very charitable in attributing credit appropriately.

Organization and Contents

Following the Introduction there is a preliminary chapter on The Theory of the Firm. The remaining eleven chapters are as follows. Chapter 1: Monopoly; Chapter 2: Product Selection, Quality, and Advertising; Chapter 3: Price Discrimination; Chapter 4: Vertical Control; Chapter 5: Short-Run Price Competition; Chapter 6: Dynamic Price Competition and Tacit Collusion; Chapter 7: Product Differentiation: Price Competition and Non-Price Competition; Chapter 8: Entry, Accomodation, and Exit; Chapter 9: Information and Strategic Behavior: Reputation, Limit Pricing, and Predation; Chapter 10: Research and Development and the Adoption of New Technologies; Chapter 11: Noncooperative Game Theory: A User's Manual.

The Introduction presents a nice review of the evolution of industrial organization and perspectives on where it is headed. It also discusses the relation between consumer surplus and the equivalent and compensating variations, and conditions under which consumer surplus generalizes to multiple goods (the path independence problem) and can be aggregated across individuals. Such issues are usually left out of industrial organization texts, and in microtheory texts the relevant material is often interspersed across different chapters. It is therefore convenient to have in one place a tight (if terse) treatment of these issues.

The Theory of the Firm chapter treats two sets of questions: 'What is a firm?' and 'Do firms maximize profit?' Tirole's discussion of the first focuses on transaction cost determinants (rather than, say, legal or tax considerations) of firm size. For example, he addresses the role of specific investments that, ex post, create a bilateral monopoly between buyer and seller. Through clear examples, he demonstrates the various inefficiencies associated with reliance on unconstrained bargaining ex post and how these inefficiencies can be mitigated by long-term contracts that specify how the trading will occur (Who specifies price? Who decides whether to trade or not?) or by joint ownership.

The profit-maximization question is prompted by separation of ownership and control. Tirole approaches shareholders' inability to costlessly monitor managers from the principal–agent framework (in its hidden-action or 'moral hazard' form), showing the tradeoff between insurance and incentives and how it is generally suboptimal to provide full incentives (i.e. to eliminate all managerial slack). The supplementary section provides a fuller treatment of the principal–agent problem (again in the hidden-action case). He then discusses the power and limitations of various mechanisms to reduce managerial slack: yardstick competition, takeover bids, product-market competition, and supervision.

This chapter is relatively difficult. Also, some of the material in it lies in the realm of contract theory and is not encountered much in the remainder of the book. (Actually, it could be very well used in a micro theory or advanced-topics course.) It is possible, therefore, to follow the time-honored tradition of assuming the existence of profit-maximizing 'firms' and plunge into the remainder of the book.

Part I, consisting of Chapters 1 through 4, examines monopoly. Chapter 1 presents the basic theory of monopoly pricing, for one and multiple products, and critically reviews some of the welfare arguments against monopoly

(*X*-inefficiency and rent seeking). For example, the absence of yardstick competition in natural-monopoly industries is mentioned briefly as a possible explanation for why monopolies may experience more *X*-inefficiency. The supplementary section addresses an important special case of demand interdependence, that of a monopolist selling durable goods. It examines the constraints placed on monopoly power by recycling (the famous Alcoa case) and by consumer expectations of lower future prices—the Coase conjecture, that price is driven to marginal cost if the monopolist cannot commit to restricting future sales.

Chapter 2 deals with questions of quality and product selection, first under complete information, then under asymmetric information. Two basic models in the study of product differentiation are presented: vertical differentiation (different qualities) and horizontal or 'spatial' differentiation (where distance is in geographic or product space). Using these models, which are encountered repeatedly in the book, he compares product selection by a monopolist with the socially efficient solution. Whereas the monopolist's quantity is too low, no clear bias exists concerning the number of products. The monopolist typically cannot capture the full consumer surplus from a given product, which suggests an underincentive to add products. But setting price above marginal cost creates an artificial demand for substitute products, which could lead a monopolist to oversupply substitutes: a planner might reduce price to marginal cost and forego adding a product to save the fixed cost. These types of welfare tradeoffs recur in Chapters 7, 8 and 10. Such second-best tradeoffs are familiar from the literature, but they are nicely illustrated in the book in concrete settings. A similar ambiguity arises for quality: the monopolist's incentive to provide quality is too low or too high as the marginal consumer values quality increments less or more than inframarginal consumers.

A distinct bias towards low quality emerges when buyers' and sellers' information about quality is asymmetric. (Many of these problems are not confined to monopoly; they can arise also when the seller faces competition.) Tirole first highlights the market failures that can arise in a one-shot interaction when the buyer cannot observe the seller's action ('moral hazard') or the seller's knowledge ('adverse selection', as in Akerlof's famous lemons paper). Then he discusses how repeat purchases can mitigate such conflicts. The supplementary section addresses repeat purchases in greater detail, e.g. the possibility of signalling one's quality through introductory offers, or that high quality is provided because of the threat of losing 'quality premia' associated with one's reputation if one cheats.

Chapter 3 moves from uniform-price monopoly to price discrimination. This chapter is one of my favorites and develops a subject that is vey mugh 'real world'. For these reasons, let me discuss it in more detail. Following Pigou, the traditional taxonomy is into first-, second- and third-degree price discrimination. (I have often wondered whether the qualifier 'degree' is inspired by how badly consumers are burned by price discrimination.) First-degree refers to perfect discrimination. Second-degree encompasses two-part tariffs—the charging of a fixed fee and a constant marginal price—or other nonlinear pricing (offering any bundles that specify a quantity and a total dollar amount). Third-degree refers to the charging of different price schedules to different groups identified by some exogenous signal. In order to make a clean separation, Tirole represents third-degree (or 'multimarket') discrimination as the charging of different *linear* prices to different groups. He then proceeds to two-part tariffs and to fully nonlinear pricing.

The various types of price discrimination are carefully motivated in terms of the information and arbitrage constraints that could give rise to them. Unconstrained commodity arbitrage mandates uniform, linear pricing. (Two-part tariffs are ruled out because no more than two buyers would pay the fixed fee. Other nonlinear pricing is ruled out because only the bundle offering the lowest average price will be bought, so the monopolist might as well charge the optimal price—the simple monopoly price.) Thus, Tirole's multimarket price discrimination can arise if commodity arbitrage is infeasible among the exogenously-identifiable groups (e.g. because they are in different geographic markets) but it is perfect within a group. After presenting the standard inverse-elasticity rule for multimarket discrimination, Tirole discusses practices designed to implement such price discrimination.

One example is vertical integration into the more elastic market to prevent resale. Another is freight absorption—the charging of delivered prices that rise by less than the transport cost. The discussion of the latter nicely illustrates how results developed earlier in the book are combined to simplify the analysis of later problems. Supposing that consumers at different locations have identical demands at a given delivered price and that delivery cost increases linearly with distance, the monopolist's preferred *delivered* price can increase by less or more than the transport cost to that location. (Exercise 1.3 shows that an increase in marginal cost can increase the monopoly price by less or more than one-for-one, depending on the particular demand function.) This, by itself, suggests that discrimination could favor more distant *or* closer consumers; but arbitrage constraints are likely to rule out discrimination in cases in which it would favor closer consumers, thereby helping explain why the typical pattern observed is freight absorption.

Two-part tariffs are examined next. Since offering the same two-part tariff to all is a rather unsophisticated scheme, some explanations are offered for its use. One possibility is that the manufacturer cannot observe the level of sales of its dealers but can observe whether a given dealer is carrying its product or not (presumably beyond some minimal level that would qualify it as a 'dealer'). Assuming that a court would enforce a franchise fee in the latter case, this could explain the use of uniform two-part tariffs. A second possibility arises when each consumer must purchase a single unit of one monopolized good in order to consume variable amounts of a second, complementary good (e.g. camera and film) and arbitrage requires both prices to be uniform across consumers. The price of the first good is essentially a fixed fee paid for the right to purchase any quantity of the variable

good at the given marginal price. With identical consumers, a monopolist would price the variable good at marginal cost and collect the resulting consumer surplus through the price of the fixed good. When demands differ, Tirole shows that (provided it is profitable to serve more consumers than just the highest-demand one) the profit-maximizing price for the variable good is above marginal cost—formalizing the old intuition that the variable good acts as a 'metering and extraction' device. Overpricing the variable good, however, can attract entry; tying the variable good to the fixed good, as in the 1936 IBM case involving calculating machines and tabulating cards, can be interpreted as an attempt to defend the two-part tariff.

Tirole next discusses fully nonlinear pricing—the offering of a menu of bundles, each specifying a quantity and total dollar amount. Suppose that commodity arbitrage is infeasible but that the monopolist knows only the distribution of consumer demands, not who is who. The monopolist can try to offer a menu of bundles that will induce the various consumer types to self-select (reveal themselves) through their purchases. In designing the bundles, the monopolist must ensure that each demand type chooses the bundle targeted to it. This threat of personal arbitrage imposes self-selection (or incentive-compatibility) constraints on the bundles offered. The case of two types of consumers is analyzed in the main text and the general case with a continuum of types appears in the supplementary section; for plausible demand structures, quantity discounts emerge. The ideas and techniques used apply also to other examples of hidden knowledge (e.g. design of auctions, or regulatory pricing, or taxation), making the supplementary section valuable beyond the specific context of price discrimination. It establishes the familiar results that consumption levels of all-but-the-highest demander are distorted, to relax the arbitrage constraint of each higher type, and that all but the lowest demander enjoy some surplus.

The model of quantity discrimination, whereby the quantity in a bundle refers to physical units, is elegantly reinterpreted as a quality model (using the vertical differentiation model introduced in Chapter 2). The analogous result now is that consumers with the highest valuation for quality purchase their socially optimal quality level, while those with lower valuation purchase suboptimal quality. That is, the monopolist magnifies the quality differentials—by reducing the low quality—so as to enforce discrimination. In a colorful example dating to Dupuit (1849), the railroad maintains open carriages with wooden benches in third class, not because third-class passengers would be unwilling to pay the extra cost of rectifying this, but in order to keep the rich in first-class where they are milked for a higher price. Additional insightful examples are offered of quality discrimination, including insurance policies with different levels of coverage, and the deliberate creation of waiting times or price dispersion (only the low-time-value, low-demand buyers take advantage of sales). The latter corresponds to tying a 'bad' to the basic good, because the high demanders are willing to pay more than low demanders to eliminate the 'bad'. Airline restrictions involving stay-over-a-weekend, 30-day advance, never-on-a-Monday,

and other such charming practices fall into this category of quality discrimination.

Chapter 4 studies vertical control. It, too, is one of my favorites. Thus far, the monopolist was assumed to sell to final consumers. Additional complications arise when instead the customers are firms. Demands of such customers typically are not independent; for instance, a distributor's demand for the manufacturer's product depends on the quantity bought by that distributor's competitors. Also, the downstream firms generally undertake actions such as advertising, promotion and other demand-increasing services that significantly affect the initial seller's reputation and ultimately profit. Finally, the scope for exerting control over such customers is greater than over final consumers, both because of their relative size (it may not pay to monitor the behavior of millions of small consumers, but it may pay to monitor one hundred distributors) and for legal reasons (a manufacturer can more easily terminate a dealer that harms the product's reputation than a consumer who misuses a product and complains vociferously to other consumers).

There are generally many reasons to exert control over such related economic stages. The literature on motives for vertical control is vast, and one must inevitably decide what to cut. Tirole focuses on motives that arise for monopoly reasons, and puts aside other considerations such as transaction costs and incomplete contracts that are present also when both stages are perfectly competitive (at least *ex ante*). Overall, I think this is a good call. A proper discussion of transaction-cost issues requires careful analysis of informational conditions and gets us into the realm of contract theory. Indeed, Tirole addresses some of these issues in the Theory of the Firm chapter. The only change possibly worth making is to pull into Chapter 4 (or at least reference) the discussion of specific investments and opportunistic behavior as providing an incentive for vertical integration or contracting; it is an important idea that might get overlooked by readers that skim the Firm chapter.

Regarding methods of vertical control, vertical restraints are usually interpreted to mean any form of contracting more complex than unrestricted sale at linear prices. The chapter begins by identifying some of the common vertical restraints (e.g. franchise fees, resale-price maintenance, territorial restrictions, exclusive dealing and tie-ins). As in Chapter 3, the discussion of which types of vertical restraints are feasible as a function of informational conditions (i.e. of what the monopolist can observe) is sophisticated. The need for such restraints is interpreted from an externality perspective—the downstream firms ignoring the impact of their decisions on each others' profit and on the monopolist's. Some of the examples discussed from this externality perspective are successive monopoly/double marginalization, input substitution (under variable-proportions technology), dealer services, and provision by the manufacturer of assets that could be diverted to other uses. For each case, Tirole discusses sufficient sets of restraints—combinations that can achieve the joint-maximizing solution for the contract-setting monopolist and its customers.

While vertical arrangements are intended to address externalities among the contracting parties, they ignore

externalities on third parties (consumers and other rival manufacturers, either actual rivals or potential entrants). One therefore should expect that competition-reducing vertical arrangements can sometimes be detrimental to overall welfare. The supplementary section provides two such examples.

Part II of the book drops the assumption that the seller is a monopolist and addresses rivalry. Game-theoretic reasoning is used more heavily in this part of the book. It is possible to follow the material (especially in the main text) without a detailed knowledge of game theory—indeed the game theory is often implicit rather than explicit—but familiarity with its main ideas certainly helps. A synopsis of the main ideas in noncooperative game theory used in industrial organization is given in Chapter 11. The choice of topics and order of presentation are excellent and the lucid exposition is particularly impressive given the difficulty of some of the material. I especially like the parable about an economist consulting a game theorist and a Bayesian statistician to arrive at Perfect Bayesian Equilibrium as the solution concept for dynamic games of incomplete information, and the discussion of refinements of this concept (to reduce the range of equilibria). This is a great chapter. It is rather terse to be covered alone, but makes a wonderful roadmap to use with a separate game theory text or in conjunction with some of the applications in this book. In general, Chapters 5, 7, 8 and 10 use one- or two-stage games of complete information; Chapters 6 and 9 use repeated games of complete and incomplete information and are generally more difficult.

Chapter 5 assumes homogeneous products and presents the Bertrand and Cournot models. In the first, firms simultaneously choose prices; in the second, quantities. Tirole first reviews the 'Bertrand Paradox' that with as few as two constant-cost firms price is driven to marginal cost (of the higher-cost firm if costs differ) and indicates why if marginal costs are not constant, and differentiable, no equilibrium (in pure strategies) exists. The Cournot model is better behaved (equilibrium exists under quite general conditions and adding identical-cost firms reduces price continuously). Tirole presents this traditional Cournot model (section 5.4) and comprehensively discusses its existence, uniqueness and convergence properties (section 5.7.1). Using Cournot examples, he illustrates the pitfalls of using concentration indices and industry profitability for making welfare inferences (section 5.5).

However, he is uncomfortable with the notion that firms literally choose quantities. Most of the rest of the chapter examines conditions under which the constant-cost Cournot equilibrium coincides with: Bertrand equilibrium if firms face exogenously given and observable rigid capacity constraints (marginal costs are inverted-L shaped); and with the equilibrium to a two-stage game in which firms would actually choose these capacities in the first stage anticipating the Bertrand competition in the second (sections 5.3 and 5.7.2).

Tirole's interpretation of Cournot competition is unorthodox. Undoubtedly, much of the popularity of the Cournot equilibrium derives from its convenient modelling properties rather than the inherent plausibility of

firms choosing quantities—the looking-under-the-lamp-post syndrome. But there are markets in which firms do literally pick quantities and accept whatever price is needed to sell those quantities (e.g. the oil market, in which producers sometimes get paid as a function of the realized price). In any case, if one accepts Tirole's view of the Cournot equilibrium as an equilibrium to a capacity-constrained price game, then (as he notes) one can salvage only the constant-cost Cournot model in this way, and only under very special conditions (e.g. capacity levels must be 'low' or 'high', not 'intermediate'). It should thus command less attention later in the book. Given that Bertrand equilibrium does not exist for most cost functions, we are sadly left with no adequate oligopoly model for homogeneous goods.

The distinction between capacity choices and subsequent pricing behavior encountered in Chapter 5 reflects the idea that some variables are easier to adjust than others. This general idea has led to numerous models in industrial organization that share the feature of a two-stage game. In the first stage, firms make irreversible investment decisions (either simultaneously or sequentially). These choices are observed by all and determine the cost or demand parameters for the second stage, in which firms compete by simultaneously choosing prices or quantities. Foreseeing that their investments will affect the subsequent competition, firms can behave strategically—invest so as to change the second-stage competition in their favor. This theme runs through Chapters 7 and 8. The two-stage game framework yields considerable insight into strategic incentives while requiring from the reader a fairly low investment in technique.

Chapter 7 abandons the assumption that firms produce homogeneous products, focusing on the interplay between price competition and product choice. Using the linear-city model introduced in Chapter 2, Tirole now considers two firms that can locate anywhere. He shows that the Bertrand equilibrium prices are lower the closer the firms' locations—closer locations tend to intensify price competition by reducing the extent of differentiation. This creates an incentive to locate away from the rival, the so-called 'principle of differentiation'. An opposing incentive is to locate at the center, so as to prevent the rival from taking the longer flank. (This incentive to cluster could arise for additional reasons, such as proximity to an input source.) The conflicting incentives are explored in a two-stage game, in which first locations (i.e. the degree of differentiation) are chosen simultaneously in anticipation of how they affect the second-stage price competition. This interesting question of whether firms will choose minimal or maximal differentiation is explored further in the supplementary section, this time using the vertical (or quality) model of differentiation.

A second issue addressed in this chapter is whether free entry results in socially excessive or insufficient product variety. Section 7.1.2 uses the spatial model except that consumers are now located around a circle (so that, unlike the center of the line, no location is inherently better), finding that free entry is socially excessive. A different representation of product differentiation is

monopolistic competition (sections 7.2 and 7.5.2). This time, free entry generates too few firms—insufficient variety. This ambiguity is unfortunately rather common in industrial organization; results are often reversed depending on whether product differentiation is represented in a spatial model or as monopolistic competition.

The assumption that investments (location choices in the preceding) are made simultaneously is relaxed in Chapter 8. With sequential investments, the first mover can affect the rival's investment choice, not only the subsequent price or output competition. That is, the investment interaction is dynamic. One example is the Spence–Dixit model (section 8.2), in which the incumbent firm (first mover) chooses a larger capital stock in order to commit itself to a low short-run marginal cost function; the rival's rational response is to scale back its size or stay out altogether. Capacity accumulation games in which timing of investment decisions is endogeneous are discussed in (the more difficult) section 8.6.1.

The two-stage game framework has been widely used in industrial organization and has led to a bewildering array of strategic postures. Section 8.3 develops a very useful and complete taxonomy of firms' investment incentives in such two-stage games. Do I want to be nice to you or mean to you? Well, if the aim is to deter your entry or induce exit, I want to be mean. If I have to live with you, it depends on whether you respond to my expected nastiness by being nice or nasty. The latter question is more commonly phrased as whether firms' instruments in the market game (the second stage) are strategic substitutes or strategic complements. In quantity competition, a higher expected output of one firm both hurts the second and typically induces it to supply lower output—which benefits the first. In price competition (with differentiated products), a low expected price of one firm hurts the rival, but this time the rival's typical response is to reduce its price—which harms the first firm. (Section 8.5 shows that this typical pattern of quantities being strategic substitutes and prices strategic complements is not universal, even if products are substitutes in demand.)

If I choose to be nasty and this requires my being big, I become a 'top dog'. If being nasty requires being small, I stay 'lean and hungry'. (An example is refraining from advertising to develop brand loyalty of some customers because this would reduce the incentive to respond aggressively to an entrant's price cuts to the nonbrand-conscious consumers.) If I want to be nice and doing so is associated with being big (heavy advertising in the preceding example), I opt to be a 'fat cat', while if being nice requires staying small I adopt the 'puppy dog ploy'. Section 8.3 illustrates this taxonomy with a wide range of business practices. These include: adoption of most-favored-customer clauses, voluntary limitation of capacity, tying, choice of systems and product compatibility, learning by doing, and vertical control.

Another important issue addressed in this chapter (section 8.1) is competition under natural monopoly cost conditions. There is a useful discussion of the distinction between fixed and sunk costs, and a demonstration of the radically different welfare properties of two models of competition: contestability and war-of-attrition. In the former, price adjustments are sluggish relative to capital adjustments (entry and exit decisions) and threat of entry forces a single firm to price at average cost—the welfare-maximizing solution subject to no subsidies. In the war of attrition, price adjustments are rapid and two firms compete in the market until one drops out; the first phase is inefficient because fixed costs are duplicated, while the second is inefficient due to monopoly pricing.

The war-of-attrition game is one polar case of timing games. Pre-emption is the other. Section 8.6.2 of the supplementary section studies pre-emption in more detail. Pre-emption possibilities arise when there is a lumpy input that introduces a zero–one decision. Examples of such lumpy inputs include plants of some minimum size, innovations, and a fixed number of locations in geographic or product space. It might then be profitable to deter entry by investing pre-emptively, because typically industry profits will be higher with a single firm than with several, what Tirole calls the 'efficiency effect' (a natural-monopoly cost structure is sufficient but not necessary, since having one firm avoids competition). Concentrating on the spatial model, he shows how occupying appropriate locations early can profitably deter in certain cases, but then discusses various reasons why, more generally, pre-emption can be unprofitable or even infeasible.

The discussion of pre-emption incentives in Chapter 8 ties in quite nicely with the first part of Chapter 10, dealing with incentives to innovate as a function of market structure. Tirole begins (section 10.1) with a review of Arrow's original comparison of the gains to a monopolist versus a firm in a competitive industry to obtain a patentable cost-reducing innovation when the alternative is no innovation by anybody. The monopolist's incentive to innovate is lower, basically because the monopolist is already earning some profit—the 're-placement effect'. Neverthless, a monopolist would be willing to outbid a potential entrant for a single innovation, because now the alternative to not innovating is to have a rival obtain the innovation (the argument below does not extend to multiple innovations). The monopolist's maximum bid is the monopoly profit with the innovation minus his share of duopoly profit when the entrant has the innovation. The entrant bids his expected duopoly profit. By the 'efficiency effect', the monopolist's bid is higher.

Section 10.2 replaces static bidding for an innovation with R&D competition in the form of patent races. With timing now important, he shows in the simplest model of a patent race (probability of success at any point depends only on current expenditure) that either effect can dominate: 'efficiency' or 'replacement'. That is, depending on the production process in R&D (for generating the innovation), the monopolist or potential entrant may be willing to spend more and thus be more likely to innovate first. There follows a discussion of richer models of patent races, e.g. choice of safe versus risky R&D technology, and the importance of observing rivals' actions when success probability depends on cumulative expenditure. Sections 10.3 and 10.4, respectively, briefly discuss the welfare effects of the patent system, and some alternative systems for promoting innovation (award system and government procurement).

The other sections of Chapter 10 discuss various issues relating to the adoption of existing technologies. One issue is whether oligopolists adopt quickly or slowly. Section 10.5 shows models in which the adoption game can take either polar form: pre-emption (adopt first) or war-of-attrition (don't adopt unless other does). Section 10.6 examines the recent literature on standardization and compatibility decisions when there exist network externalities. The supplementary section (10.8) examines patent licensing. This chapter is worthwhile both on its own merits and for enforcing intuitions developed earlier.

In all the above, pricing is noncooperative. Chapter 6 explores possibilities for collusion. The chapter begins (section 6.1) with a brief, informal discussion of factors relevant to collusion, e.g. detection lags and cost asymmetries. Next is a critical discussion of static approaches to repeated interaction, such as the kinked demand curve story and conjectural variations (section 6.2). Such approaches sustain cooperation by assuming that firms expect particular reactions to their own action and deriving the equilibrium given these assumed beliefs. As Tirole forcefully argues, the problem is that reactions involve actions in real time; until the explicit dynamic game is written down one cannot tell if the imposed beliefs make sense or not. The alternative methodology proposed is that of repeated games.

He begins with infinitely repeated games ('supergames') in which firms choose prices simultaneously each period, having observed all past choices (section 6.3). Using this simple framework, he re-examines some of the structural factors relevant to collusion in terms of their impact on the gains from cheating and the future costs. Matters are more complex when rivals' prices are not fully observable even *ex post*; a firm now is not sure if a loss in its sales is attributable to rivals' cheating or to reduced demand (the Green–Porter analysis). Correspondingly, firms must be more circumspect in their punishment strategies, for fear of mistakenly 'blowing up the world'. Of course, this also affects the range of equilibria (perfect collusion generally cannot be sustained). A fuller discussion of secret price cuts appears in the supplementary section 6.7.1.

As Tirole notes, an unsatisfactory feature of simultaneous-move supergames is that strategies condition future responses on bygones, events that do not affect profit. (Sections 6.4 and 6.7.2 consider alternating moves: prices are fixed for two periods so a past price cut does affect current profit.) For instance, the 'grim strategy' sustains cooperation by threatening low prices forever if the rival cuts price once, even if the rival has long raised its price. This methodology is at odds with the principle that bygones are bygones. If one insists on this principle here, by requiring that current choices be independent of irrelevant information, then the cooperative equilibria disappear (if the one-shot game has a unique equilibrium). If one allows bygones to matter, much of economics needs to be re-examined.

One way that rivals' past choices can be made relevant is by relaxing the assumption of complete information. Now, a firm has some uncertainty about a rival's 'type', e.g. about the rival's cost, and makes inferences about the type from the rival's past choices. History will matter not only in a self-fulfilling expectations way but because it affects agents' beliefs about each other's underlying characteristics; this may well be most people's intuition for why past choices should matter. The inference process provides incentives to 'invest in disinformation'—act to change rivals' expectations about one's type in a way that induces more favorable behavior. Thus, a firm might charge a high price in early periods so as to mimic the behavior of a high-cost firm, leading a rival to expect a higher future price and therefore to raise its own price (recall that typically prices are strategic complements). Section 6.5 outlines such a model. Chapter 9 further develops dynamic games under asymmetric information, focusing on exclusionary rather than collusive behavior: limit pricing to deter entry, or predation to deter entry, induce exit, or buy-out rivals on favorable terms. (A firm now wants to appear tough, rather than soft as in the collusion case.)

The asymmetric information framework is both more realistic and addresses the bygones criticism (it also allows cooperation for part of a game even when the horizon is finite). Unfortunately, the problem of multiple equilibria, endemic in complete-information supergames, persists; what signals are sent depends on how the sender expects various signals to be interpreted, and there is wide latitude in the profiles of mutually consistent signals and beliefs. Also, results are very sensitive to the particular informational asymmetry, e.g. whether about demand or rival's cost. So while dynamic games of asymmetric information can rationalize phenomena that traditionally experienced tension in the framework of optimization and equilibrium, they are more difficult and at this stage their predictive power is weak.

Alternative Sequences

The book's chapters are organized around topics; but within a given chapter, the type of models employed can differ considerably depending on the issue at hand, and a given model is often encountered in different chapters. For someone familiar with the general models and techniques, proceeding according to topics is the preferred approach. A well-trained economist can pick up Chapter 4, say, and acquire a good understanding of work on vertical control. For pedagogical purposes, it might sometimes be more convenient to introduce a particular model, then illustrate the different uses to which it can be put.

Regarding order of topics, the author motivates the break between monopoly and oligopoly (Parts I and II) as designed to postpone game-theoretic considerations, an objective that is generally met if one avoids the supplementary sections. However, there are also advantages to introducing game theory early—to give students sustained and gradual exposure to it rather than forcing them to sink or swim when arriving at Part II. In any event, for those wishing to introduce oligopoly and game theory early, one possibility is to establish some of the main ideas for homogeneous products, then proceed to product differentiation. Applications involving asymmetric information could be grouped separately, and

similarly those involving repeated games. Here is one possible order.

Begin with the Introduction, skip or skim The Theory of the Firm chapter, then cover the main text in Chapter 1. From there, proceed to Chapter 5 and cover sections 5.1 and 5.2 (the Bertrand Paradox and an overview of possible resolutions), then section 5.4 on the traditional Cournot model (and, depending on interests, also section 5.5 on concentration indices and industry profitability). The Cournot model is easy for students, since it involves monopoly behavior along a residual demand, and—like it or not—is widely used. By now, the themes of how competition destroys profit and how it affects welfare will have been introduced. One could proceed to Chapter 10, on R&D, and cover sections 10.1 through 10.4 (skimming perhaps 10.2 on patent races). These sections develop the important ideas of the efficiency and replacement effects, and the welfare tradeoffs between monopoly pricing and product availability (i.e. the innovation); they also make good use of Cournot examples.

Having shown the incentive to prevent competition, one could now develop the role of commitments. The Dixit–Spence model from section 8.2 can be used. The Coase conjecture (section 1.5.2) can also be done now. It is an important idea that reinforces understanding of backward induction and shows the value of commitments in a non-oligopoly context (the monopolist benefits from committing not to increase future sales—by leasing rather than selling, for instance). Alternatively, this material could be covered under price discrimination as an example of how being able to engage in intertemporal price discrimination harms the monopolist (consumers expect future price cuts, hence reduce their offering price today). The game theory used implicitly thus far corresponds to sections 11.1, 11.2 and 11.3.1 (games and strategies, Nash Equilibrium, and subgame–perfect equilibrium for games of perfect information). These could be done concurrently if one wishes to make the game theory explicit.

One could turn now to differentiated products, Chapter 2, covering sections 2.1 and 2.2 (product space and product selection); the welfare tradeoffs examined are similar to those in Chapter 10. I would defer issues of asymmetric information and repeat purchase (section 2.3 through 2.6), and go to Chapter 7, which examines product differentiation still under symmetric information. There, cover spatial competition (section 7.1) and perhaps also monopolistic competition (sections 7.2 and 7.5.2). Those interested in the question of maximal-or-minimal differentiation could also do vertical differentiation (section 7.5.1). Having developed Bertrand competition with differentiated products, one could proceed to Chapter 8 and do section 8.3, 8.4 and 8.5. These treat strategic incentives in two-stage games, and give numerous interesting examples. The supplementary subsection 8.6.2 on pre-emption also is worthwhile and not difficult. Section 8.1 (Fixed Costs: Natural Monopoly and Contestability) could also be done now (it is important and should be covered somewhere).

Now turn to Price Discrimination, covering all the material in the text and perhaps also the supplementary

section 3.5.1 on Nonlinear Prices, then proceed to Vertical Control (Chapter 4), covering all the text. The supplementary section 4.6.2 is also quite accessible, and might be covered by those interested in vertical foreclosure.

The remaining uncovered material largely concerns applications of asymmetric information and repeated games: Chapters 6 and 9, and the supplementary section (2.6) of Chapter 2. Sections 6.1 through 6.3 involve repeated games with complete information and could be covered now along with the game theory section 11.3.2 (Games of 'Almost Perfect' information). The quality premia model from section 2.6.2.1 is very similar. Asymmetric information games are trickier. The simplest are static games of incomplete information (section 11.4), in which players move once simultaneously so issues of inferences do not arise (the design of nonlinear pricing treated in Chapter 4 is such a game). More complicated are dynamic games of asymmetric information, in which the possibility exists for influencing others' beliefs about one's type through one's actions (11.5). Again, one could start with the simpler, two-period case in which first the informed party undertakes an action and then the informed responds (the introductory-offers model of quality signalling in section 2.6.1.2 is a particularly simple example, since there are only two qualities, i.e. 'types'). More complicated are the Reputation Games, in which repeated choices are made over time and beliefs are revised based on these choices (as in section 2.6.2 and much of Chapters 6 and 9).

Conclusions

Do not be misled by the innocent-looking cover—this book will never make bedtime reading. Superbly as it is written, the inherent complexity of the material requires studying it carefully with pencil and paper. It is therefore not for everyone; but to the reader interested in learning about an exciting field, it is hard to find time better spent.

Some might find the range of topics and models presented an embarrassment of riches. They can leave the reader shell shocked, groping for 'the answer' or 'the model'. Indeed, a common frustration expressed at modern theoretical microeconomics is that there seems to be a model for every story—'anything goes and nobody knows'. But a main purpose of theory is to identify possibilities and warn against the pitfalls of blind assertions. So it would be silly to fault the book for tackling intricacies head on rather than sweeping them under the rug. Complaining about the range of issues addressed is a bit like the king's complaint in *Amadeus* that 'there are too many notes—just take out a few'.

It is true that occasionally the coverage devoted to a topic is commensurate more with its difficulty than its economic importance. But this is because Tirole has chosen to summarize the state of the field rather than play benevolent dictator: blame the message if you will, not the messenger! He tells it like it is and does so admirably. It is partly up to the reader to judge which factors are of primary importance and which are more tangential.

Finding flaws with the book is hard; one must really stretch. But ,since a reviewer must offer some critical remarks to retain credibility, let me try.

Chapters 6 and 9 probably belong together. They both use repeated games, section 6.7.3 on Folk Theorems (about the set of perfect equilibria for discount factors near one) apply to both, and the material in section 6.5 is very similar to that in Chapter 9.

The Theory of the Firm chapter, in the author's words, 'is at worst a hodgepodge of received ideas; at best it is a road map of the relevant contributions'. It is certainly more of the latter, and contains a lot of good stuff—the discussion of inefficiencies caused by unconstrained bargaining, and the principal–agent treatment, to cite but two examples. Also, the exposition is characteristically lucid. Still, the chapter is tough going. The main culprit is the unsettled state of the area; another is the broad range of issues covered and, as Tirole warns, the fast clip at which this is done. The chapter could be split into two, one focusing on what is a firm and the other on the profit-maximization hypothesis. The current treatment of the two questions is fairly separated anyhow; unbundling them would allow for a more leisurely development and help readers self-select according to interest. For my taste, the literature on what is a firm is rather confusing at this stage of its development. It addresses diverse themes, whose connections are not always obvious (rather reminiscent of the blind men and the elephant). The

material on profit maximization seems in better shape and introduces useful general ideas (such as the principal–agent framework). Making it a separate chapter would also permit expanding the discussion of important, interesting topics such as corporate takeovers that fall between the cracks of industrial organization and finance.

Finally, Chapter 8 also is rather long. The themes are intimately linked, but plowing through the vast amount of material gets tiring. One option is to break up the chapter into one focusing on the two-stage game structure (sections 8.2 through 8.5) and another in which timing of investment decisions is endogenous (sections 8.1 and 8.6). Also, some of the material in the supplementary section 8.6.2 is central and simple enough to be moved into the main text—the discussion of pre-emption incentives and of the differences between the two polar timing games, pre-emption and war-of-attrition.

All this, however, is really nit picking. The clue, echoing Sherlock Holmes, lies in the dog's not barking. This book is an amazing accomplishment, and should take pride of place on the shelf of every economist.

MARIUS SCHWARTZ
Department of Economics,
Georgetown University,
Washington DC 20057