

Tirole's Industrial Regulation and Organization Legacy in Economics*

Drew Fudenberg

Harvard University, Cambridge MA 02138, USA
drew.fudenberg@gmail.com

Abstract

Jean Tirole was awarded the 2014 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel for his analysis of market power and regulation. This paper provides an overview of some of that work, and of his related contributions to game theory.

Keywords: Game theory; industrial organization; mechanism design; oligopoly theory; patents; platform competition; regulation

JEL classification: C72; C73; D21; D42; D43; D82; K21; L12; L13; L41; L42

I. Introduction

“Je suis un chercheur; je ne suis pas capable de parler de tout et n’importe quoi.”¹

Jean Tirole was awarded the 2014 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel for his analysis of market power and regulation. In this paper, I try to convey the main ideas of that work, along with some of Jean’s positive and methodological contributions to the study of imperfect competition and its implications for industrial organization. This narrow focus emphasizes my connections with Jean, and leaves out his many important contributions in such fields as asset pricing, behavioral economics, and organizational economics, which on their own would constitute a stellar career, as well as his arguably Nobel-level work on banking and corporate finance.² Jean’s phenomenal energy and breadth are reflected in the fact that he has distinct and influential collaborations with

*I thank Glenn Ellison, Mira Frick, Ryota Iijima, Eric Maskin, Lones Smith, Patrick Rey, Jean-Charles Rochet, Al Roth, Nathalie Tirole, and Jean Tirole for helpful comments. NSF grant SES-1258665 provided financial support.

¹Jean explaining why he will not offer opinions on arbitrary topics: “I am a researcher, and not able to talk about everything and anything.” (See <http://www.canalplus.fr/c-divertissement/c-le-petit-journal/pid7560-vu.html?vid=1148121>.) While quite a reasonable position, Jean’s areas of competence are much larger than those of most economists I know.

²For a broader and more detailed overview of Jean’s work, see the excellent scientific report compiled by the Economic Sciences Prize Committee (Royal Swedish Academy of Sciences, 2014); I relied on it extensively when writing this paper. Another sign of what this paper

each of Philippe Aghion, Roland Benabou, Mathias Dewatripont, Oliver Hart, Bengt Holmstrom, Paul Joskow, Jean-Jacques Laffont, Josh Lerner, Eric Maskin, Patrick Rey, and Jean-Charles Rochet. Any of them could have been asked to write this paper, and each would have their own take on Jean's story; I would like to share a little of my own.

I first met Jean in 1978 when we both started graduate school at Massachusetts Institute of Technology. At the time, theorists such as Dasgupta, Dixit, Spence, and Stiglitz were exploring the implications of dynamic issues, such as commitment and timing for such industrial organization topics as patent races and preemptive investment.³ Simultaneously, game theorists were developing tools that seemed natural for studying these problems, such as sequential equilibrium (Kreps and Wilson, 1982), and showing how to use formal game theoretic models to study dynamic competition, as in Milgrom and Roberts (1982) on limit pricing. Neither game theory nor mechanism design was then a standard part of the economics curriculum, but Eric Maskin taught an advanced class on these topics.⁴ Just as important, Eric spent a lot of additional time with us in a reading class where we read a number of soon-to-be classic papers, including all of the game theory papers mentioned above. Jean's attraction to game theory, and particularly dynamic games, was quick to take root. During graduate school, he analyzed the strategic aspects of capacity expansion, learning-by-doing, and bargaining with incomplete information (Fudenberg and Tirole, 1983a,b,c), and began working on the dynamic oligopoly models that became Maskin and Tirole (1987, 1988a,b).⁵

Since then, Jean has been a leader in applying game theory and mechanism design to analyze how firms set prices, make investment decisions, etc., and how to design rules and regulations that lead to better outcomes. As the Nobel committee noted, "No other scholar has done more to enhance our understanding of industrial organization in general, and of optimal policy interventions in particular." He did this by identifying important economic problems, developing and extending the appropriate game theoretic tools, and applying them to derive important conclusions

leaves out is the fact that the 2015 Nemmers Prize conference in his honor, "Liquidity, Bubbles, and Crises", focused on Jean's work on banking and financial markets as opposed to my focus on regulation and industrial organization.

³ See, for example, Spence (1977) and Dixit (1980) on strategic capacity expansion, and Dasgupta and Stiglitz (1980) on patent races.

⁴ The syllabus for the class was the minmax theorem, Nash equilibrium (including existence in discontinuous games, and applications to implementation theory), the core, Shapley value, bargaining set, and other cooperative solution concepts.

⁵ In addition to all of this, Jean found time to write a paper on the impossibility of speculation with rational expectations (Tirole, 1982) and a paper on fixed-price equilibria (Maskin and Tirole, 1984).

and results. After Jean moved to Toulouse in 1991, he and Laffont began working with Electricité de France and France Telecom, and those connections would start a fruitful feedback loop between research and its application. However, it is worth bearing in mind that, in Jean's case, the research interest came first.

The move to Toulouse brings up another major part of Jean's contribution to economics, namely helping to build an excellent economics group in a French university. The first step towards this program was the establishment of the Institut d'Economie Industrielle (IDEI), which Jean helped Laffont found in 1990. Jean has played a lead role in the growth of the IDEI, and in the creation and success of the Toulouse School of Economics, which has become one of the best economics groups in Europe. This institution-building has many positive externalities for the profession, and Jean is rightly very proud of it. The time and effort involved make his research accomplishments all the more impressive.

Finally, while the Nobel Prize is not awarded for textbooks, Jean's contribution in this domain are far-reaching and important. To quote the Prize Committee, "After more than 25 years, his ground-breaking 1988 textbook remains the best road-map to the field. If the book is becoming outdated in a few areas, this is largely due to Jean's own subsequent work and the work he has inspired. Jean's 1993 book, co-authored with Jean-Jacques Laffont, presented a unified framework which has deeply influenced how economists think about regulation." These are only two of Jean's 12 books; I say more about a few of them below.

II. Regulation as Applied Mechanism Design

The Regulation of Monopolies

Many governments use regulations to reduce the distortions that would be caused by monopoly pricing, and there is a long tradition of work that asks what sorts of regulations should be used. Ramsey (1927) and Boiteux (1956), among others, studied how to set prices to maximize total social surplus (i.e., consumer surplus plus firm surplus plus taxpayer surplus, ignoring distributional concerns) given a break-even constraint in static settings where the regulator knows both the production technology and the demand function. Here, the optimal policy is to set the proportional mark-up $(p_i - c_i)/p_i$ equal to θ/η_i , where η_i is the elasticity of demand of good i and $\theta \in [0, 1]$ reflects the tightness of the break-even constraint or the social cost of providing subsidies. Thus, there is less price distortion on goods where mark-ups cause more deadweight loss, and the ratio between the regulated mark-up and the monopoly mark-up is the same for every good. In practice, though, regulators might not know the firm's production

cost, and this sort of cost-plus pricing does not provide the correct incentives for cost-reducing investment. This might be why regulated firms have often been governed by “rate of return” regulation, which sets prices so that the firm earns a “reasonable” return on its investments.

Some of the early work on the theory of regulation, such as Averch and Johnson (1962), analyzed how to set the parameters of commonly used regulatory rules. This did provide some insights on the expected distortions that these rules cause, but because these rules were taken as given, without an explanation of when or why they would be used, the analysis was incomplete. Might there be other simple and practical rules that addressed the underlying incentive problems and led to a better outcome?

Baron and Myerson (1982) pioneered the modern study of regulation – analyzing the fully optimal direct mechanism in a setting with all of the objectives and constraints made explicit – in their analysis of regulation with adverse selection, where the monopolist’s cost is unknown but fixed and cannot be altered by the firm’s actions.⁶ Sappington (1982) allowed for moral hazard in addition to adverse selection, supposing that the firm can carry out unobserved research to reduce its production cost, but restricted attention to linear contracts.

Using Cost Information to Regulate Firms. Like Sappington (1982), Laffont and Tirole (1986) suppose that the monopoly firm is subject to moral hazard as well as adverse selection, allowing for either perfect or noisy observation of costs. They then use the revelation principle to characterize the fully optimal menu of contracts. Myerson (1982) had shown that the revelation principle applies in these problems, but did not explicitly analyze one; Laffont and Tirole (1986) appear to be the first to have provided an explicit solution to one.

In the simplest version of their model, a firm – identified with a single agent – produces a single output q at cost $c = (\beta - e)q$, where $e \geq 0$ is the firm’s effort, and β is the firm’s efficiency parameter, which is known to the firm but not to the regulator.⁷ (Laffont and Tirole (1986) allow cost to also depend on a mean-zero cost shock that is independent of

⁶ Vogelsang and Finsinger (1979) and Loeb and Magat (1979) had previously incorporated asymmetric information into models of regulation, but each made key simplifications that limited the impact of their work. Vogelsang and Finsinger (1979) studied a dynamic process of rate adjustment when the firm responds completely myopically, and Loeb and Magat (1979) assumed that there is no social cost to the rents of the firm; here, it is easy to implement the efficient pricing rule by paying the firm the whole consumer surplus.

⁷ The identification of the firm with a single agent is a standard assumption in most of the regulation and industrial organization literature, which ignores the internal structure of the firm. However, see Chapter 7 of Tirole (2006) for a discussion of how that structure can interact with behavior in the product market.

β and does not depend on q and e , but this makes no difference for most of the analysis because both the firm and the regulator are risk neutral and transfers can be arbitrarily large.) The good provides consumer surplus $S(q)$, and to simplify, the paper assumes that the good is not marketed but provided at no cost to consumers.⁸ The planner observes and reimburses the cost c incurred by the firm, and pays, in addition, a net monetary transfer t so that the total payment is $c + t$. The utility level of the firm's manager is then $U = t - \psi(e)$, where ψ measures the disutility of effort.

The regulator's objective is to maximize the sum $S + U$ of consumer and producer surplus, minus the social cost $(1 + \lambda)(c + t)$ of the funds transferred to the firm, where $\lambda > 0$ reflects the shadow cost of the distortionary taxation. Thus, the regulator would prefer to make small transfers, but must balance this with concerns for efficient effort and socially optimal provision of the good. The regulator knows all of the parameters of the model, except for the firm's type β ; it can set the output level q and it will observe the firm's cost c .

Using the revelation principle, the regulator's problem is to design a direct mechanism under which each report $\hat{\beta}$ is assigned an effort level $e(\hat{\beta})$, anticipated per-unit cost $c(\hat{\beta})$, and output $q(\hat{\beta})$ that satisfy the incentive-compatibility constraints that each type of firm prefers to announce truthfully and to take the assigned effort level, and that also satisfy the interim individual rationality constraints that all types are willing to participate in. Under some technical assumptions,⁹ Laffont and Tirole (1986) showed that it is sufficient to consider only the local incentive constraint. From there, they proceeded to show that the optimal mechanism can be implemented by having the firm pick from a menu of contracts, where instead of specifying a quantity, effort level, and transfer, each contract specifies a quantity, a lump-sum payment, and the share of the realized costs that the regulator will reimburse. Moreover, the menu is such that the most efficient firm (type $\underline{\beta}$) chooses a fixed-price contract without any cost sharing at all, so that it chooses the cost-minimizing level of effort given its assigned production. As in Mirrlees (1971) and subsequent work, when only the local incentive constraints bind, there is "no distortion at the top". The planner could assign such contracts to all types, and indeed would have

⁸ In a footnote, Laffont and Tirole (1986) explain that modeling private sales just requires adding the social value of the associated revenue to the objective function. This case was studied in the working paper version of Laffont and Tirole (1986), and is explored at length in Laffont and Tirole (1993).

⁹ All of the functions above, as well as the density of types, are continuously differentiable, S is concave, ψ is convex, and the density satisfies the monotonic hazard rate property. Additional assumptions are used to ensure that it is not optimal to shut down the firm or to send the marginal cost of production to zero, that the full information problem is sufficiently convex, and that the first-order condition for the planner's maximization is sufficient.

to do so if costs were not observed, as in Baron and Myerson (1982), but this would increase the size of the payments needed to ensure truthful reporting. Instead, the optimal contract provides cost sharing and lower lump-sum payments to higher types. Thus, in general, the effort level is less than the optimum, costs are higher, and production is lower than with perfect information. In addition, the optimal contract moves toward a fixed-price contract when demand increases, because the importance of per-unit cost reduction is higher when more is produced.

In subsequent work, Laffont and Tirole extended this model to analyze a range of related regulatory issues, still within the static framework. Laffont and Tirole (1990a) and Laffont and Tirole (1993, Chapter 3) study multi-product firms, and show that when cost satisfies a separability condition, the relative prices obey the Ramsey–Boiteux formulae, so that the regulator does not distort the price structure to try to extract rents; rent extraction is addressed through the regulation of the rate of return. This result provides a foundation for the widespread use of regulations that impose aggregate caps but let the firm allocate its constrained exercise of market power according to demand elasticities. Laffont and Tirole (1993, Chapter 4) extend the model to the (indirect) regulation of quality when quality is not verifiable and so cannot be included in the terms of the mechanism. Laffont and Tirole (1990c) and Laffont and Tirole (1993, Chapter 6) study “cream skimming” by a firm that faces two different markets or types of consumer.

With hindsight, given Jean's later work on platform competition and two-sided markets, the most evocative static extension might be Laffont and Tirole (1993, Chapter 5) and Laffont and Tirole (1994) on access pricing. This involves a regulated multiproduct firm, some of whose products (e.g., access to a network of optical or electricity cables) are sold to other firms, who then compete with the regulated firm for consumers. If the firm's cost of providing access to others is the same as the cost of using the network itself, the only way the firm could discourage access is by saying that the cost of the network is high. This would lead the regulator to reduce planned supply, which reduces the firm's incentive to exaggerate costs. However, the firm would be very keen to exaggerate any incremental cost that is needed solely to provide access to others. Typically, the regulator will choose to only partially offset this incentive, so the optimal mechanism will tend to have higher access prices and less competition than with full information.

The Dynamics of Regulation. While in Laffont and Tirole (1986) and much of Laffont and Tirole (1993) they studied a static problem, in Laffont and Tirole (1988, 1990b) and Laffont and Tirole (1993, Chapters 9 and 10) they studied dynamic versions of their basic (1986) model, with the simplification that output is simply 0 or 1 (as if demand is very inelastic, or if the firm is the sole plausible supplier of an indivisible public good,

such as a bridge or train line). Baron and Besanko (1984) had shown that when the regulator can commit to future rules, and the firm's type is the same in every period, the optimal dynamic scheme is simply to commit to enforcing the optimal static scheme in each period, much as the optimal pricing policy for a monopoly seller of a durable good is to commit to a fixed price (Stokey, 1981).¹⁰ In practice, such commitment can be difficult to achieve, which motivated Laffont and Tirole to consider other scenarios. Laffont and Tirole (1988) suppose that the regulator is only able to offer short-term contracts, so that not only can the regulator not commit, but also the firm cannot either, because it has the option of rejecting the current contract. The regulator's lack of commitment gives rise to a ratchet effect (as in Weitzman, 1976): when determining the contract to offer in the last period of the relationship, the regulator will use all of the information he has acquired previously, but this makes the firm less willing to reveal information in earlier periods.

Freixas *et al.* (1985) considered this sort of problem in a pure adverse selection setting with a restriction to linear contracts. Laffont and Tirole (1988) considered the equilibria with an unrestricted contract space in a two-period version of the cost-observation model of Laffont and Tirole (1986). Here, the static optimum cannot be implemented, because in addition to the usual problem in static models of deterring the more efficient types from claiming to be inefficient, the regulator needs to worry that a less efficient firm will claim to be efficient, collect a large lump sum (which is the optimal contract for an efficient agent in the static model), and then refuse the second-period contract. That is, the incentive constraints can bind in both directions, which makes them difficult to characterize. Moreover, no period-one mechanism can induce full revelation of types – there will always be some degree of pooling, so after period one the regulator will, in general, not know the firm's true type.¹¹ Laffont and Tirole show that when there is little uncertainty about the firm's type (in the sense that the type distribution is a small interval), then either almost all types pool together in the first period (so that the lack of commitment leads to an almost complete breakdown of the information revelation obtained in the static case) or the equilibrium takes a very complex, non-partition form, which seems unlikely to occur in practice.

¹⁰ They also considered the case where the type can change from period to period, a topic that has attracted increased attention in recent years.

¹¹ If there were a separating equilibrium, a less efficient firm, which claims to be more efficient than it is in the first period, would choose to exit in the second period, while a more efficient firm, which claims to be less efficient, would earn a positive second-period surplus. The conclusion then follows from the constraints that neither firm wants to announce the other's type.

In Laffont and Tirole (1990b), the regulator and firm can enter into a long-term contract, which either of them can enforce, so the regulator can commit not to use first-period information to extract all of the firm's second-period rents, and the firm can commit itself not to "take the money and run". However, the two parties are unable to jointly commit not to tear up the contract and sign a new one. That is, the contract must be "renegotiation-proof", in the sense of Dewatripont (1989). Here, there are contracts that lead to the full separation of types, but this is not optimal for the regulator, and the best contract involves some pooling. Laffont and Tirole (1990b) establish this result when the prior is a smooth density over a continuum of types, but the detailed analysis is restricted to the two-type case. Here, the second-period contract is conditionally optimal – meaning that it is optimal in the continuation game that starts in the second period – although the rent, given the good type, depends on second-period beliefs and hence on first-period play. In the first period, the principal offers two contracts, one without any cost sharing (as in the contract for the most efficient type in Laffont and Tirole, 1986). The bad type accepts the other contract (which does have some cost sharing) and the good type randomizes, so that the equilibrium is semi-separating: the acceptance of the fixed-price contract reveals the firm's type, but the acceptance of the contract with cost sharing does not. Thus, as in the case of short-term contracts only, there is some pooling of types to reduce the regulator's ex-post incentive to extract rents from the firm.¹²

Regulatory Capture. Stigler (1971) pointed out that regulatory agencies might be "captured" by those most affected by their regulations, and showed that trucking regulations in the 1930s seemed to have been heavily influenced by the competing railroad industry. Peltzman (1976) took a step towards formalizing the model, but left the behavior of consumers and firms exogenous. He did not incorporate the informational asymmetries and other constraints that make the design of regulations non-trivial and allow regulated firms to obtain the rents, and thus give them an incentive to distort the process. To model regulatory capture, Laffont and Tirole (1991) and Laffont and Tirole (1993, Chapter 11) combined their basic cost-observation model with the three-tier structure that Tirole (1986) introduced to model collusion in general hierarchies. Here, as in the basic model, the firm has cost function $C = (\beta - e)q$, where β is private information, and has only two possible values, $\underline{\beta} < \bar{\beta}$. Congress takes the role

¹² For the same reason, the optimal renegotiation-proof contracts are semi-separating in the Hart and Tirole (1988) model of renegotiation-proof contracts for the repeated rental of a durable good, and in the Fudenberg and Tirole (1990) model of renegotiation-proof contracts, when the agent faces moral hazard, and contracts can be renegotiated after the agent has made the effort decision but before all of its effects are realized.

of the regulator in the base model, and has the same payoff function, and observes the firm's realized cost. In addition, there is a supervisor who, with some probability, learns the firm's cost parameter β . The supervisor then sends a report of its information to Congress. Importantly, when the supervisor does receive a signal, it is verifiable, as in Grossman (1981) and Milgrom (1981); the supervisor cannot falsely report that the firm's type is high ($\bar{\beta}$) when it is low (β). However, unlike in those papers, there is a chance that the supervisor is uninformed (receives a null signal), and the supervisor can falsely claim to be uninformed.¹³ In the absence of collusion with interested parties, none of this matters, and the regulator reports truthfully. However, if the supervisor learns that the firm's cost is low, the firm has an incentive to bribe the supervisor to suppress that information. Foreseeing this, Congress knows it must pay the supervisor a bonus when the supervisor reports that the type is low. Because the size of the bonus depends on the low type's gain from reporting high, the optimal mechanism when these bribes are possible will sacrifice efficiency to reduce this gain, and thus to reduce the transfers needed to generate honest reporting by the supervisor. In particular, the optimal regulatory scheme now further distorts the effort and output of the inefficient firm by using a contract with lower powered incentives (i.e., more cost sharing, so that the transfer is less responsive to the outcome), as in Tirole (1986). This can be seen as a defense or explanation for the, perhaps surprising, prevalence of such contracts.¹⁴

Regulation of Oligopolies

Vertical Restraints. Vertical relationships between firms, where one firm sells to another, often use contractual restrictions, such as resale price maintenance or exclusive territories. These restrictions have sparked a large number of anti-trust cases, and a sizable body of economic literature.

As with the analysis of price regulation, Jean's approach was to explicitly model the contracting problems that prevent participants from obtaining their preferred outcome. In Rey and Tirole (1986), a monopoly upstream

¹³ In practice, supervisors and regulators cannot simply choose to not send a report, so reporting the null message should be thought of as leaving some relevant information out of the report. In Grossman's signaling model, an unraveling argument shows that failure to report will be interpreted as revealing the worst possible information; this conclusion extends to sufficiently small probabilities that the sender is uninformed.

¹⁴ The extent of this distortion depends on the probability that the supervisor is uninformed, and vanishes as this probability goes to zero. Laffont and Tirole also consider an extension of the model with "interest groups" who can spend resources to influence the supervisor's report, and show that the ability to do so can make the groups worse off when it is foreseen by the mechanism designer.

firm sells to a number of geographically dispersed retailers. The contracting problem arises because the (location-specific) cost and demand are uncertain, and retailers observe these local shocks but the monopolist does not. In addition, the monopolist is unable to observe actual sales of each retailer, so to avoid arbitrage it must set the same variable price to all of them.¹⁵ In this case, the monopolist can effectively only use two-part pricing.

In the absence of vertical restraints, the resulting Bertrand competition leads the retail price to be independent of local shocks to demand, and to respond fully to local cost shocks, so that the equilibrium franchise fee (the fixed part of the monopoly price) is zero. Rey and Tirole show that the monopolist does better by assigning each retailer an exclusive region in exchange for a positive franchise fee, provided that the retailers are not too risk-averse. This allows the retail price to respond to demand shocks, and to respond less to cost shocks, so that more consumer surplus can be extracted. When retailers are risk-neutral, the monopolist gains from exclusive dealing, even though it lowers social welfare; the same is true with linear demand when retailers are not too risk-averse. This shows that the private optimality of exclusive dealing is not *prima facie* evidence that it is socially desirable.

Instead of asymmetric information, Hart and Tirole (1990) consider the problems posed by the upstream monopolist's inability to make binding commitments. They assume a known demand curve, so that with full commitment, the monopolist can obtain the monopoly profit with a "forcing" contract, which specifies the quantity that each retailer will sell, charging them a lump sum equal to their earnings. Hart and Tirole's insight is that this contract is not an equilibrium if the monopolist can offer secret price cuts or otherwise renegotiate with any individual retailer, because the two of them can extract profit from the others by slightly increasing that retailer's output.¹⁶ However, if the monopolist can credibly commit to an exclusive dealing arrangement with a single retailer (which was ruled out by assumption in Rey and Tirole, 1986), or can vertically integrate with one of the retailers, then the problem with secret side-deals goes away, and the upstream firm receives the monopoly profit – as it does if the government enforces a rule that the monopolist cannot discriminate between retailers.

¹⁵ The monopolist also wants (or is required) to sell to all of the retailers, so it cannot auction off rights to all of the markets to one retailer.

¹⁶ The easiest way to see this is perhaps to suppose that the monopolist does not actually produce anything, but is a mediator who is empowered to set output limits for each retailer. The retailers can collude on the monopoly output if the mediator can make binding commitments, but without commitment power, the collusive agreement would unravel to Cournot competition.

Telecommunications. In the 1960s and 1970s, telecommunications in the US and many other countries was run by either a private but regulated monopoly or a state-run firm. However, since the early 1980s, the trend has been towards privatization and oligopoly. Because telecoms need to cooperate to send messages from one network to another, this raises a number of new regulatory issues related to access charges (e.g., when one firm has a monopoly on a trunk or intercontinental line), and to interconnection or termination fees (when a consumer calls a customer of another network). Laffont *et al.* (1998a, 1998b) developed a formal model of two-way access, based on the assumptions that the receiver does not pay for calls (the caller's company pays termination fees);¹⁷ Laffont *et al.* (2001) extend this to allow charges to both sides. Laffont and Tirole (2000) consider the optimal regulation of termination and access fees, and argue that the consumer side of the market is best regulated with price caps.

Joint Marketing and Patent Pools. Suppose that n firms each have one patent, which they could sell separately. When, and under what conditions, should they be allowed to form a patent pool that sells all of the patents as a single bundle? Shapiro (2000) pointed out that the static Cournot model implies that pools are anti-competitive (i.e., lead to higher prices and lower sales) when the patents are perfect substitutes, and pro-competitive (i.e., lower prices and higher sales) if the goods are perfect complements. Lerner and Tirole (2004) generalize this analysis to other specifications of utility and demand. In the base version of their model, the patents are symmetric, and the consumer of type θ derives benefit $V(m) + \theta$ from owning m licenses, with V at least weakly increasing.¹⁸ Here, perfect complements correspond to the case where there is no benefit to owning any fewer than all of the licenses, while perfect substitutes can be defined as $V(n) = V(n - 1)$ because this is enough for Bertrand competition to lead to zero price. However, as Lerner and Tirole point out, when V is strictly increasing, patents will be complements at very low prices, because any user who buys one patent will want to buy all of them, and yet patents

¹⁷ Armstrong (1998) independently developed a similar model.

¹⁸ Perhaps surprisingly, the fact whether V is concave or convex plays a very small role in the analysis, although Proposition 4 shows that it does matter for the existence of asymmetric equilibria. The assumed additive separability between consumer type and the number of patents means that the pool will only offer the package of all patents, as opposed to using menus of various sizes to price discriminate; it also leads to the clean characterization of the Bertrand equilibrium when the firms set prices independently. Later sections of Lerner and Tirole (2004) maintain this assumption but consider asymmetric patents, patent holders who can also be licensees, and the impact of the pool on innovation. Boutin (2014) and Lerner and Tirole (2015) extend the analysis to asymmetric equilibria and more general value functions.

will be substitutes when prices are high, because users will only buy some of them so that the patents are in competition.

To analyze the effect of a patent pool, Lerner and Tirole (2004) compare it to the symmetric pure-strategy equilibrium of the Bertrand pricing game, which they show is unique. Generically, a firm that is considering raising its price will only be concerned with one of two effects: either the “competition margin binds”, meaning that the consumers purchase from some other firm instead, so that total demand is constant, and market share changes; or the “demand margin binds”, meaning that the consumers who purchase any patent buy all of them, and increasing the price lowers total sales without changing the market share. Because the firms do not internalize this effect of price on overall demand, patent pools lower prices and increase welfare whenever the demand margin binds, as they do in the case of perfect complements. In contrast, when the competition margin binds, whether the pool would raise or lower welfare depends on whether the equilibrium price (which depends on the properties of $V(m)$ for $m < n$ but not the distribution of θ) is above or below the pool's price (which depends only on $V(n)$ and the distribution of θ). Because regulators might not know enough about the demand side of the market to know which case applies, Lerner and Tirole (2004) propose the following elegantly simple policy: allow firms to collaborate on the pricing of bundles, but require that they also be free to offer standalone licenses. In this way, firms will only be tempted to offer a standalone if the pool leads to higher prices.

Rey and Tirole (2013) explore the robustness of this result to markets where repeated interactions between the firms can facilitate tacit collusion. Here, joint marketing agreements can be welfare-decreasing, even with standalone licenses; however, the agreements are socially beneficial when unbundling is imposed (i.e., when the pool is required to offer the individual elements at prices whose sum cannot exceed that of the bundle).

Financial Markets. Jean has written extensively on the regulation of banks and financial markets; I summarize that work only very briefly here. Dewatripont and Tirole (1994) used a contracting framework to analyze optimal regulation to protect depositors from excessive risk taking by managers. In later work, Jean focused instead on the fact that the interconnections of financial institutions means that a bankruptcy of one firm can have negative impacts on other firms. Rochet and Tirole (1996) used the Holmström and Tirole (1998) model of intermediation and liquid assets to study how government insurance of deposits reduces the incentives of banks to monitor each other, and to study what sorts of regulations might help restore them. This anticipated many of the problems that would hit the financial system in the 2008 crisis. Farhi and Tirole (2012) showed that financial institutions might use too much short-term debt, and correlate their

investment returns, more than is socially optimal; perversely, banks like this correlation if bailouts only occur when many banks are threatened at once. Farhi and Tirole show that optimal policies combine liquidity requirements with a commitment to lower interest rates in a crisis. As in other mechanism design problems, the optimum sacrifices some ex-post efficiency to increase ex-ante performance.

III. Industrial Organization

Strategic Investment

The study of entry and deterrence has a long history in the study of oligopolistic markets. In order to decide whether or not to enter a market that is currently dominated by an established firm, a potential entrant must consider how that firm would respond if entry occurred. As the Economic Sciences committee noted, before its reconstitution on the foundation of game theory, the industrial organization literature used various ad hoc assumptions about this response. For example, if the potential entrant believes that entry will not cause an incumbent firm to change its output, the incumbent can prevent entry by choosing a pre-entry output level just high enough that entry would be unprofitable. Of course, if the entrant anticipated that lower post-entry prices would induce the incumbent to reduce its output post-entry, then entry might nonetheless occur, and indeed, in typical settings, the incumbent's implicit threat to maintain output after entry is not credible. Spence (1977) pointed out that (in the absence of substantial depreciation) having a substantial pre-entry capital stock is a credible commitment to have a large capital stock after entry. So, if the incumbent firm's capital level lowers the profitability of entry, the former might want to over-accumulate capital. To complete the model, Spence assumed that post-entry output would either equal capacity or be large enough that the price would equal the incumbent's marginal cost. Dixit (1980) made the more plausible assumption of Cournot equilibrium (i.e., Nash in quantities) in the second period, and so was the first to present a model of entry deterrence in a subgame-perfect equilibrium. Here, the incumbent's first-period capital stock lowers the payoff to entry because it lowers the incumbent's marginal cost, and so increases the incumbent firm's post-entry output.

Mobility Deterrence. Spence (1977) and Dixit (1980) modeled entry as a discrete event, but as Caves and Porter (1977) argued, entry is often accomplished in stages, so strategic investment can be used not only to deter entry but also more generally to deter mobility. Because mobility deterrence is an ongoing process, it is best modeled as a game with an

infinite or indefinite horizon; the equilibrium to this game can then be used as the second period to consider the decision to enter and strategic pre-entry investment. Spence (1979) analyzed this strategic investment in a continuous-time model, where flow payoffs are determined by Cournot competition given the marginal costs associated with the capital stocks $(K_1(t), K_2(t))$, there is no depreciation so that $dK_i(t)/dt$ is firm i 's investment $I_i(t)$, and the cost of investment is constant up to a fixed upper bound; he concluded that the equilibrium is for the lead firm to invest as quickly as possible to some capital level, and then to stop. Because this level was chosen knowing the follower's response, the result was much like the equilibrium in a static Stackelberg game, as opposed to the Cournot–Nash levels that would arise if the two firms simultaneously chose steady-state capital levels.

Fudenberg and Tirole (1983a) re-analyzed the Spence (1979) model more formally.¹⁹ As with many student papers, our original impetus was simply to correct a mistake, but as we worked on the model, we also saw that there was a continuum of “early-stopping” equilibria that Spence had overlooked. Here, both firms stop investing before the Spence levels are reached, and make higher profits – the follower firm is permanently deterred from investing up to its reaction curve by the leader's credible threat to invest up to the Spence level if the follower “steps over the line”.²⁰ As in the differential games literature, our analysis assumed that each player's investment at time t depended only on the payoff-relevant state $(K_1(t), K_2(t))$, and that it is optimal for each firm to do when the other one does. Thus, the outcomes we derived were also Markov perfect equilibria (MPE) in the sense of Maskin and Tirole (1987), and one implication of the paper is that games can have multiple MPE, and that these can be Pareto-ranked.²¹

Fat Cats and Puppy Dogs. The first wave of work on strategic investment focused on the incentive to invest more in productive capital to deter or reduce the scale of entry. Subsequently, Schmalensee (1982) developed a model of informative advertising in which entry deterrence always involves advertising less than if the threat of entry were absent. His work led

¹⁹ This was Jean's first paper to use game theory. His first publication, Tirole and Guesnerie (1981), came from the work he did in France for his master's degree.

²⁰ In the early-stopping equilibria, the value functions are discontinuous in the state. Iijima and Kasaharay (2015) show that value functions are continuous when the state is subject to Brownian motion, and they use this to show that there is a unique equilibrium in the finite-horizon stochastic differential version of the game.

²¹ Note that the assumption that the flow payoffs depend only on the current capital stock is itself a form of Markov restriction. The formal definition of state space equilibrium is not equivalent to MPE, because the former did not define the more general class of strategies where actions can depend on uncountably many past actions. Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) develop ways to handle the technical difficulties involved.

Fudenberg and Tirole (1984) to develop a taxonomy of factors that lead to over- or under-investment, both when entry will be deterred and when it will not.

In their model, and in the similar and independent work of Bulow *et al.* (1985), the incumbent (firm 1) moves first and chooses a capital stock K_1 , which could, for example, be productive machinery, advertising, or investment in R&D; the key is that the level of K_1 is fixed and will not change if entry occurs. Firm 2 sees K_1 and then chooses whether or not to enter. If it does, the two firms simultaneously choose actions (x_1, x_2) , which could be quantities, prices, or further R&D spending.²² If no entry occurs, firm 1's payoff is $\pi^1(K_1)$ (which incorporates the optimal choice of x_1 by a monopolist) and firm 2's payoff is 0. Under entry, the two firms' payoffs are $\pi^1(K_1, x_1, x_2)$ and $\pi^2(K_1, x_1, x_2)$, where we expect π^2 to be independent of K_1 in some cases (such as the Spence–Dixit investment model) but not in other cases (e.g., when K_1 represents pre-entry advertising). In this setting, a subgame-perfect equilibrium is a triple (K_1^*, x_1^*, x_2^*) , where (x_1^*, x_2^*) are functions of K_1 , and $(x_1^*(K_1), x_2^*(K_1))$ is a Nash equilibrium for each value of K_1 . This second-period equilibrium is assumed to be unique, and continuously differentiable, and both it and firm 1's choice of K_1 are assumed to be characterized by the first-order conditions. We also assume that entry is not blockaded; that is, the capital stock K_1 that firm 1 would choose in the absence of firm 2's entry threat would lead firm 2 to enter.

Then, firm 1's problem is to choose K_1^* to obtain the larger of the best entry-detering payoff,

$$\max_{K_1} (\pi^1(K_1) | \pi^2(K_1, x_1^*(K_1), x_2^*(K_1)) \leq 0),$$

and the best payoff with entry accommodation,

$$\max_{K_1} (\pi^1(K_1, \pi_1^*(K_1), x_2^*(K_1)) | \pi^2(K_1, x_1^*(K_1), x_2^*(K_1)) \geq 0).$$

To determine the best way for firm 1 to deter entry, consider the total derivative of π^2 with respect to K_1 ,

$$\frac{d\pi^2}{dK_1} = \pi_{K_1}^2 + \pi_1^2 \frac{dx_1^*}{dK_1},$$

where subscripts K_1 , 1, and 2 denote partial derivatives with respect to K_1 , x_1 , and x_2 , respectively, and $\pi_2^2 = 0$ from the envelope theorem. Here, the term $\pi_{K_1}^2$ is the direct effect of K_1 on firm 2's profit, and $\pi_1^2(dx_1^*/dK_1)$

²² One might expect both investment and production decisions to be made post-entry, as opposed to the one post-entry decision in the model here. The key for strategic investment to have a role is that pre-entry and post-entry investments are not perfect substitutes.

is the strategic effect – the effect that comes from the fact that pre-entry changes in K_1 will change firm 1's post-entry behavior. Fudenberg and Tirole say that investment makes a firm “tough” if $d\pi^2/dK_1$ is negative, and “soft” if it is positive; as they point out, this determines whether the optimal entry-detering level is more or less than the levels in the “open-loop equilibrium”, in which firms set time paths of their actions at the start, and are unable to respond in period 2 to period 1 play. Specifically, when firm 1 is tough, the way to deter entry is with a “top dog” strategy of investing more than in open-loop equilibrium. When firm 1 is soft, the way it should deter entry is to invest less than in the open-loop benchmark, because staying “lean and hungry” leads to more aggressive post-entry play. This case can arise in models of advertising. Firms with a large enough captive market might be less willing to respond to entry with price cuts when price discrimination is not possible.

If, instead, firm 1 chooses to allow entry, what matters is the effect of K_1 on firm 1's post-entry profit, and the relevant first-order condition is

$$\frac{d\pi^1(K_1, x_1^*(K_1), x_2^*(K_1))}{dK_1} = \pi_{K_1}^1 + \pi_2^1 \frac{dx_2^*}{dK_1},$$

so that the strategic effect is now $\pi_2^1(dx_2^*/dK_1)$. To relate this to whether investment makes firm 1 tough or soft for the purpose of entry deterrence, suppose that $\pi_2^1\pi_1^2 > 0$, so that the second-period actions of the two firms have the same sort of effect on the other. Then, algebra shows that the sign of the strategic effect here is the same as for entry deterrence if increasing x_i makes firm j want to decrease x_j (so that the second-period reaction curves slope down, and the actions are strategic substitutes in the terminology of Bulow *et al.*, 1985), while the effect has the opposite sign than for entry deterrence if the reaction curves slope up and the goods are strategic complements.

Summing up, when second-period actions are strategic substitutes, and investment makes the firm tough, as in the Dixit–Spence model of strategic investment, the strategic incentive is to overinvest, both to deter entry and to accommodate it (i.e., the top dog strategy). With strategic complements (e.g., cost-reducing investment followed by price competition), the same strategy is optimal to deter entry, but now, to accommodate entry, firm 1 wants to stay a small and friendly “puppy dog” to encourage a higher second-period action by firm 2. When investment makes the firm soft, the best way to deter entry is to stay “lean and hungry”. Staying lean also generates the desired post-entry accommodation response with strategic substitutes, which fits R&D models in the spirit of Reinganum (1983), but with strategic complements, as in the Schmalensee advertising model, when firm 1 accommodates entry, it wants to encourage higher actions by firm 2,

and the way to do this is by becoming a “fat cat”. Tirole (1988, Section 8.4) and Shapiro (1989) develop a number of other applications and extensions of this taxonomy, including topics such as learning by doing, most-favored-customer clauses, multimarket oligopoly, and tying versus bundling.

Price Rigidities and Markov Perfect Equilibrium

While capital stocks are slow to change, and so form fairly long-run commitments, firms also typically face some costs in changing their prices or production plans. These costs might be small, but they can nonetheless have a large impact on the time paths of prices and output, as shown by Maskin and Tirole (1987, 1988a, 1988b). For most of the analysis, Maskin and Tirole consider a game of perfect information where firms alternate moves, as in the finite-horizon model of Cyert and Degroot (1970).²³ However, because Maskin and Tirole drop the (unrealistic) assumption of a commonly known last period, there are a great multitude of repeated-game-like equilibria when players are patient. To rule out these equilibria, Maskin and Tirole introduce the concept of MPE, which refines the set of subgame-perfect equilibria to those in which strategies depend only on the payoff-relevant state, which here is the action chosen by the other firm in the previous period.

In Maskin and Tirole (1988a), firms choose quantities in each period, with firm 1 choosing q_1 in odd-numbered periods and firm 2 choosing q_2 in even-numbered periods, so that Markov strategies correspond to reaction functions $R_i(q_j)$. Firms pay a fixed cost f whenever output is non-zero, and also face constant marginal cost c . Demand is linear, $p = 1 - q_1 - q_2$, so the per-period payoff of a firm with positive output is $q_i(1 - q_1 - q_2 - c) - f \equiv \pi_i(q_1, q_2) - f$, and the fixed cost is such that a monopolist finds it profitable to operate, but two firms would not.²⁴ Maskin and Tirole show that there is a unique symmetric Markov perfect equilibrium. When the discount factor is close enough to 1, this symmetric equilibrium has the following, particularly simple, form. Let q^* be the largest root of $(1 - \delta)\pi(q, q) + \delta\pi(q, 0) = f$. This is the optimal entry-detering output level, because any higher output would yield lower profits, and lower output would not deter entry. Thus, given the two-period commitment (and the restriction to MPE), output acts as an entry deterrent, as in the classic limit-pricing models of, for example, Bain (1956). However, here, as the period length grows shorter so that $\delta \rightarrow 1$, $\pi(q^*, 0) \rightarrow f$, and

²³ Later, they show that, at least in some cases, firms whose actions are locked in for two periods will choose to alternate, and that the model also describes the behavior of firms whose revision opportunities arrive at Poisson intervals.

²⁴ $8f > (1 - c)^2 > 4f$.

so the incumbent gains very little advantage from very short-run commitments, as in the contestability argument of Baumol *et al.* (1982). However, Maskin and Tirole (1988a) also show that when the discount factor and fixed costs are high enough, the game has exactly two other, asymmetric, MPE, where one firm's reaction curve is identically 0 and the other firm takes over the market starting from any state, and that in these equilibria, the incumbent's payoff does not converge to 0, so the conclusion that short-run commitments have little value relies on the symmetry assumption.

Maskin and Tirole (1988b) consider alternating-move duopoly when firms set prices in each period as opposed to quantities, and prices are constrained to lie on a discrete grid. Here, there are both "kinked-demand curve" equilibria and "Edgeworth-cycle" equilibria. In the former, there is a steady-state focal price, which is sustained by strategies that say any price cut will trigger a sequence of further cuts; although the price does eventually return to the steady state after a deviation, the lost profits in the meantime outweigh the gain from deviation. Moreover, when the grid is fine enough and the discount factor sufficiently high, any price in a certain interval can be focal. In the Edgeworth-cycle equilibria, the price never settles down, and instead equilibrium predicts a recurring series of gradual price cuts followed by sharp price increases. These cycles have a long history in industrial organization (Edgeworth, 1925; Hall and Hitch, 1939; Sweezy, 1939), but Maskin and Tirole (1988b) were the first to derive them formally. Since then, Edgeworth cycles have been studied in both field data (starting with Eckert, 2002; Noel, 2007) and laboratory experiments (Morgan *et al.*, 2006).

Innovation and Technology Adoption

In order to study industry dynamics and to design patent, copyright, and related policy instruments, it is very important to understand technology development, adoption, and diffusion. Jean has contributed to this understanding in a variety of different settings.

Patent Races. An extensive body of literature has studied the effectiveness of patents and their impact on market structure and competition. At the end of the 1970s, economists began to explicitly model the dynamics of patent competition, starting with a memoryless model, in which each period's expenditure either leads to a patentable breakthrough or else has no effect.²⁵ Fudenberg *et al.* (1983) modeled R&D as a race to the finish line, where the probability of a discovery at a given date depends on past as well as current expenditures. They show that if progress towards the finish line is a

²⁵ See Reinganum (1984) for a survey of those models.

deterministic function of expenditure, and if firms can very quickly observe and respond to the expenditures of their opponents, then a firm that is even slightly ahead can “coast” at the monopoly pace, safe in the knowledge that it will be able to see and respond to any surges by the others before they can catch up. However, this ε -preemption result needs to be modified if the observation lags are not trivial, as in a discrete-time model with simultaneous moves. Here, if the state is such that it is possible (and not too expensive) for the lagging firm to leapfrog into the lead before the leader notices, the equilibrium prescribes intense but stochastic investment by both firms, which persists until one of them pulls far enough ahead that the other drops out. Leapfrogging can also occur without observation lags if there are multiple, observed, stages to the race, completing a stage is a stochastic function of the time path of expenditure, and a firm must complete the first step before starting the second. Here, a firm that is behind at the start can stay in the race in the hope of leapfrogging the leader, and again a very small lead is not enough to preempt competition, with the most intense competition occurring when the race is the closest.²⁶

Open-Source Software. A substantial amount of software is now written without relying on patent protection. Lerner and Tirole (2002) explain this with a model of career concerns, in the style of Holmström (1999). Contributing to open-source software might serve to signal one’s programming ability, and thus lead to job offers, and might also generate psychic rewards from being acknowledged by one’s peers. Drawing on Holmström’s work, they argue that “the signaling incentive is stronger when the performance of the signaling task is more visible to the relevant audience, and when the outcome of the task is more sensitive to the effort exerted and to the ability of the agent.” As they point out, the need for an audience leads to multiple equilibria: programmers want to work on projects that will attract the attention of many other programmers. Lerner and Tirole also argue that there are three reasons that signaling incentives might be higher in open-source than in proprietary software: (i) it is easier for outsiders to evaluate open-source software because the source code for proprietary software is usually not made public; (ii) open-source programmers are self-employed and take full responsibility for the success of a subproject, so their ability is more tightly linked to the outcome; (iii) open-source programmers are more mobile as they have less firm-specific human capital. Then, they compare their theoretical predictions with case studies of Apache, Linux,

²⁶ Harris and Vickers (1985) obtained ε -preemption in a perfect information model with a much more general R&D technology. Budd *et al.* (1993) and Horner (2004) analyze the more complex behavior that arises when the state of competition is measured by the difference between the number of patents obtained by each firms.

Perl, and Sendmail. They find support for the signaling hypothesis, as did Hann *et al.* (2013), who looked at three different software projects operating within the Apache Software Foundation.

Adoption of New Technologies. When a new technology first becomes available, it can be quite costly to adopt. Over time, adoption becomes cheaper, but waiting too long to adopt puts a firm at a competitive disadvantage. Thus, the adoption of a technology becomes a game of timing. Earlier work by Scherer (1967) and Reinganum (1981a,b) studied the open-loop equilibria of these games, where firms commit to adoption times in advance. Fudenberg and Tirole (1985) considered the subgame-perfect equilibria of such games; most of that paper supposes that there are only two firms, and I maintain that assumption here. In their model, the flow payoffs of the firms are symmetric, and depend on calendar time, and also on the number of firms who have adopted the technology; adopting at time t incurs a time- t cost of $c(t)$. Firms know the current state – how many firms have already adopted – and, at any time, any firm who has not yet adopted can immediately do so. That is, the model is cast in continuous time, so that the observation lags are negligible. Thus, if the first adoption occurs at time t , the other firm faces a single-agent decision problem, and (from subgame perfection) chooses the optimal adoption time in this continuation game. Payoffs are assumed to be strictly quasi-concave, so this follower time is unique; it cannot be before t but it can be infinite as firms are not forced to adopt. Knowing how the follower will respond, it is then possible to compute the continuation payoffs $L(t)$, $F(t)$, and $J(t)$ of the firms as a function of the time of the first adoption, and whether firm i was the leader (the first to adopt), follower, or whether both firms jointly adopted; the assumptions on the underlying payoff functions imply that these continuation payoffs are continuous in the time of first adoption.

Fudenberg and Tirole (1985) consider only the preemption case where $L(t) < F(t)$ for a non-empty interval $[0, \bar{T}_1)$, $L(t) > F(t)$ for a non-empty interval (\bar{T}_1, \bar{T}_2) , and $L(t) = F(t) = J(t)$ for a possibly empty interval $[\bar{T}_2, \infty)$.²⁷ Here, the open-loop equilibria are for one firm to adopt at $T^* \equiv \arg \max_t L(t)$, and the nature of the subgame-perfect equilibria depends on whether the maximal first-mover payoff $L(T^*)$ is more or less than the highest payoff to joint adoption $\max_t J(t)$. In the first case, the equilibrium is unique and involves diffusion (i.e., the two firms adopt at different times) and rent dissipation; the equilibrium path is for one firm to adopt at $\bar{T} < T^*$, where $L(t)$ intersects $F(t)$ from below. The first adopter would prefer to adopt at a somewhat later date if it could still be guaranteed

²⁷ Fudenberg and Tirole (1991b, Chapter 4) give a more complete taxonomy of timing games.

to be the first, but this is not consistent with equilibrium, because any delay would allow the other firm to adopt first.²⁸ Rent dissipation remains an equilibrium outcome when $L(T^*)$ is less than the highest joint adoption payoff, but here there are also delayed joint-adoption equilibria that are better for the firms. In these equilibria, each firm would like to adopt earlier, holding the adoption times of the others fixed, but delayed adoption is enforced by the knowledge that adopting early would lead the other firm to do the same.

The heuristic discussion above ignores a technical issue. While the equilibrium paths described above are intuitive, they require a form of randomization that cannot be supported as subgame-perfect equilibria using the distributional strategy space that had been previously used to model mixed strategies in timing games. This is because that space used a single continuous time mixed strategy to represent the limits of sequences of discrete-time mixed strategies that have very different implications. To formally justify the equilibria, Fudenberg and Tirole constructed a richer strategy space that permits a continuous-time representation of the relevant discrete-time limits.

Two-Sided Markets

Many markets can usefully be viewed as having two sides, whose interactions are mediated by one or more platforms or intermediaries – not only online platforms, such as Ebay or Uber, but also employment agencies, dating services, credit cards, newspapers, and computer operating systems. This leads to a form of network effect. However, unlike in the work of Katz and Shapiro (1985, 1986) and Farrell and Saloner (1985, 1986), which model the externalities between users on the same side of a network (as with, for example, instant messaging programs), the externality is across sides.²⁹

²⁸ As an example, suppose the two firms are engaged in Bertrand competition with a homogeneous good, so that each makes zero profit initially, and adopting the innovation lowers the constant average cost of production from \bar{c} to \underline{c} . Then, once one firm adopts the innovation, the other decides to never adopt, so that $F(t)$ is identically 0, while if the cost of adoption eventually becomes low, L is maximized at the time t where $\bar{c} - \underline{c} = rc(t) - c'(t)$, and the leader earns a rent from adopting. In contrast, in the equilibrium, the leader's profit is zero so the rent is entirely dissipated; preemptive pressure pushes up the time of adoption to the date when it first breaks even. For more on rent dissipation, see Posner (1975) and Fudenberg and Tirole (1987).

²⁹ Caillaud and Jullien (2003) analyze two-sided markets with a continuum of ex-ante identical agents on each side, and preferences that come from seeking a match on the other side. In some two-sided markets, such as dating services or auctions, there can be negative externalities between users on the same side, so that when the populations are finite, a user who moves on to a particular platform makes that platform worse for all users on that

In pioneering work, Rochet and Tirole (2003, 2006) analyzed price-setting by competing platforms when there are positive externalities between the two sides and no externalities between agents on the same side.³⁰ They pointed out that practices, which might look like predatory pricing in a one-sided market, need not be predatory in markets with two sides, as even a monopolist might want to offer very low prices to one side in order to attract more users, and to extract a higher price from the other side.³¹

In their model, agents in the populations $i = S, Y$ (*Sellers* and *buyers*) receive a gross benefit of $B_k^i + b_k^i N_k^j$ from using platform k , where $j \neq i$ is the other side of the market, N_k^j is the number of j -users on platform k , and the membership benefits B_k^i and per-use benefits b_k^i are random variables that are i.i.d. across agents in the same population, and are private information. By assumption, the platforms use a two-part pricing schedule, with fixed charge A_k^i and per-use charge a_k^i . The number of uses on platform k is proportional to the product $N_k^S N_k^Y$, and firms incur marginal cost c per use, as well as fixed costs C^i per user.³²

Rochet and Tirole (2006) show that the monopoly solution to this problem has the following simple form. Let $p^i = a^i + (A^i - C^i)/N^j$ be the firm's average revenue per i -user, net of the fixed costs, and let $p = p^S + p^Y$. Then, the mark-up over the marginal cost of service follows the standard Lerner formula, with $(p - c)/p$ equal to the inverse of the elasticity of demand. The allocation of this total price between the two sides depends on the nature of demand; for example, when the fixed costs and fixed benefits are both zero, and the firm uses only per-use pricing, the mark-up on side i sets $[p^i - (c - p^j)]/p^i$ equal to the inverse elasticity for good i . Here, the effective cost of providing the service to side i is reduced by the extra revenue p^j this brings from users on the other side, and the fact that this effective cost can be negative can lead to low or even negative prices.

side, including himself. This can lead to a great multiplicity of equilibria, as in Ellison and Fudenberg (2003), and subsequent work on when markets tip.

³⁰ Spiegler (2000) considers surplus extraction when agents on each side of the market are identical, and the platforms can use general contracts instead of two-part prices. Here, instead of the trade-off between surplus extraction and surplus maximization studied in Rochet and Tirole, there is an equilibrium in which one of the two platforms appropriates all of the surplus.

³¹ Some advertising-supported newspapers are free, and credit-card users can face a negative price if they pay no annual fees, avoid late charges, and receive cash or merchandise credits.

³² Rochet and Tirole (2003) assumed that the membership benefit and the fixed cost per user are both identically equal to 0; Rochet and Tirole (2006) generalized this to cover the model of Armstrong (2006), but only covered the monopoly case. In their Section 5, Rochet and Tirole (2006) discuss extensions that include allowing agents on one side to care about the quality or type of agents on the other side or to care about the number of agents on their own side.

Rochet and Tirole (2003) go on to consider the case of competing platforms, where agents have a choice between “singlehoming” (using only one platform, such as one credit card or dating site) or “multihoming” and using both. This issue has important implications for the design of regulations and anti-trust policy, and has sparked a sizable body of literature.³³

IV. Game Theory and Mechanism Design

The last element of my selection from Jean’s work is his contributions to other aspects of game theory and mechanism design. Maskin and Tirole (2001) gave a general definition of MPE that shows how to identify the relevant notion of a state from the extensive form of the game.³⁴ MPE have since been widely applied in theoretical analyses in many areas of economics, and used in empirical studies of oligopoly following the work of Pakes and McGuire (1994) and Ericson and Pakes (1995).³⁵

The sequential equilibrium of Kreps and Wilson (1982) says that assessments are consistent if they are the limit of a sequence of totally mixed strategy profiles and associated beliefs. This condition is elegantly stated but is complicated to apply to games of any length. Perhaps for this reason, various authors throughout the 1980s instead used variants of perfect Bayesian equilibrium (PBE) to impose some restrictions on off-path play in dynamic games of incomplete information. Fudenberg and Tirole (1991a) gave a formal definition of this concept in multistage games with observed actions using four quantifier-free conditions such as “no signaling what you don’t know”. They showed that all four conditions are implied by sequential equilibrium, and are equivalent to it when each player has only two possible types. They also showed that the consistency requirements of sequential equilibria require that each player keeps track of the relative probabilities of the types of other players that have posterior probabilities of zero. This requires additional effort on the part of both the players and the analyst, and might not always be plausible.

Inspired by the work of Rubinstein (1982) and Stahl (1972) on non-cooperative models of bargaining with complete information, Fudenberg

³³ In addition to Caillaud and Jullien (2003) and Rochet and Tirole (2003), see, for example, Armstrong (2006), Armstrong and Wright (2007), and Hermalin and Katz (2007).

³⁴ Earlier work by Shapley (1953) and Sobel (1971) had established the existence of MPE in finite stochastic games, where the state is explicitly defined as part of the model.

³⁵ Testing the MPE restriction on field data seems very difficult. In the lab, it is well established that behavior in repeated-game experiments is typically very far from the repeated play of the static equilibrium if players are sufficiently patient, which is what MPE predicts if the subjects act to maximize their expected earnings. (See the discussion and meta-analysis in Dal Bó and Frechette (2014)). In other, more complicated, experiments, MPE do seem to be a good guide to observed behavior, and work such as Vespa (2014) tries to better understand just when this is the case.

and Tirole (1983c) and Fudenberg *et al.* (1985) studied two-person extensive-form bargaining with incomplete information. Fudenberg and Tirole (1983a) showed how incomplete information could lead to bargaining breakdowns and delays in a two-period model. Fudenberg *et al.* (1985) extended the analysis to an *a priori* infinite horizon, and showed that when there is common knowledge of gains from trade (the gap case), bargaining ends in a finite number of steps. They used this to give the first game-theoretic proof of the Coase conjecture (Coase, 1972) that the monopolist is unable to extract a rent when the period between successive offers shrinks to zero, a result that was extended to the stationary equilibria of the no-gap case (and a more general set of type distributions) by Gul *et al.* (1986).³⁶

Maskin and Tirole (1990, 1992) study mechanism design in a principal–agent setting, where the principal has private information and the agent's action is verifiable. Here, as emphasized by Myerson (1983), a key issue is that the proposal of a contract reveals information to the agent, and so can influence both the agent's willingness to accept it, and his behavior if he does.³⁷ Maskin and Tirole (1990) considered the case of independent private values, where the principal's type is independent of the type of the agent, and neither player's payoff function depends on the type of the other. They show that the principal typically earns an information rent, and that the equilibria correspond to the Walrasian equilibria of an economy where the traders are different types of principal, who exchange the slack on the agent's individual rationality and incentive compatibility constraints. Maskin and Tirole (1992) studied mechanism design when both players' payoff functions depend on the type of the principal. They show that the equilibrium set consists of all allocations that weakly Pareto dominate the "RSW allocation",³⁸ and use this to help explain the difference between the Spence signaling model (where the worker chooses effort before the

³⁶ Sobel and Takahashi (1983) analyzed equilibrium in linear strategies when buyer types are exponentially distributed, which is a particular case of no gaps. Ausubel and Deneckere (1989) show that the Coase conjecture fails in the no-gap case, and moreover the limit of the seller's equilibrium payoff can be arbitrarily close to the monopoly profit. Hart and Tirole (1988) show that the same dynamics arise (in the two-type case) when the monopolist makes a repeated series of sales, or when it contracts for the lease of a durable good but cannot commit not to renegotiate the contract.

³⁷ They build on the work of Myerson (1983), which posed the general problem of mechanism design by an informed principal when there are multiple agents, each of which might have some private information, and the work of Holmström and Myerson (1983), which developed the notion of interim efficiency and of "durable" agreements.

³⁸ The RSW allocation is so-called because, in insurance markets, it corresponds to the zero-profit separating allocation that figures prominently in Rothschild and Stiglitz (1976) and Wilson (1977). More generally, it is the allocation that maximizes the payoff of each type of principal within the class of incentive-compatible allocations, which ensure that the

contract is signed) and models of insurance, such as Rothschild and Stiglitz (1976).

Finally, Jean deserves much of the credit for the fact that we wrote our textbook on game theory (Fudenberg and Tirole, 1991b). I would like to think that this book helped to advertise and spread game theoretic methods; without his urging, I would not have thought of writing it. Jean's courage and his vision are the main reasons he has had such a huge impact on economics, both in and out of academia.

References

- Armstrong, M. (1998), Network Interconnection in Telecommunications, *Economic Journal* 108, 545–564.
- Armstrong, M. (2006), Competition in Two-Sided Markets, *RAND Journal of Economics* 37, 668–691.
- Armstrong, M. and Wright, J. (2007), Two-Sided Markets, Competitive Bottlenecks and Exclusive Contracts, *Economic Theory* 32, 353–380.
- Ausubel, L. M. and Deneckere, R. (1989), Reputation in Bargaining and Durable Goods Monopoly, *Econometrica* 57, 511–531.
- Averch, H. and Johnson, L. L. (1962), Behavior of the Firm under Regulatory Constraint, *American Economic Review* 52, 1052–1069.
- Bain, J. (1956), *Industrial Organization*, Wiley, New York.
- Baron, D. and Besanko, D. (1984), Regulation and Information in a Continuing Relationship, *Information Economics and Policy* 1, 447–470.
- Baron, D. and Myerson, R. (1982), Regulating a Monopolist with Unknown Costs, *Econometrica* 50, 911–930.
- Baumol, W., Panzar, J., and Willig, R. (1982), *Contestable Markets and the Theory of Industrial Structure*, Harcourt Brace Jovanovich, New York.
- Bergin, J. and MacLeod, W. B. (1993), Efficiency and Renegotiation in Repeated Games, *Journal of Economic Theory* 61, 42–73.
- Boiteux, M. (1956), Sur la gestion des monopoles publics astreints à l'équilibre budgétaire, *Econometrica* 24, 22–40. (Published in English: (1971), On the Management of Public Monopolies Subject to Budgetary Constraints, *Journal of Economic Theory* 3, 219–240.)
- Boutin, A. (2014), Screening for Good Pools: Consequences of Independent Licensing, Manuscript, Université Libre de Bruxelles.
- Budd, C., Harris, C., and Vickers, J. (1993), A Model of the Evolution of Duopoly: Does the Asymmetry between Firms Tend to Increase or Decrease?, *Review of Economic Studies* 60, 543–573.
- Bulow, J., Geanakoplos, J., and Klemperer, P. (1985), Multimarket Oligopoly: Strategic Substitutes and Complements, *Journal of Political Economy* 93, 488–511.
- Caillaud, B. and Jullien, B. (2003), Chicken and Egg: Competition among Intermediation Service Providers, *RAND Journal of Economics* 24, 309–328.
- Caves, R. E. and Porter, M. E. (1977), From Entry Barriers to Mobility Barriers: Conjectural Decisions and Contrived Deterrence to New Competition, *Quarterly Journal of Economics* 91, 241–261.
- Coase, R. H. (1972), Durability and Monopoly, *Journal of Law and Economics* 15, 143–149.

agent expects to receive at least his utility from the default option uniformly over the agents' beliefs about the principal's type.

- Cyert, R. M. and DeGroot, M. H. (1970), Multiperiod Decision Models with Alternating Choice as a Solution to the Duopoly Problem, *Quarterly Journal of Economics* 84, 410–429. Bayesian Analysis and Duopoly Theory, *Journal of Political Economy* 78, 1168–1184.
- Dal Bó, P. and Frechette, G. (2014), On the Determinants of Cooperation in Infinitely Repeated Games: A Survey, Mimeo (available at <http://ssrn.com/abstract=2535963>).
- Dasgupta, P. and Stiglitz, J. (1980), Uncertainty, Industrial Structure, and the Speed of R&D, *Bell Journal of Economics* 11, 1–28.
- Dewatripont, M. (1989), Renegotiation and Information Revelation over Time: The Case of Optimal Labor Contracts, *Quarterly Journal of Economics* 104, 589–619.
- Dewatripont, M. and Tirole, J. (1994), A Theory of Debt and Equity: Diversity of Securities and Manager-Shareholder Congruence, *Quarterly Journal of Economics* 109, 1027–1054.
- Dixit, A. (1980), The Role of Investment in Entry Deterrence, *Economic Journal* 90, 95–106.
- Eckert, A. (2002), Retail Price Cycles and Response Asymmetry, *Canadian Journal of Economics/Revue canadienne d'économique* 35, 52–77.
- Edgeworth, F. Y. (1925), *Papers Relating to Political Economy*, Volumes I, II, and III, Royal Economic Society.
- Ellison, G. and Fudenberg, D. (2003), Knife-Edge or Plateau: When Do Market Models Tip?, *Quarterly Journal of Economics* 118, 1249–1278.
- Ericson, R. and Pakes, A. (1995), Markov Perfect Industry Dynamics: A Framework for Empirical Work, *Review of Economic Studies* 62, 53–82.
- Farhi, E. and Tirole, J. (2012), Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts, *American Economic Review* 102, 60–93.
- Farrell, J. and Saloner, G. (1985), Standardization, Compatibility, and Innovation, *RAND Journal of Economics* 16, 70–83.
- Farrell, J. and Saloner, G. (1986), Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation, *American Economic Review* 76(5), 940–955.
- Freixas, X., Guesnerie, R., and Tirole, J. (1985), Planning under Incomplete Information and the Ratchet Effect, *Review of Economic Studies* 52, 173–191.
- Fudenberg, D., Gilbert, R., Stiglitz, J. E., and Tirole, J. (1983), Preemption, Leapfrogging and Competition in Patent Races, *European Economic Review* 22, 3–31.
- Fudenberg, D. and Tirole, J. (1983a), Capital as a Commitment: Strategic Investment to Deter Mobility, *Journal of Economic Theory* 31, 227–250.
- Fudenberg, D. and Tirole, J. (1983b), Learning by Doing and Market Performance, *Bell Journal of Economics* 14, 522–530.
- Fudenberg, D. and Tirole, J. (1983c), Sequential Bargaining with Incomplete Information, *Review of Economic Studies* 50, 221–247.
- Fudenberg, D. and Tirole, J. (1984), The Fat-Cat Effect, the Puppy-Dog Ploy, and the Lean and Hungry Look, *American Economic Review* 74, 361–366.
- Fudenberg, D. and Tirole, J. (1985), Preemption and Rent Equalization in the Adoption of New Technology, *Review of Economic Studies* 52, 383–401.
- Fudenberg, D. and Tirole, J. (1987), Understanding Rent Dissipation: On the Uses of Game Theory in Industrial Organization, *American Economic Review* 77, 176–183.
- Fudenberg, D. and Tirole, J. (1990), Moral Hazard and Renegotiation in Agency Contracts, *Econometrica* 58, 1279–1319.
- Fudenberg, D. and Tirole, J. (1991a), Perfect Bayesian Equilibrium and Sequential Equilibrium, *Journal of Economic Theory* 53, 236–260.
- Fudenberg, D. and Tirole, J. (1991b), *Game Theory*, MIT Press, Cambridge, MA.
- Fudenberg, D., Levine, D. K., and Tirole, J. (1985), Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information, in A. Roth (ed.), *Game Theoretic Models of Bargaining*, Cambridge University Press, Cambridge.

- Gul, F., Sonnenschein, H., and Wilson, R. (1986), Foundations of Dynamic Monopoly and the Coase Conjecture, *Journal of Economic Theory* 39, 155–190.
- Grossman, S. J. (1981), The Informational Role of Warranties and Private Disclosure about Product Quality, *Journal of Law and Economics* 24, 461–483.
- Hall, R. and Hitch, C. (1939), Price Theory and Business Behaviour, *Oxford Economic Papers* 2, 12–45.
- Hann, I. H., Roberts, J. A., and Slaughter, S. A. (2013), All Are Not Equal: An Examination of the Economic Returns to Different Forms of Participation in Open-Source Software Communities, *Information Systems Research* 24, 520–538.
- Harris, C. J. and Vickers, J. S. (1985), Perfect Equilibrium in a Model of a Race, *Review of Economic Studies* 52, 193–209.
- Hart, O. and Tirole, J. (1988), Contract Renegotiation and Coasian Dynamics, *Review of Economic Studies* 55, 509–540.
- Hart, O. and Tirole, J. (1990), Vertical Integration and Market Foreclosure, *Brookings Papers on Economic Activity: Microeconomics 1990*, 205–286.
- Hermalin, B. E. and Katz, M. L. (2007), The Economics of Product-Line Restrictions with an Application to the Network Neutrality Debate, *Information Economics and Policy* 19, 215–248.
- Holmström, B. (1999), Managerial Incentive Problems: A Dynamic Perspective, *Review of Economic Studies* 66, 169–182 (originally published in 1982).
- Holmström, B. and Myerson, R. B. (1983), Efficient and Durable Decision Rules with Incomplete Information, *Econometrica* 51, 1799–1819.
- Holmström, B. and Tirole, J. (1998), Private and Public Supply of Liquidity, *Journal of Political Economy* 106, 1–40.
- Horner, J. (2004), A Perpetual Race to Stay Ahead, *Review of Economic Studies* 71, 1065–1088.
- Iijima, R. and Kasaharay, A. (2015), Gradualism and Equilibrium Uniqueness in Stochastic Differential Games, Mimeo.
- Katz, M. L. and Shapiro, C. (1985), Network Externalities, Competition, and Compatibility, *American Economic Review* 75(3), 424–440.
- Katz, M. L. and Shapiro, C. (1986), Technology Adoption in the Presence of Network Externalities, *Journal of Political Economy* 94, 822–841.
- Kreps, D. and Wilson, R. (1982), Sequential Equilibria, *Econometrica* 50, 863–894.
- Laffont, J.-J. and Tirole, J. (1986), Using Cost Observation to Regulate Firms, *Journal of Political Economy* 94, 614–641.
- Laffont, J.-J. and Tirole, J. (1988), The Dynamics of Incentive Contracts, *Econometrica* 56, 1153–1175.
- Laffont, J.-J. and Tirole, J. (1990a), The Regulation of Multiproduct Firms, *Journal of Public Economics* 43, 1–66.
- Laffont, J.-J. and Tirole, J. (1990b), Adverse Selection and Renegotiation in Procurement, *Review of Economic Studies* 75, 597–626.
- Laffont, J.-J. and Tirole, J. (1990c), Optimal Bypass and Cream Skimming, *American Economic Review* 80, 1042–1061.
- Laffont, J.-J. and Tirole, J. (1991), The Politics of Government Decision Making: A Theory of Regulatory Capture, *Quarterly Journal of Economics* 106, 1089–1127.
- Laffont, J.-J. and Tirole, J. (1993), *A Theory of Incentives in Procurement and Regulation*, MIT Press, Cambridge, MA.
- Laffont, J.-J. and Tirole, J. (1994), Access Pricing and Competition, *European Economic Review* 38, 1673–1710.
- Laffont, J.-J. and Tirole, J. (2000), *Competition in Telecommunications*, MIT Press, Cambridge, MA.

- Laffont, J.-J., Rey, P., and Tirole, J. (1998a), Network Competition: I. Overview and Nondiscriminatory Pricing, *RAND Journal of Economics* 29, 1–37.
- Laffont, J.-J., Rey, P., and Tirole, J. (1998b), Network Competition: II. Price Discrimination, *RAND Journal of Economics* 29, 38–56.
- Laffont, J.-J., Marcus, S., Rey, P., and Tirole, J. (2001), Internet Peering, *American Economic Review* 91(2), 287–291.
- Lerner, J. and Tirole, J. (2002), Some Simple Economics of Open Source, *Journal of Industrial Economics* 50, 197–234.
- Lerner, J. and Tirole, J. (2004), Efficient Patent Pools, *American Economic Review* 94, 691–711.
- Lerner, J. and Tirole, J. (2015), Standard-Essential Patents, *Journal of Political Economy*, forthcoming.
- Loeb M. and Magat, W. (1979), A Decentralized Method for Utility Regulation, *Journal of Law and Economics* 22, 399–404.
- Maskin, E. and Tirole, J. (1984), On the Efficiency of Fixed Price Equilibrium, *Journal of Economic Theory* 32, 317–327.
- Maskin, E. and Tirole, J. (1987), A Theory of Dynamic Oligopoly, III: Cournot Competition, *European Economic Review* 31, 947–968.
- Maskin, E. and Tirole, J. (1988a), A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs, *Econometrica* 56, 549–569.
- Maskin, E. and Tirole, J. (1988b), A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles, *Econometrica* 56, 571–599.
- Maskin, E. and Tirole, J. (1990), The Principal–Agent Relationship with an Informed Principal: The Case of Private Values, *Econometrica* 58, 379–409.
- Maskin, E. and Tirole, J. (1992), The Principal–Agent Relationship with an Informed Principal, II: Common Values, *Econometrica* 60, 1–42.
- Maskin, E., and Tirole, J. (2001). Markov Perfect Equilibrium: I. Observable Actions, *Journal of Economic Theory* 100, 191–219.
- Milgrom, P. R. (1981), Good News and Bad News: Representation Theorems and Applications, *Bell Journal of Economics* 12, 380–391.
- Milgrom, P. and Roberts, J. (1982), Limit Pricing and Entry under Incomplete Information, *Econometrica* 50, 443–460.
- Morgan, J., Orzen, H., and Sefton, M. (2006), An Experimental Study of Price Dispersion, *Games and Economic Behavior* 54, 134–158.
- Myerson, R. (1982), Optimal Coordination Mechanisms in Generalized Principal–Agent Problems, *Journal of Mathematical Economics* 10, 67–81.
- Myerson, R. B. (1983), Mechanism Design by an Informed Principal, *Econometrica* 51, 1767–1797.
- Noel, M. D. (2007), Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market, *Journal of Industrial Economics* 55, 69–92.
- Pakes, A. and McGuire, P. (1994), Computing Markov Perfect Nash Equilibrium: Numerical Implications of a Dynamic Differentiated Product Model, *Rand Journal of Economics* 25, 555–589.
- Pelzman, S. (1976), Toward a More General Theory of Regulation, *Journal of Law and Economics* 19, 211–248.
- Posner, R. A. (1975), The Social Cost of Monopoly and Regulation, *Journal of Political Economy* 83, 807–828.
- Ramsey, F. P., (1927), A Contribution to the Theory of Taxation, *Economic Journal* 37, 47–61.

- Reinganum, J. (1981a), On the Diffusion of New Technology: A Game-Theoretic Approach, *Review of Economic Studies* 153, 395–406.
- Reinganum, J. (1981b), Market Structure and the Diffusion of New Technology, *Bell Journal of Economics* 12, 618–624.
- Reinganum, J. F. (1983), Uncertain Innovation and the Persistence of Monopoly, *American Economic Review* 73(4), 741–748.
- Reinganum, J. (1984), Practical Implications of Game Theoretic Models of R&D, *American Economic Review* 74, 61–66.
- Rey, P. and Tirole, J. (1986), The Logic of Vertical Restraints, *American Economic Review* 76, 921–939.
- Rey, P. and Tirole, J. (2013), Cooperation vs. Collusion: How Essentiality Shapes Co-opetition, IDEI Working Paper 801.
- Rochet, J.-C. and Tirole, J. (1996), Interbank Lending and Systemic Risk, *Journal of Money, Credit and Banking* 28, 733–762.
- Rochet, J.-C. and Tirole, J. (2003), Platform Competition in Two-Sided Markets, *Journal of the European Economic Association* 1, 990–1029.
- Rochet, J.-C. and Tirole, J. (2006), Two-Sided Markets: A Progress Report, *RAND Journal of Economics* 35, 645–667.
- Rothschild, M. and Stiglitz, J. (1976), Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Incomplete Information, *Quarterly Journal of Economics* 90, 629–649.
- Royal Swedish Academy of Sciences (2014), *Jean Tirole: Market Power and Regulation, Scientific background on the Sveriges Riksbank Prize in Economic Sciences in memory of Alfred Nobel 2014*, compiled by the Economic Sciences Prize Committee (available at http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2014/advanced.html).
- Rubinstein, A. (1982), Perfect Equilibrium in a Bargaining Model, *Econometrica* 50, 97–109.
- Sappington, D. (1982), Optimal Regulation of Research and Development under Imperfect Information, *Bell Journal of Economics* 13, 354–368.
- Schmalensee, R. (1982), Product Differentiation Advantages of Pioneering Brands, *American Economic Review* 72(3), 349–365.
- Scherer, F. M. (1967), Research and Development Resource Allocation under Rivalry, *Quarterly Journal of Economics* 81, 359–394.
- Shapiro, C. (1989), Theories of Oligopoly Behavior, in R. Schmalensee and R. Willig (eds.), *Handbook of Industrial Organization*, Vol. 1, North-Holland, Amsterdam, 329–414.
- Shapiro, C. (2000), Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting, *Innovation Policy and the Economy* 1, 119–150.
- Shapley, L. (1953), Stochastic Games, *Proceedings of the National Academy of Sciences USA* 39, 1095–1100.
- Simon, L. K. and Stinchcombe, M. B. (1989), Extensive Form Games in Continuous Time: Pure Strategies, *Econometrica* 57, 1171–1214.
- Sobel, M. J. (1971), Noncooperative Stochastic Games, *Annals of Mathematical Statistics* 42, 1930–1935.
- Sobel, J. and Takahashi, I. (1983), A Multistage Model of Bargaining, *Review of Economic Studies* 50, 411–426.
- Spence, A. M. (1977), Entry, Capacity, Investment and Oligopolistic Pricing, *Bell Journal of Economics* 8, 534–544.
- Spence, A. M. (1979), Investment Strategy and Growth in a New Market, *Bell Journal of Economics* 10, 1–19.
- Spiegler, R. (2000), Extracting Interaction-Created Surplus, *Games and Economic Behavior* 30, 142–162.
- Stahl, I. (1972), *Bargaining Theory*, Economic Research Institute, Stockholm.

- Stigler, G. (1971), The Economic Theory of Regulation, *Bell Journal of Economics* 2, 3–21.
- Stokey, N. L. (1981), Rational Expectations and Durable Goods Pricing, *Bell Journal of Economics* 12, 112–128.
- Sweezy, P. (1939), Demand under Conditions of Oligopoly, *Journal of Political Economy* 47, 568–573.
- Tirole, J. (1982), On the Possibility of Speculation under Rational Expectations, *Econometrica* 50, 1163–1182.
- Tirole, J. (1986), Hierarchies and Bureaucracies, *Journal of Law, Economics, and Organization* 2, 181–214.
- Tirole, J. (1988), *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.
- Tirole, J. (2006), *The Theory of Corporate Finance*, Princeton University Press, Princeton, NJ.
- Tirole, J. and Guesnerie, R. (1981), Tax Reform from the Gradient Projection Viewpoint, *Journal of Public Economics* 15, 275–293.
- Vespa, E. I. (2014), Cooperation in Dynamic Games: An Experimental Investigation, Working Paper, University of California, Santa Barbara.
- Vogelsang, I. and Finsinger, J. (1979), A Regulatory Adjustment Process for Optimal Pricing by Multiproduct Monopoly Firms, *Bell Journal of Economics* 10, 157–171.
- Weitzman, M. (1976), The New Soviet Incentive Model, *Bell Journal of Economics* 7, 251–257.
- Wilson, C. (1977), A Model of Insurance Markets with Incomplete Information, *Journal of Economic Theory* 12, 167–207.