Tirole's Industrial Regulation and Organization Legacy in Economics

Citation

Fudenberg, Drew. 2015. Tirole's Industrial Regulation and Organization Legacy in Economics. *The Scandinavian Journal of Economics* 117, no. 3: 771–800. doi:10.1111/sjoe.12117.

Published Version

doi:10.1111/sjoe.12117

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:27303657>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Tirole's Industrial Regulation and Organization Legacy in Economics

Drew Fudenberg

March 21, 2015

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.

1. 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. This essay will try to convey the main ideas of that work, along with some of Tirole’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 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 essay, 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 MIT. 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

¹ Tirole explaining why he won’t offer opinions on arbitrary topics: “I am a researcher, and not able to talk about everything and anything”: <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.

² For a broader and more detailed overview of Jean’s work, see the Economic Sciences Prize Committee’s excellent scientific report; I relied on it extensively when writing this essay. Another sign of what this essay 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 IO.

racers 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 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), (1983b), (1983c)), 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 IO 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 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

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 fixed-price equilibria (Maskin and Tirole (1984)).

between research and its application, but 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 IDEI (Institut d'Economie Industrielle), 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 groundbreaking 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 Tirole's own subsequent work and the work he has inspired. Tirole'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 twelve books; I say more about a few of them below.

2. Regulation as Applied Mechanism Design

2A: 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 asking what sorts of regulations should be used. Ramsey (1927), Boiteux (1956) and others studied how to set prices to maximize total social surplus (that is, 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 markup $(p_i - c_i) / p_i$ 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 markups cause more deadweight loss, and the ratio between the regulated markup and the monopoly markup is the same for every good. In practice, though, regulators may 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 may 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-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 since 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 can't be altered by the firm's actions. ⁶ Sappington (1982) allowed for moral hazard in

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.

addition to adverse selection, as well, supposing that the firm can do unobserved research to reduce its production cost, but restricted attention to linear contracts.

Using Cost Information to Regulate Firms

Like Sappington, 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) appears 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 (The paper allows cost to also depend on a mean-zero cost shock that is independent of 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.

⁷ The identification of the firm with a single agent is a standard assumption in most of the regulation and industrial organization literatures, which gloss over the internal structure of the firm, but see Chapter 7 of Tirole (2006) for a discussion of how that structure can interact with behavior in the product market. One way to interpret the firm's "effort" here is as a (restraint on) wasteful perks for management.

⁸ In a footnote, the paper explains that modelling 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, and is explored at length in Laffont and Tirole (1993).

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 prefer to announce truthfully and take the assigned effort level, and also satisfy the interim individual rationality constraints that all types are willing to participate. Under some technical assumptions,⁹ Laffont and Tirole (1986) showed that it is sufficient to consider only the local incentive constraint, and 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 to do so if costs were not observed, as in Baron and Myerson, but this would increase the

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 monotone hazard rate property. Additional assumptions are used to ensure that it is neither optimal to shut down the firm nor send the marginal cost of production to 0, that the full information problem is "sufficiently convex," and that the first-order condition for the planner's maximization is sufficient.

size of the payments needed to ensure truthful reporting. Instead, the optimal contract provides cost sharing and lower lump sum payments to higher types, so 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 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 doesn't 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, and Laffont and Tirole (1990c and 1993 Chapter 6) study "cream skimming" by a firm that faces two different markets or types of consumers.

In hindsight, given Jean's later work on platform competition and two-sided markets, the most evocative static extension may be Laffont and Tirole (1993 Chapter 5 and 1994) on access pricing, which involves a regulated multi-product 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, and this would lead the regulator to reduce planned supply, which

reduces the firm's incentive to exaggerate cost. On the other hand 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 Laffont and Tirole (1986) and much of Laffont and Tirole (1993) studied a static problem, Laffont and Tirole (1988, 1990b, 1993 Chapters 9 and 10) 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 may be difficult to achieve, which motivated Laffont and Tirole to consider other scenarios. Laffont and Tirole (1988) supposes that the regulator is only able to offer short-term contracts, so that not only can the regulator not commit, the firm cannot either, as 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.

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

Freixas, Guesnerie, and Tirole (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-1 mechanism can induce full revelation of types- there will always be some degree of pooling, so after period 1 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.

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

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

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. The paper establishes this result when the prior is a smooth density over a continuum of types, but the detailed analysis restricts 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- though 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-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”: accepting the fixed-price contract reveals the firm’s type, but accepting the contract with cost sharing does not. Thus, as in the case of only short-term contracts, 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 1930’s 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, and 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, 1993 chapter 11) combined their basic “cost observation” model with the three-tier structure that

¹² For the same reason, the optimal renegotiation-proof contracts are semi-separating in the Hart-Tirole (1988) model of renegotiation-proof contracts for the repeated rental of a durable good, and in the Fudenberg-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.

Tirole (1986a) 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 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 ($\underline{\beta}$). 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, but if the supervisor learns that the firm's cost is low, the firm has an incentive to bribe her to suppress that information. Foreseeing this, Congress knows it must pay the supervisor a bonus when she reports that the type is low, and since 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 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 (that is, more cost sharing, so that the transfer is less

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 unravelling argument shows that failure to report will be interpreted as revealing the worst possible information; that conclusion extends to sufficiently small probabilities that the sender is uninformed.

responsive to the outcome) as in Tirole (1986a), which can be seen as a defense or explanation for the perhaps surprising prevalence of such contracts.¹⁴

2B 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 economic literature.

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

¹⁴ The extent of this distortion depends on the probability that the supervisor is uninformed, and vanishes as this probability goes to 0. 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.

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

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, and the same is true with linear demand when retailers aren't 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 that 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, as 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), 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.

Telecommunications

In the 1960's and 70's, telecommunications in the US and many other countries was run by either a private but regulated monopoly or a state-run firm, but since the early 80's the trend has been towards privatization and oligopoly. Because telecoms need to cooperate to send

¹⁶ The easiest way to see this is may be to suppose that the monopolist doesn't 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 absent commitment power the collusive agreement would unravel to Cournot competition.

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 interconnection or termination fees (when a consumer calls a customer of another network). Laffont, Rey and Tirole (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); 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 1 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 (2001) pointed out that the static Cournot model implies that pools are anti-competitive (lead to higher prices and lower sales) when the patents are perfect substitutes, and pro-competitive (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.¹⁷ Here perfect complements corresponds to the case where there is no benefit to owning any fewer than all of the licenses, while perfect substitutes can be defined as $V(n) = V(n-1)$ since this is enough

¹⁷ Perhaps surprisingly, whether V is concave or convex plays a very small role in the analysis, though Proposition 4 shows that it does matter for the existence of asymmetric equilibria. The assumed additive separability between consumer type and 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 the paper 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

for Bertrand competition to lead to 0 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 substitutes when prices are high, because users will only buy some of the patents so that the patents are in competition.

To analyze the effect of a patent pool, Lerner and Tirole compare it to the symmetric pure-strategy equilibrium of the Bertrand pricing game, which they show is unique. Generically, a firm 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 else the “demand margin binds,” meaning that the consumers who purchase any patent buy all of them, and increasing price lowers total sales without changing 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). Since regulators may 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 stand-alone licenses. That way, firms will only be tempted to offer a stand-alone 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 stand-alone licenses; however, the agreements are socially beneficial when unbundling is imposed, that is, 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 will only summarize that work 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 facts that the interconnections of financial institutions means that a bankruptcy of one firm can have negative impacts on others. 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 studies what sorts of regulations might help restore them; it anticipated many of the problems that would hit the financial system in the 2008 crisis. Farhi and Tirole (2012) showed that financial institutions may 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.

3. Industrial Organization

3A Strategic Investment

The study of entry and deterrence has a long history in the study of oligopolistic markets. 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 absence of substantial depreciation) having a substantial pre-entry capital stock is a credible commitment to have a large capital stock after entry, and so if the incumbent firm's capital level lowers the profitability of entry, the former may 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) modelled entry as a discrete event, but Caves and Porter (1977) argued, entry is often accomplished in stages, so strategic investment may be used not only to deter entry but more generally to deter "mobility." Because mobility deterrence is an ongoing process, it is best modelled 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 stop. As 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 (1980) model more formally.¹⁸ 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.”¹⁹ 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 it is optimal for each firm to do when the other one does, so that the outcomes we derived were also Markov

¹⁸ This was Jean’s first paper to use game theory. His first publication, Guesnerie and Tirole (1981), came from the work he did in France for his master’s degree.

¹⁹ In the early-stopping equilibria the value functions are discontinuous in the state. Iijima and Kasaharaya (2015) show that value functions are continuous when the state is subject to Brownian motion, and use this to show that there is a unique equilibrium in the finite-horizon stochastic differential version of the game.

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 these may be Pareto-ranked.²⁰

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 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, Geanakoplos, and Klemperer (1985), the incumbent (firm 1) moves first and chooses a capital stock K_1 , which could for example be productive machinery, advertising, 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.²¹ 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 others, for example when K_1 represents pre-entry advertising. In this setting, a subgame-perfect

²⁰ 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.

²¹ 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 not be perfect substitutes.

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, 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 that the capital stock K_1 that firm 1 would choose in the absence of 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, $\max_{K_1} (\pi^1(K_1, x_1^*(K_1), x_2^*(K_1)) | \pi^2(K_1, x_1^*(K_1), x_2^*(K_1)) \geq 0)$.

To determine the best way for firm 1 to deter entry, consider the total derivative of π^2 with respect to K_1 , namely (using subscripts K_1 , 1, and 2 to denote partial derivatives with respect to K_1, x_1 , and x_2 , and noting that $\pi_2^2 = 0$ from the envelope theorem)

$$\frac{d\pi^2}{dK_1} = \pi_{K_1}^2 + \pi_1^2 \frac{dx_1^*}{dK_1}. \text{ Here the term } \pi_{K_1}^2 \text{ is the direct effect of } K_1 \text{ on firm 2's profit, and}$$

$$\pi_1^2 \frac{dx_1^*}{dK_1} \text{ is the strategic effect- the effect that comes from the fact that pre-entry changes in } K_1$$

will change firm 1's post-entry behavior. Fudenberg and Tirole say that investment makes a

firm "tough" if $\frac{d\pi^2}{dK_1}$ 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 as 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 may 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 $\pi_2^1 \frac{dx_2^*}{dK_1}$. 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, Geanakoplos, and Klemperer (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- the “top dog” strategy. With strategic complements (for example 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.

3B 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 may be small, but they can nonetheless have a large impact on the time paths of prices and output, as shown in 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),²² but 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 Markov Perfect Equilibria (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.

²² They later 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.

In the (1988a) paper, firms choose quantities 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.²³ Maskin and Tirole show that there is a unique symmetric MPE. When the discount factor δ 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 e.g. Bain (1956), but here as the period length grows shorter so that $\delta \rightarrow 1$, $\pi(q^*, 0) \rightarrow f$, and so the incumbent gains very little advantage from very short run commitments, as in the “contestability” argument of Baumol, Panzar, and Willig (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 commitments have little value relies on the symmetry assumption.

Maskin and Tirole (1988b) consider alternating-move duopoly when firms set prices each period as opposed to quantities, and prices are constrained to lie on a discrete grid. Here there are

²³ $8f > (1 - c)^2 > 4f$.

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, and although price eventually does 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 this was the first paper to derive them formally. Since then, Edgeworth cycles have been studied in both field data (starting with Eckert (2002) and Noel (2007)) and laboratory experiments (Morgan, Orzen, and Sefton (2006)).

3C Innovation and Technology Adoption

Understanding technology development, adoption, and diffusion is very important for studying industry dynamics and for designing patent, copyright, and related policy instruments, and Tirole has contributed to this understanding in a variety of different settings.

Patent Races

A long literature has studied the effectiveness of patents and their impact on market structure and competition, and at the end of the 1970’s 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.²⁴ Fudenberg, Gilbert, Stiglitz, and Tirole (1983) modelled 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

²⁴ See Reinganum (1984) for a survey of those models.

show that if progress towards the finish line is a deterministic function of expenditure, and 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 may 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.²⁵

Open-source software

A substantial fraction 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 Holmstrom (1982): Contributing to open source software may serve to signal one’s programming ability, and thus lead to job offers, and may also generate psychic rewards from being acknowledged by one’s peers. Drawing on Holmstrom’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

²⁵ Harris and Vickers (1985) obtained ε -preemption in a perfect information model with a much more general R&D technology. Budd, Harris, and Vickers (1993) and Horner (2004) analyze the more complex behavior that arise when the state of competition is measured by the difference between the number of patents obtained by each firms.

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 may 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) The open source programmer is her own boss and takes full responsibility for the success of a subproject, so her ability is more tightly linked to the outcome; and (iii) Open-source programmers are more mobile as they have less firm-specific human capital. They then compare their theoretical predictions with case studies of Apache, Linux, Perl and Sendmail, and find support for the signaling hypothesis, as did Hann, Roberts and Slaughter (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 may 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) had 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 the paper supposes there are only two firms and I will 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 one 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 quasiconcave 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, one can then compute the continuation payoffs $L(t), F(t)$ and $J(t)$ of the firms as a function of the time of first adoption and whether firm i was the leader (first to adopt), follower, or if both firms jointly adopted; the assumptions on the underlying payoff functions imply 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 nonempty interval $(\bar{T}_1, \bar{T}_2)$, and $L(t) = F(t) = J(t)$ for a possibly empty interval $[\bar{T}_2, \infty)$.²⁶ 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” (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 to be the first, but this is not consistent with equilibrium, because any delay would allow the other firm to adopt first.²⁷ This rent-dissipation outcome is also an equilibrium in the second case, but there are also

²⁶ Fudenberg and Tirole (1991, Chapter 4) give a more complete taxonomy of timing games.

²⁷ As an example, suppose the two firms are engaged in Bertrand competition with a homogenous good, so that each makes 0 profit initially, and adopting the innovation lowers the constant average cost of production from $\bar{c}$ to $\underline{c}$.

“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 glosses over 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 spaces” that had been previously used to model mixed strategies in timing games, 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.

3D 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 like Ebay or Uber, but also employment agencies, dating services, credit cards, newspapers, and computer operating systems. This leads to a form of “network effect,” but unlike in the work of Katz and Shapiro (1985, 1986) and Farrell and Saloner (1985, 1986), which models the externalities between users on the same side of a network (as with say instant messaging programs) the externality is across sides.²⁸

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).

²⁸ 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 or auctions, there can be negative externalities between users on the same side, so when the populations are finite a

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.²⁹ They pointed out that practices that 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 extract a higher price from the other side.³⁰

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.³¹

Rochet and Tirole (2006) show that the monopoly solution to this problem has the following simple form: Let $p^i = a^i + \frac{A^i - C^i}{N^j}$ 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

user who moves onto a particular platform makes that platform worse for all users on that 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.”

²⁹ Spiegler (2000) considers surplus extraction when agents in each side of the market are identical and the platforms can use general contracts instead of two-part prices. Here, instead of the tradeoff 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.

³⁰ 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.

³¹ Rochet-Tirole (2003) assumed that the membership benefit and the fixed cost per user are both identically equal to 0; the 2006 paper generalized this to cover the model of Armstrong (2006), but only covers the monopoly case. Section 5 of Rochet-Tirole (2006) discusses extensions including allowing agents on one side to care about the quality or type of agents on the other or to care about the number of agents on their own side.

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 0 and the firm uses only per-use pricing, the markup 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.

Rochet and Tirole (2003) go on to consider the case of competing platforms, where agents have a choice between “singlehoming” (using only one platform, e.g. one credit card or dating site) or “multihoming” and using both of them. This issue has important implications for the design of regulations and anti-trust policy, and has sparked a sizable literature.³²

4. 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 Markov perfect equilibrium (MPE) that shows how to identify the relevant notion of a “state” from the extensive form of the game. MPE has 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).³³

³² In addition to Caillaud and Julien (2003) and Rochet and Tirole (2003), see e.g. Armstrong (2006), Julien (2006), Armstrong and Wright (2007), and Hermalin and Katz (2007).

³³ Testing the MPE restriction on field data seems very difficult. In the lab it is well established that behavior in repeated games 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 Bo and Frechette (2015)). In other, more complicated, experiments, MPE does 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.

Kreps and Wilson's (1982) sequential equilibrium 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 complicated to apply to games of any length. Perhaps for this reason, various authors throughout the 80's 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 multi-stage games with observed actions using four quantifier-free conditions such as "no signaling what you don't know," 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 show 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 0, which requires additional effort on the part of both the players and the analyst and may not always be plausible.

Inspired by the work of Rubinstein (1982) and Stahl (1972) on non-cooperative models of bargaining with complete information, Fudenberg and Tirole (1983) and Fudenberg, Levine and Tirole (1985) studied two-person extensive-form bargaining with incomplete information. The first paper showed how incomplete information could lead to bargaining breakdowns and delays in a two-period model. The second paper 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; it 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 0, a result that was extended to the stationary

equilibria of the no-gap case (and a more general set of type distributions) by Gul, Sonnenschein and Wilson (1986).³⁴

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 his willingness to accept it and his behavior if he does.³⁵

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,"³⁶ and use this to help explain the difference between the Spence signaling model (where the worker chooses effort before the contract is signed) and models of insurance such as Rothschild-Stiglitz (1976).

³⁴ 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.

³⁵ They build on the work of Myerson (1983), which had posed the general problem of mechanism design by an informed principal when there are multiple agents, each of which may have some private information, and Holmstrom and Myerson (1983), which developed the notion of interim efficiency and of "durable" agreements.

³⁶ So-called because in insurance markets it corresponds to the zero-profit separating allocation that figures prominently in Rothschild-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 that ensure the agent at least the utility he gets from the default option no matter what his beliefs are.

Finally, Jean deserves much of the credit for our having written our (1991b) game theory textbook, which I'd like to think helped 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 J. Wright (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, 32, 353-38.
- Averch, H. and L.L. Johnson (1962). Behavior of the firm under regulatory constraint. *American Economic Review* 52, 1052-1069.
- Bain, J. (1956). *Industrial Organization*. New York: Wiley.
- Baron, D. and D. Besanko (1984). Regulation and information in a continuing relationship. *Information Economics and Policy* 1, 447-470.
- Baron, D. and R. Myerson (1982). Regulating a monopolist with unknown costs. *Econometrica* 50, 911-930.
- Baumol, W., J. Panzar, and R. Willig (1982). *Contestable Markets and the Theory of Industrial Structure*. New York: Harcourt, Brace, Jovanovich.
- Berry, S., and A. Pakes (1993). Some applications and limitations of recent advances in empirical industrial organization: Merger analysis. *American Economic Review* 83, 247-252.
- Boutin, A. 2014 "Screening for Good Pools: Consequences of Independent Licensing." Manuscript, Université Libre de Bruxelles.
- Bulow, J., J. Geanakoplos and P. Klemperer (1985). Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political Economy* 93, 488-511.
- Caillaud, B. and B. Jullien (2003) Chicken and egg: Competition among intermediation service providers. *RAND Journal of Economics* 24, 309-328.
- Coase, R. H. (1972). Durability and monopoly. *Journal of Law and Economics* 15, 143-149.

- Dal Bó, E. (2006). Regulatory capture: A review. *Oxford Review of Economic Policy* 22, 203-225.
- Dal Bo, P. and G. Frechette (2014). On the determinants of cooperation in infinitely repeated games: A survey. Mimeo.
- Dasgupta, P., and J. Stiglitz, (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.
- Diamond, P. and J. Mirrlees (1971). Optimal taxation and public production I: Production efficiency. *American Economic Review* 61, 8-27.
- 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'économie*, 35, 52-77.
- Ellison, G. and D. Fudenberg (2003). Knife-edge or plateau: When do market models tip? *Quarterly Journal of Economics* 118, 1249-1278.
- Ericson, R. and A. Pakes (1995). Markov perfect industry dynamics: A framework for empirical work. *Review of Economic Studies* 62, 53-82.
- Farhi, E. and J. Tirole (2012a). Collective moral hazard, maturity mismatch, and systemic bailouts. *American Economic Review* 102, 60-93.
- Farrell, J., and G. Saloner, G. (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics*, 70-83.
- Farrell, J., and G. Saloner (1986). Installed base and compatibility: Innovation, product preannouncements, and predation. *The American economic review*, 940-955.
- Freixas, X., R. Guesnerie and J. Tirole (1985). Planning under incomplete information and the ratchet effect. *Review of Economic Studies* 52, 173-191.
- Fudenberg, D., R. Gilbert, J.E. Stiglitz and J. Tirole (1983). Preemption, leapfrogging and competition in patent races. *European Economic Review* 22, 3-31.
- Fudenberg, D. and J. Tirole (1983a). Sequential bargaining with incomplete information. *Review of Economic Studies* 50, 221-247.
- Fudenberg, D. and J. Tirole (1983b). Capital as a commitment: Strategic investment to deter mobility. *Journal of Economic Theory* 31, 227-250.

Fudenberg, D. and J. Tirole (1983c). Learning by doing and market performance. *Bell Journal of Economics* 14, 522-530.

Fudenberg, D. and J. Tirole (1984). The Fat-Cat Effect, the Puppy-Dog Ploy, and the Lean and Hungry Look. *American Economic Review* 74, 361-366.

Fudenberg, D. and J. Tirole (1985). Preemption and rent equalization in the adoption of new technology. *Review of Economic Studies* 52, 383-401.

Fudenberg, D. and J. Tirole (1987). Understanding rent dissipation: On the uses of game theory in industrial organization. *American Economic Review* 77, 176-183.

Fudenberg, D. and J. Tirole (1990). Moral hazard and renegotiation in agency contracts. *Econometrica* 58, 1279-1319.

Fudenberg, D. and J. Tirole (1991a). Perfect Bayesian Equilibrium and Sequential Equilibrium. *Journal of Economic Theory* 53, 236-260.

Fudenberg, D. and J. Tirole (1991b). *Game Theory*. Cambridge, MA: MIT Press.

Gul, F., H. Sonnenschein and R. Wilson (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.

Hann, I. H., J.A. Roberts and S.A. Slaughter (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.

Hart, O. and J. Tirole (1988). Contract renegotiation and Coasian dynamics. *Review of Economic Studies* 55, 509-540.

Hart, O. and J. Tirole (1990). Vertical integration and market foreclosure. *Brookings Papers on Economic Activity* vol. 1990: "Microeconomics", Brookings Institution, Washington, pp. 205-286.

Holmström, B. (1982,1999). Managerial incentive problems: A dynamic perspective. *Review of Economic Studies* 66, 169-182.

Horner, J. (2004). A perpetual race to stay ahead. *Review of Economic Studies* 71, 1065-1088.

Iijima, R. and A. Kasaharaya (2015) Gradualism and Equilibrium Uniqueness in Stochastic Differential Games, mimeo.

- Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American economic review*, 424-440.
- Katz, M. L., & Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of political economy*, 822-841.
- Kreps, D. and R. Wilson (1982). Sequential equilibria. *Econometrica* 50, 863-894.
- Laffont, J. J., Marcus, S., Rey, P., & Tirole, J. (2001). Internet peering. *American Economic Review*, 287-291.
- Laffont, J-J. and J. Tirole (1986). Using cost observation to regulate firms. *Journal of Political Economy* 94, 614-641.
- Laffont, J-J. and J. Tirole (1987). Auctioning incentive contracts. *Journal of Political Economy* 95, 921-937.
- Laffont, J-J. and J. Tirole (1988). The dynamics of incentive contracts. *Econometrica* 56, 1153-1175.
- Laffont, J-J. and J. Tirole (1990a): The regulation of multiproduct firms. *Journal of Public Economics* 43, 1-66.
- Laffont, J-J. and J. Tirole (1990b). Adverse selection and renegotiation in procurement. *Review of Economic Studies* 75, 597-626.
- Laffont, J-J. and J. Tirole (1990c). Optimal bypass and cream skimming. *American Economic Review* 80, 1042-1061.
- Laffont, J-J. and J. Tirole (1991). The politics of government decision making: A theory of regulatory capture. *Quarterly Journal of Economics* 106, 1089-1127.
- Laffont, J-J. and J. Tirole (1993). *A Theory of Incentives in Procurement and Regulation*. Cambridge, MA: MIT Press.
- Laffont, J-J. and J. Tirole (1994). Access pricing and competition. *European Economic Review* 38, 1673-1710.
- Laffont, J-J. and J. Tirole (1996). Creating competition through interconnection: Theory and practice. *Journal of Regulatory Economics* 10, 227-256.
- Laffont, J-J. and J. Tirole (2000). *Competition in Telecommunications*. Cambridge, MA: MIT Press.
- Laffont, J-J., P. Rey, and J. Tirole (1998a). Network competition: I. Overview and nondiscriminatory pricing. *RAND Journal of Economics* 29, 1-37.

Laffont, J-J., P. Rey, and J. Tirole (1998b). Network competition: II. Price discrimination. *RAND Journal of Economics* 29, 38-56.

Lerner, J. and J. Tirole (2002). Some simple economics of open source. *Journal of Industrial Economics* 50, 197-234.

Lerner, J. and J. Tirole (2004). Efficient patent pools. *American Economic Review* 94, 691-711.

Lerner, J. and J. Tirole (2015). Standard-essential patents, forthcoming in the *Journal of Political Economy*.

Loeb M. and W. Magat (1979). A decentralized method for utility regulation. *Journal of Law and Economics* 22, 399-404.

Maskin, E. and J. Tirole (1987). A theory of dynamic oligopoly, III: Cournot competition. *European Economic Review* 31, 947-968.

Maskin, E. and J. Tirole (1988a). A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica* 56, 549-569.

Maskin, E. and J. Tirole (1988b). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica* 56, 571-599.

Maskin, E. and J. Tirole (1990). The principal-agent relationship with an informed principal: The case of private values. *Econometrica* 58, 379-409.

Maskin, E. and J. Tirole (1992). The principal-agent relationship with an informed principal, II: Common values. *Econometrica* 60, 1-42.

Maskin, E., and J. Tirole (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. *The Bell Journal of Economics*, 380-391.

Milgrom, P. and J. Roberts (1982). Limit pricing and entry under incomplete information. *Econometrica* 50, 443-460.

Morgan, J., H. Orzen and M. Sefton (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. (1991). *Game Theory*. Cambridge, MA: Harvard University Press.

Noel, M.D. (2007). Edgeworth price cycles: Evidence from the toronto retail gasoline market. *The Journal of Industrial Economics*, 55, 69-92.

Pakes, A. and P. McGuire (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. (1982). A dynamic game of R and D: Patent protection and competitive behavior. *Econometrica* 50, 671-688.

Reinganum, J. (1984). Practical implications of game theoretic models of R&D. *American Economic Review* 74, 61-66.

Rey, P. and J. Tirole (1986). The logic of vertical restraints. *American Economic Review* 76, 921-939.

Rey, P. and J. Tirole (2013) Cooperation vs. Collusion: How Essentiality Shapes Co-opetition,” mimeo.

Rochet, J-C. and J. Tirole (2003). Platform competition in two-sided markets. *Journal of the European Economic Association* 1, 990-1029.

Rochet, J-C. and J. Tirole (2006). Two-sided markets: A progress report. *RAND Journal of Economics* 35, 645-667.

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*.

Sappington, D. (1982). Optimal regulation of research and development under imperfect information. *Bell Journal of Economics* 13, 354-368.

- 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. *Handbook of industrial organization*, 1, 329-414.
- Shapiro, C. (2001). Navigating the patent thicket: Cross licenses, patent pools, and standard setting. In *Innovation Policy and the Economy, Volume 1* (pp. 119-150). MIT press.
- Shapley, L. (1953). Stochastic games. *Proceedings of the National Academy of Sciences USA* 39, 1095-1100.
- 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, Ran. "Extracting interaction-created surplus." *Games and Economic Behavior* 30.1 (2000): 142-162.
- Stigler, G. (1971). The economic theory of regulation. *Bell Journal of Economics* 2, 3-21.
- 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. (1985). Asset bubbles and overlapping generations. *Econometrica* 53, 1499-1528.
- Tirole, J. (1986a). Hierarchies and bureaucracies. *Journal of Law, Economics, and Organization* 2, 181-214.
- Tirole, J. (1986b). Procurement and renegotiation. *Journal of Political Economy* 94, 235-259.
- Tirole, J. (1988). *The Theory of Industrial Organization*. Cambridge: MIT Press.
- Tirole, J. (2006). *The Theory of Corporate Finance*. Princeton, NJ: Princeton University Press.
- Tirole, J. (2012). Overcoming adverse selection: How public intervention can restore market functioning. *American Economic Review* 102, 29-59.
- Volgelsang, I. and J. Finsinger (1979). A regulatory adjustment process for optimal pricing by multiproduct monopoly firms. *Bell Journal of Economics* 10, 157-171