"GENERAL INTRODUCTION" TO *INFORMATION ECONOMICS*

MICHAEL R. BAYE AND DAVID E. M. SAPPINGINGTON

PRELIMINARY DRAFT, SEPTEMBER 2013

Information economics can be viewed as the study of economic issues in settings where uncertainty or asymmetric knowledge of prevailing conditions plays a central role. Such study underlies many inquiries in a vast array of fields. It is perhaps not surprising, then, that a search on *Google Scholar* for "information economics" returns more than 2.5 million results and over 28,000 exact matches.

Space constraints limit us to including in these four volumes only a miniscule fraction of the already expansive and growing literature on information economics. To help refine our choice among the many important and influential articles in this literature, we have focused on contributions in microeconomic analysis. Thus, seminal contributions to macroeconomics, including the pioneering work of Muth (1961) and Sargent and Wallace (1975) on rational expectations, do not appear in these volumes.

Volume 1 begins with some technical, foundational articles that address such issues as how to represent utility in the presence of uncertainty, how to measure riskiness and risk aversion, and how to define relevant equilibrium concepts in settings with incomplete or imperfect information. Volume 1 concludes with a collection of influential articles that illustrate, for example, how informational asymmetries complicate market interactions and impede market performance. Volume 2 focuses on articles that provide more detailed examinations of how information (or the lack thereof) affects resource allocation within firms and other organizations and that analyze potential solutions to moral hazard and adverse selection problems. Volume 3 presents articles that consider how common mechanisms such as auctions and bargaining are employed to determine the allocation of scarce resources when buyers and sellers have limited information. The articles in Volume 4 examine the impact of incomplete and imperfect information on a range of issues in industrial organization, including industry pricing, innovation, collusion, and entry deterrence.

The remainder of this introduction provides a brief overview of the articles that appear in the four volumes. We attempt to trace the evolution of the relevant literatures and to identify some of the many important related articles that, unfortunately, we could not include in these volumes. We note at the outset that some of the papers in these volumes (e.g., Weitzman [7], Vickrey [11], Stigler [12], Spence [15], and Myerson [41]) have influenced multiple areas of information economics. Consequently, the ensuing discussion does not mirror exactly the order in which the papers appear in the four volumes.

MODELING RISK, UNCERTAINTY, AND BELIEFS

In order to characterize an individual's choice among commodities that entail uncertain payoffs (e.g., financial stocks), the individual's preferences must be specified. Von Neumann and Morgenstern (1947)'s expected utility hypothesis suggests how to model these preferences. Friedman and Savage [2] review the formulation, foundations, and interpretation of this hypothesis. The authors also identify the postulates that underlie the expected utility hypothesis. In essence, the postulates are that: (i) individuals have complete and consistent (transitive) orderings of all outcomes that might arise; (ii) preferences are continuous; and (iii) rankings of outcomes are not affected by the presence of irrelevant alternatives.

Hadar and Russell [5] show that all individuals may prefer some lotteries to others. Formally, a lottery is a set of possible payoffs, $x \in [\underline{x}, \overline{x}] \equiv X$, and associated density and distribution functions, $f(x)$ and $F(x)$. Hadar and Russell show, for example, that all individuals whose utilities are increasing in x prefer lottery $f(x)$ to lottery $g(x)$ if $f(x)$ dominates $g(x)$ in the sense of first-order stochastic dominance, i.e., if $F(x) \leq G(x)$ for all $x \in X$. Intuitively, if the higher payoffs are systematically more likely in the sense that the probability of receiving at least a given payoff is systematically higher under $f(x)$ than under $g(x)$, then all individuals will prefer to face the $f(x)$ lottery.[1]

The systematic preference identified by Hadar and Russell [5] holds for all individuals, regardless of how averse they might be to the risk associated with the lotteries. Pratt [3,4] provides a formal means to characterize the extent to which individuals are averse to risk. Pratt demonstrates, for example, that the amount of wealth (w) a risk-averse individual with utility function $u(w)$ is willing to forego in order to avoid a small, actuarially fair gamble is proportional to $r(w) = -u''(w)/u'(w)$.

Rothschild and Stiglitz [6] consider how to compare the variation in the outcomes that can arise in different settings.[2] The authors show that the following three conditions for random variable Y to be "more variable" than random variable X are equivalent: (i) $Y = X + Z$, where Z is a random variable whose expected value, given the realization of X is zero (for all X); (ii) all risk averse individuals prefer to face payoffs that reflect the realizations of variable Y than variable X when the two variables have the same expected value; and (3) the density function for Y is derived from the density function for X by moving some probability weight from its center to the tails.

In some decision environments, an observable signal x about an unobserved variable θ is available. For example, x might represent an individual's observed performance and θ might represent his unobserved effort. Milgrom [8] determines when a higher realization of x constitutes systematically stronger evidence of a higher underlying realization of θ. This will be the case when $f(x|\theta)$, the conditional density of signal x given θ, satisfies the monotone likelihood ratio property (MLRP), i.e., when $f(x|\theta_H)/f(x|\theta_L)$ is monotonically increasing in x for every $\theta_H > \theta_L$.

In many settings of interest, an individual undertakes an action in order to signal that he possesses some desirable characteristic. For instance, in Spence's [15] model of the job market, a worker with exceptional capabilities may undertake extensive education in order to convince prospective employers of his superior skills.[3] To determine the outcomes that arise in such signaling games, one must specify: (i) what uninformed parties (e.g., prospective employers) infer

---

1 Levhari et al. (1975) extend Hadar and Russell's [5] analysis to the case of multidimensional random variables.

2 Rényi (1961) proposes measures of the entropy, or uncertainty, in a system. Blackwell (1951) discusses how to compare the informativeness of information structures.

3 See Riley (2001) for an informative review of the signaling literature.

from observed behavior (e.g., the level of education an individual obtains); (ii) the actions that uninformed parties will pursue in light of their inferences; and (iii) the actions that informed parties will undertake in order to shape the beliefs of uninformed parties and thereby influence their subsequent actions.

Kreps and Wilson [9] propose the sequential equilibrium concept for this purpose. Under a sequential equilibrium, beliefs induced by out-of-equilibrium actions are required to be the limit of a sequence of beliefs formed using Bayes' Rule in the presence of mixed strategies where all possible actions are undertaken with strictly positive probability. Because it can be cumbersome to identify the relevant sequence of beliefs and the corresponding mixed strategies, Cho and Kreps [10] suggest an alternative approach. The authors propose the Intuitive Criterion which states that uninformed parties believe there is no chance that an informed party would take an out-of-equilibrium action when the informed party could not possibly have benefited from the action (relative to his equilibrium action), regardless of the inference that uninformed parties draw from the action.

INFORMATION AND THE FUNCTIONING OF MARKETS

Limited information can impede the performance of markets and induce behavior that would not otherwise arise. Akerlof [13], for example, considers a setting in which each potential seller of a product (e.g., used automobiles) knows the quality of his product. In contrast, potential buyers of the product cannot observe the quality of any particular product before purchase. Risk-neutral potential buyers will purchase a product as long as its price is less than the average quality of the products offered for sale. However, potential sellers will only offer their product for sale if its quality does not exceed the prevailing market price. Therefore, the average quality of the products offered for sale will be less than the prevailing price. Consequently, there may be no market clearing price above the minimum quality.

Klein and Leffler [19] offer a related insight. The authors observe that firms will have strong incentives to reduce their production costs by delivering low-quality products to consumers who cannot fully assess product quality before purchasing an item.[4] This is the case even if consumers never purchase again from a firm that sells a low-quality product. Furthermore, the threat of entry by alternative suppliers may fail to discipline incumbent suppliers because the entrants will face corresponding incentives to deliver low-quality products to consumers.

Rothschild and Stiglitz [16] analyze a setting in which risk-averse individuals are privately informed about the probability that they will suffer an accident. Risk neutral, profit-maximizing insurance companies compete to serve consumers. If an equilibrium exists, it will be one in which low-risk consumers credibly signal their lower accident probability by choosing a (lower cost) contract that provides only partial insurance against the loss from the accident. If the fraction of low-risk customers is sufficiently high, though, an equilibrium may not exist. In this case, a firm may be able to deviate profitably from the putative equilibrium by offering a single contract that attracts both low-risk and high-risk consumers.

---

[4] See Dulleck and Kerschbamer (2006) for an overview of a related literature on credence goods.

Grossman and Stiglitz [18] show that an equilibrium may not exist in a market for risky securities when the price of the security reflects nearly perfectly the information about asset value that a trader acquires at small personal cost. In this setting, a trader may find it profitable to become informed if no other trader is informed, but an informed trader would prefer to avoid the cost of acquiring information that benefits all traders. More generally, the authors show that the equilibrium price of a risky asset will reflect partially, but not fully, asset value information that informed traders acquire at personal cost. If the price of the asset were to fully reveal this information, no party would be willing to incur the cost required to learn about the underlying value of the security.

Stiglitz and Weiss [20] analyze a setting in which lenders are unable to fully assess the credit-worthiness of borrowers. In this setting, lenders may not raise the interest rate in order to reduce an excess demand for loans. Instead, lenders may impose credit rationing. This is the case because a higher interest rate may deter the most credit-worthy borrowers (who are most likely to generate the returns required to repay the loan) while attracting less credit-worthy borrowers (who are relatively unlikely to ultimately secure the resources required to repay the loan). Higher interest rates may also induce borrowers to undertake projects that are unduly risky.

In contrast to settings where certainty prevails, prices and quantities may not be equivalent instruments in the presence of limited information. Weitzman [7] analyzes a setting where, before a supplier learns his production cost, a policymaker must set either a unit price for a commodity or specify a particular level of the commodity that must be supplied. In the former case, the firm supplies the profit-maximizing level of the commodity after learning his production cost. Weitzman shows that the relative merits of setting a price and specifying a quantity vary with the relative curvatures of the prevailing benefit and cost functions.

Williamson [14] suggests that the problems created by limited information sometimes may be ameliorated by moving transactions within the boundaries of an organization. Williamson observes that within an organization, more harmonized interests can prevail, additional control instruments (including promotion, demotion, and resource allocation) typically are available, and superior access to relevant information (before, during, and after relevant transactions are consummated) often prevails.

CONTRACTING WITH MORAL HAZARD

Holmstrom [21] analyzes the canonical moral hazard problem in which a principal seeks to motivate a risk-averse agent to supply unobserved effort. Higher levels of effort increase the likelihood that higher levels of the performance that the principal values will be realized.

The principal and the agent share the same information about the stochastic relationship between the agent's unobserved effort and his observed performance. The principal designs a contract that specifies the payment she will deliver to the agent for each possible performance realization. The agent accepts the contract as long as it provides an expected utility in excess of his (known) reservation utility. The agent then chooses his effort supply. After the agent's realized performance is observed, the principal delivers the corresponding payment specified in the contract.[5]

To motivate the agent to supply personally costly effort, the principal must link the agent's payment to his realized performance. Due to the stochastic relationship between effort and performance, this linkage imposes risk on the agent. The principal can elicit more effort from the agent by linking his payment more closely to his realized performance. However, such linkage imposes greater risk on the agent, and he will require more generous compensation to bear this risk.

Holmstrom characterizes the optimal contract for the principal in settings where the agent's effort choice is fully characterized by the point at which his expected marginal utility of effort is zero. Rogerson [23] provides conditions under which this "first-order approach" necessarily identifies the optimal contract for the principal.[6] The two primary conditions are the monotone likelihood ratio property (MLRP) and diminishing expected returns to the agent's productive effort.[7]

If the agent were risk-neutral and not wealth constrained, the principal could sell the operation to the agent for its expected value, given the surplus-maximizing effort supply (e^*). This policy would both induce the agent to deliver e^* and eliminate the agent's rent.[8] However, this policy is not feasible when the agent has little or no initial wealth. Innes [24] analyzes a moral hazard problem in which the principal and agent are both risk neutral and the agent has no initial wealth. Innes identifies conditions under which the optimal contract is a debt contract whereby: (i) the agent makes a fixed payment to the principal as long as realized profit exceeds a threshold value; and (ii) the principal receives the entire realized profit when this profit is below the threshold value.[9]

Holmstrom [22] analyzes how to motivate a team of agents to work diligently. Only the aggregate performance of the team is observable, not the individual contributions of any individual team member.[10] Holmstrom shows it is not possible to induce all members of the team to deliver the surplus-maximizing level of effort when the entire output produced by the team must be

---

5 Holmstrom's [21] model considers a single, non-repeated interaction between the principal and agent. See Bolton and Dewatripont (2005) for an insightful discussion of the literature on repeated moral hazard.

6 Mirrlees (1999) shows that the first-order approach to solving the principal's problem may not always be appropriate. He also shows how a principal can sometimes approximate the outcome that she achieves when she can observe the agent's effort.

7 See Jewitt (1988), Sinclair-Desgagné (1994), and Conlon (2009) for additional sufficient conditions. Grossman and Hart (1983) avoid the first-order approach by considering settings in which the agent can choose his preferred effort from a finite set of possible efforts.

8 See Shavell (1979).

9 Poblete and Spulber (2012) provide a corresponding conclusion under less stringent conditions than those considered by Innes.

10 Itoh (1991) and Che and Yoo (2001), among others, analyze settings where the individual performance of team members is observable.

distributed to its members.[11] In essence, the penalties required to deter "free-riding" are not available when the team cannot credibly threaten to destroy some of the realized surplus should this surplus fall below the surplus-maximizing level.

Holmstrom and Milgrom [25] consider how to motivate a risk-averse agent to perform multiple tasks.[12] The authors show that when it is difficult to measure success on one task accurately, an agent may be induced to devote more effort to this task by reducing her marginal reward for success on other tasks that are more readily measured. For example, a home remodeler may be motivated to devote attention to enhancing the quality of his work) (which is highly valued but difficult to measure) by not explicitly rewarding him for timely completion of the job (which is more readily measured, but perhaps of less importance).

TOURNAMENTS

Lazear and Rosen [26] analyze the design of reward structures in which payments only reflect the relative performance (i.e., the ordinal ranking) of agents rather than cardinal measures of their performance. Cardinal structures can be problematic when output is difficult to measure or when its measurement is subject to misrepresentation (because, for example, a principal might understate observed performance in order to justify lower payments to agents).

The authors show that a large incremental reward for winning the tournament can be optimal even when the equilibrium efforts and performances of winners and losers are very similar. The large incremental reward for the best performance can motivate diligent effort by all participants in the tournament.[13] The authors also conclude that a systematic ranking of the performance of tournaments and piece-rate contracts (contract under which the payment to an agent depends only on his own realized performance) is not possible when the agents are risk averse. This is the case because even though the payment structure is restricted under tournaments, tournaments provide valuable insurance against the extreme payment variation that can arise under a piece-rate contract. In essence, by rewarding agents on the basis of their relative performance, tournaments correct automatically for large shocks that affect the performance of all agents. In doing so, tournaments effectively eliminate variation in payment for factors beyond the agents' control, and thereby provide useful insurance for risk-averse agents.[14,15]

Nalebuff and Stiglitz [27] provide corresponding conclusions in a setting where agent i's output is $Q_i = \theta e_i + \varepsilon_i$, where θ is a common productivity shock, e_i is agent i's effort, and ε_i is

---

[11] Different conclusions can arise when the agents are risk averse (Rasmusen, 1987) or when team members can post forfeitable bonds before joining the team (McAfee and McMillan, 1991).

[12] Bond and Gomes (2009) analyze a related setting in which a risk-neutral agent has limited wealth.

[13] Thus, promotion to the top position in an organization may optimally be accompanied by a substantial increase in salary, even though the selected individual may exhibit skills and accomplishments very similar to those of his former peers. Waldman (1984, 2013) offers an alternative explanation for such policies.

[14] Green and Stokey (1983) provide corresponding insights.

[15] Chen (2003) observes, though, that tournaments can provide incentives for participants to sabotage the performance of other participants.

an idiosyncratic shock for agent $i \in \{1, \dots, n\}$. The authors also show that a pure-strategy equilibrium may not exist is this setting when there is little noise (so the variance of ε is small).[16] This is the case because, starting from the common interior effort supply at a proposed equilibrium, an agent can guarantee that he wins the contest by increasing his effort by a small amount, and a small prize for winning the tournament can induce the surplus-maximizing effort. However, when the prize for winning is small, the best response of the other agent can be to deliver no effort (since he thereby avoids all effort costs and forfeits only a small prize).

CONTRACTING WITH ADVERSE SELECTION

In the standard moral hazard model, the principal and the agent share the same information about the prevailing environment at the start of their interaction. In contrast, in the standard model of adverse selection, the agent has better information than the principal about a key element of the prevailing environment from the outset of their interaction. Therefore, while investigations of moral hazard problems often focus on the trade-off between providing stronger incentives and imposing greater risk on a risk-averse agent, the standard investigation of an adverse selection problem considers how to limit the rent the agent commands from his privileged information.

Mirrlees [28] considers a setting in which each citizen (i.e., each agent) is privately informed about his labor productivity. Citizens find it costly to supply labor, but the income that his labor generates enables a citizen to purchase a valued commodity. The government (i.e., the principal) can observe each citizen's realized income, but cannot monitor the number of hours that a citizen worked in order to generate that income. The government designs an income tax to maximize some function of the expected utility of citizens. Mirrlees demonstrates that it is often optimal to set the tax rate at a level that induces the least able workers not to work at all. Furthermore, even though a high tax rate can substantially reduce labor supply without securing a great deal of income to redistribute to less capable citizens, a relatively high tax rate can be optimal when there is substantial variance in the dispersion of ability.[17]

Mussa and Rosen [17] analyze an adverse selection problem in which a monopolist can sell products of varying quality. The monopolist's unit cost of production increases at an increasing rate with product quality. A consumer with marginal valuation $\theta \in [\underline{\theta}, \overline{\theta}]$ derives utility θq from each unit of a product of quality q he consumes. The monopolist cannot observe a consumer's valuation, but knows the density of these valuations across consumers. The authors show that the monopolist optimally reduces below the surplus maximizing level the quality of the products purchased by consumers with the lower valuations, $\theta < \overline{\theta}$. These distortions reduce the attraction

[16] Technically, the model converges to an all-pay auction with complete information as the noise vanishes. As discussed below, all-pay auctions have been used to model innovation games.

[17] Seade (1977) shows that the optimal marginal income tax rate is zero for the highest earner quite generally. Diamond and Saez (2011) explain why relatively progressive tax rates may be valuable in practice.

of the low-quality products to consumers who value quality more highly, and so the monopolist can charge these consumers more for the higher-quality products they purchase.[18]

Baron and Myerson [63] analyze a setting in which a regulator seeks to maximize a weighted average of the consumer surplus and profit generated in a regulated industry, subject to ensuring that the firm secures nonnegative profit. The monopoly supplier in this industry is privately informed about its constant marginal cost of production (c), whereas the firm's fixed cost of production and industry demand are common knowledge.

The firm's incentive in this setting is to exaggerate c in an attempt to convince the regulator that it requires substantial revenue in order to operate profitably. To limit this incentive, the regulator reduces below the surplus-maximizing level the output she authorizes the firm to produce when the firm claims to have a high marginal cost of production. The reduced output and corresponding lower authorized revenue limit the profit the firm anticipates from exaggerating c.[19]

In contrast to Baron and Myerson [63], Laffont and Tirole [29] consider a setting in which the firm's realized cost is observable. However, the effort the firm supplies to reduce its realized unit cost is unobservable, as is the firm's innate unit cost (which is the firm's cost in the absence of cost-reducing effort). Because realized cost is observable, there is no benefit to reducing output below the surplus-maximizing level, given the realized cost. However, it is often optimal to induce the firm to supply less than the surplus-maximizing level of cost-reducing effort. Doing so limits the firm's ability to substitute an innate cost advantage for cost-reducing effort, and thereby limits the firm's rent.

Lewis and Sappington [30] extend Baron and Myerson's [63] analysis to allow the monopoly supplier to be privately informed about both its fixed cost (F) and its marginal cost of production (c). These two costs are inversely related, so the firm no longer necessarily has a systematic incentive to exaggerate c. The authors show that the countervailing incentives to exaggerate c and to understate F that arise in this setting can induce outputs in excess of surplus-maximizing levels and to regions of "pooling" in which output does not vary with the realization of c.[20]

Martimort [65] extends these basic models of adverse selection to consider cases in which multiple principals simultaneously offer contracts to a common agent. For example, manufacturers might offer contracts to a common retailer who is privately informed about the demand for the manufacturers' products. Externalities among contracts lead to distortions that can differ from those that arise with a single principal. The nature of the distortions depends in part on whether the products of the manufacturers are complements or substitutes.

---

[18] Maskin and Riley (1984) extend Mussa and Rosen's [17] analysis to allow consumers to purchase more than one unit of the monopolist's product.

[19] Guesnerie and Laffont (1984) characterize the solution to a broad range of adverse selection problems, including those in which standard, simplifying regularity conditions are not satisfied.

[20] Maggi and Rodriguez-Clare (1995) and Jullien (2000) provide important generalizations. Lewis and Sappington (1988) show that extensive pooling can arise when a monopolist is privately informed about industry demand rather than industry production costs.

The adverse selection literature examines many other variations of these basic models,[21] including settings in which: (i) the agent's private information is multidimensional;[22] (ii) the principal, not the agent, has private information about the prevailing environment;[23] (iii) the information structure is endogenous;[24] and (iv) the interaction between the principal and agent is repeated over time.[25]

When the principal and agent interact repeatedly, the principal may be unable to commit herself to not use against the agent in subsequent interactions information about the prevailing environment that she infers from earlier interactions. For example, the principal may naturally increase future performance thresholds when she observes that the agent is able to achieve challenging targets in early interactions. Recognizing that the principal will engage in such ratcheting of performance standards, the agent may intentionally limit the success that he achieves in early interactions, as Freixas et al. [64] report.[26]

CONTRACTUAL INCOMPLETENESS

The models of moral hazard and adverse selection identified above presume that contracts are complete in the sense that every possible outcome of the interaction between the principal and agent is anticipated and fully accounted for in the contract that governs their relationship. In practice, though, contracts often are incomplete, i.e., they fail to fully specify all relevant contingencies.

Contracts may be incomplete in practice for many reasons. For instance, bounded rationality may preclude *ex ante* recognition of all relevant outcomes in a contractual relationship (Simon, 1981). Alternatively, it may be too time-consuming or too costly to specify all relevant outcomes, even if they can be anticipated (Williamson, 1975; Dye, 1985). One party to a contract also may decline to mention a relevant contingency in order to avoid unfavorable inferences by other parties to the contract (Spier, 1992). In addition, variables that parties to a contract can readily manipulate may intentionally be omitted from the contract (Allen and Gale, 1992).[27]

Bernheim and Whinston [32] offer another explanation for contractual incompleteness. The authors consider a setting in which it is not possible to contract explicitly on an important action that one party to a contract might undertake (e.g., an employee's effort). In this setting, it may be optimal to intentionally omit from the contract a detailed specification of the actions that

---

21 Armstrong and Sappington (2007) review the literature on regulatory policy design in the presence of limited information.

22 See, for example, Armstrong and Rochet (1999), Rochet and Stole (2003), and Basov (2005).

23 See Maskin and Tirole (1990, 1992), Beaudry (1994), and Severinov (2008), for example.

24 See, for example, Townsend (1979), Baron and Besanko (1984b), Dye (1986), Crémer and Khalil (1992), Khalil (1997), Lewis and Sappington, (1997), Crémer et al. (1998a,b), and Szalay (2009).

25 Baron and Besanko (1984a) and Krishna et al. (2013) consider settings where the principal can commit at the outset to the policies that she will pursue in all subsequent interactions with the agent.

26 See Laffont and Tirole (1988; 1993, chapter 9), for example, for additional analyses of repeated adverse selection problems when the principal has limited commitment powers.

27 See Aghion and Holden (2011) for an insightful review of the literature on incomplete contracts.

another party must undertake (e.g., the tasks that the employer will assign to the employee). Such intentional contractual incompleteness can leave the employer with an effective means to discipline the employee, should the employee deliver an undesirably low level of effort. For example, the employer may assign the employee to an unpleasant task if he fails to deliver the desired level of effort when he is assigned to a more pleasant task. Thus, when a contract cannot explicitly preclude undesired behavior, contractual incompleteness can help to deter such behavior by admitting credible punishments should the undesired behavior arise.

Grossman and Hart [31] observe that contractual incompleteness can have important implications for asset (e.g., firm) ownership. The authors view the ownership of an asset as incorporating the right to determine how the asset will be deployed in situations where its deployment is not explicitly specified in a contract (perhaps because of the prohibitive cost of specifying all relevant contingencies). The authors observe that when ownership and the associated exercise of control rights enables an individual to secure a large share of the surplus generated by the firm, then assigning ownership to the individual whose non-contractible investment has the greatest impact on this surplus can induce desirable investment patterns.

AUTHORITY AND COLLUSION IN ORGANIZATIONS

Most of the models of adverse selection identified above consider settings in which only the agent can observe a key element the environment, e.g., the agent's production cost. Tirole [33] considers a setting in which a supervisor can also sometimes observe the agent's cost. If he chooses to do so, the supervisor can share his knowledge of the agent's cost with the principal. However, the agent may bribe the supervisor to conceal evidence that he has a low cost, and thereby secure more generous reimbursement from the principal. Tirole analyzes how the presence of the supervisor affects both the optimal contract between the principal and the agent and the agent's equilibrium performance.[28]

Aghion and Tirole [34] extend Grossman and Hart's [31] analysis of ownership and control by analyzing a distinction between formal and real authority within an organization. The authors note that although a principal may have the formal authority to make a decision within an organization (e.g., to exercise control rights), the principal may delegate the actual right to make this decision to an agent. By doing so, the principal can create incentives for the agent to acquire valuable planning information. The authors show that a principal is particularly likely to cede decision-making authority to an agent when the interests of the principal and agent are relatively similar, when the agent is more adept than the principal at acquiring information, and when the outcome of the decision is of particular consequence for the agent.[29]

---

[28] Other important analyses of collusion within organizations include Laffont and Martimort (1997, 1998, 2000) and Faure-Grimaud et al. (2003).

[29] Other studies of optimal delegation of decision-making authority include Holmstrom (1984), Alonso and Matouschek (2008), Bester and Krähmer (2008), and Koessler and Martimort (2012). Mookherjee (2006) provides a useful discussion of the literature on delegation and organizational structure.

CONTRACT RENEGOTIATION

Even when contracting frictions permit only highly restrictive contracts, the freedom to renegotiate the terms of a contract as valuable planning information arrives can sometimes motivate efficient investment decisions. Hart and Moore [35] consider a setting in which an initial contract between a buyer and a seller can only specify a payment if the buyer purchases a product from the seller and another payment if no transaction occurs. After signing the initial contract, the buyer makes a noncontractible investment that affects her valuation (v) of the product and the seller makes a noncontractible investment that affects his cost (c) of supplying the product. After observing the realizations of v and c, the parties can renegotiate the original contract.

The authors identify conditions under which the surplus-maximizing investments and transactions can be ensured.[30] The surplus-maximizing outcome does not arise in general, though, because of a fundamental investment externality. The externality arises because, for example, if the buyer reduces her investment below the surplus maximizing level, her valuation declines. This reduced valuation and corresponding reduced willingness to pay for the product reduce the seller's incentive to invest.

More generally, the possibility of contract renegotiation can reduce welfare below the level that is secured when contract renegotiation is precluded. This is the case because it is sometimes optimal, when feasible, for a principal to commit herself not to use against the agent in future interactions information that the principal gleans from earlier interactions. The freedom to renegotiate contracts can undermine this valuable commitment.[31]

AUCTIONS

Auctions are commonly employed to transfer objects from sellers to buyers in the presence of limited information. Auctions differ with regard to payment rules and the knowledge of the bids of rival buyers ("bidders"). To illustrate, the high-bidder pays his bid in a first-price auction but pays the second-highest bid in a second-price auction. Furthermore, rivals' bids are observed in English auctions but unobserved in sealed-bid and Dutch auctions.[32] Vickery [11] shows that the Dutch and first-price sealed bid auctions are strategically equivalent. He also demonstrates that a bidder optimally bids his true value of the item in a second-price auction.[33] For the case of symmetric bidders with independent, uniformly distributed valuations, Vickrey proves that the English, Dutch, first price, and second-price auctions all generate the same expected revenue for the seller.

Myerson [36] considers an environment where a seller has incomplete information about n bidders' valuations (or value estimates) for a single object. These valuations are independent, and their densities are continuous, but are otherwise unrestricted. Myerson employs the revelation

---

[30] See Hermalin and Katz (1991) for a related observation.

[31] See, for example, Dewatripont (1988, 1989), Aghion et al. (1994), and Dewatripont and Maskin (1995).

[32] See McAfee and McMillan (1987a), Milgrom (1989), and Klemperer (2002) for insightful surveys of auction theory.

[33] Similar mechanisms have been independently constructed to solve information and free-rider problems in the provision of public goods; see Clarke (1971), Groves (1969, 1973) and Groves and Ledyard (1977).

principle (see below) to derive a revenue equivalence theorem for auctions in this broad class of environments. Riley and Samuelson [37] independently derive a similar result using a direct (calculus-based) proof for the symmetric case with strictly increasing bid functions. A key implication of these two papers is that the standard auctions studied by Vickery [11] are *not* optimal auctions from the seller's standpoint. Optimality requires the seller to commit to a reserve price (or minimum acceptable bid) that strictly exceeds the seller's own valuation of the item. Consequently, an auction that maximizes the seller's expected revenue is not generally Pareto efficient. Even when bidders are symmetric *ex ante*, there is a positive probability that the item will not be awarded to the party that values it most highly.

Milgrom and Weber [38] consider auction environments where bidders have *affiliated* rather than stochastically independent valuations (or value estimates). Intuitively, the valuations of two bidders are affiliated if a high valuation for one bidder increases the likelihood that the other bidder's valuation is high. For the case in which the joint density $f(z_1, z_2, \ldots, z_n)$ of the valuations (z_i) of n bidders is strictly positive and twice continuously differentiable, z_i and z_j are affiliated if and only if $\partial^2 \ln f(\cdot) / \partial z_i \partial z_j \geq 0$. Milgrom and Weber show that in symmetric environments where risk-neutral bidders are privately informed about their own valuations, a seller earns higher expected revenue in an English auction than a second-price auction, and lower expected revenue in a first-price auction. Additionally, the authors show that the seller enhances expected revenue by disclosing any private information that he might have to the bidders. These findings reflect the *linkage principle*: Auction forms and disclosure policies that link the price paid to variables that are correlated with the winner's valuation (e.g., all other bids in an English auction, the second-highest bid in a second-price auction, and seller information in the case of disclosure) increase the seller's expected revenues.[34]

Search engines such as *Google* and *Bing* employ auctions to allocate scarce advertising space. Advertisers effectively bid for the positions in which their ads will appear following a keyword search.[35] Since consumers tend to click ads that are closer to the top of the screen, higher screen positions are more valuable to advertisers than lower positions. Edelman et al. [40] examine the generalized second-price auction (GSP), which is used to price such advertisements. Each advertiser submits the amount it is willing to pay if a consumer clicks its ad following a keyword search. The engine then positions the ads according to the order of the bids. If a user clicks an ad in position j, the advertiser does not pay his own bid but rather the amount bid by the advertiser in position $j + 1$. Notice that in the special case where only one ad is ultimately placed on the screen, this mechanism corresponds to the second-price (Vickrey) auction (where the dominant strategy of a player is to bid his true value of the item). The GSP essentially generalizes the second-price auction to the case of multiple items (positions). Unlike the Vickery auction, the GSP does not generally have an equilibrium in dominant strategies and bidders do not optimally bid their true valuations.

---

[34] Other issues examined in this literature include bidding rings (McAfee and McMillan, 1992), budget-constrained bidders (Che and Gale, 1998), multi-unit auctions (Perry and Reny, 1999), uncertainty regarding the number of bidders (McAfee and McMillan, 1987b), and entry (Levin and Smith, 1994; McAfee and McMillan, 1987c).

[35] See also Varian (2007) and Athey and Ellison (2011).

Bulow and Klemperer [39] suggest that sellers often prefer standard auctions (without reserve prices) to bargaining or optimal auctions (with reserve prices) because standard auctions encourage participation by more bidders than these alternative mechanisms. Roughly, the authors show that attracting an additional bidder in an auction generates more revenue than setting an optimal reserve price but attracting one less bidder. Since an auction with an optimal reserve price is an optimal mechanism, it then follows that an auction with $N + 1$ bidders dominates an optimal auction with N bidders, and hence, any alternative mechanism with N bidders (including optimal negotiation).

BARGAINING

In many bargaining environments, individuals have incomplete information—that is, they lack information about each other's preferences. For example, a car dealer does not typically know the maximum amount a consumer is willing and able to pay for a new car and the consumer does not know the minimum price the dealer will accept. Mechanism design and the revelation principle are powerful tools for analyzing these and related environments.

In a direct mechanism, players independently report their valuations and the mechanism determines allocations and payments. A direct mechanism is *incentive compatible* if, in a Bayesian-Nash equilibrium, players report their valuations truthfully. Myerson [41] establishes the *revelation principle*: For a general class of games of incomplete information that includes auctions and bargaining, there is no loss of generality in restricting attention to incentive compatible direct mechanisms. In the context of bargaining, this means that one may analyze the outcomes of an arbitrary bargaining game by considering a mechanism in which each player truthfully reports his valuation to an *arbiter* who then determines the allocation and payments. Myerson also shows that when players bargain with incomplete information about one another's valuations, there is generally a positive probability that the players will fail to reach an agreement.

Myerson and Satterthwaite [42] show that if an outside party subsidizes the bargaining process, negotiation can lead to outcomes that are *ex post* efficient. They also construct the minimum lump sum subsidy required to induce such outcomes. Watson (1998) provides an interesting model of bargaining under incomplete information in which players make alternating offers.

SEARCH

Following Stigler's [12] seminal work on the economics of information, economists have devoted considerable attention to analyzing wage and price dispersion. To find better opportunities, workers must search for higher-paying jobs and consumers must search for lower product prices.[36,37] Because such search is costly, rational individuals generally will not obtain full information about the wages offered and prices charged by different firms.

McCall [43] employs these insights to develop a model of job search in which a worker knows the (cumulative) distribution of wages, $F(w)$, but incurs a search cost of c to generate each

---

36 Lippman and McCall (1976) and Rogerson et al. (2005) survey the job search literature. Rothschild (1973) and Baye et al. (2006) provide overviews of the literature on product search.

37 Roth (1989) considers the related problem of matching with incomplete information.

job offer. A worker utilizing a sequential search strategy compares the expected benefit of generating another job offer with the associated cost. This gives rise to a reservation wage, w^*, defined implicitly by

$$\int_{w^*}^{\infty} (w - w^*)\, dF(w) = c\,.$$

This decision rule leads the worker to reject any job offer with a wage below w^* (since the expected benefit of generating another offer exceeds the cost of another search) and to accept any higher wage offer. This framework has proved useful for distinguishing between discouraged workers and dropouts from the frictionally unemployed.[38]

An analogous stopping rule arises when a consumer (with unit demand) sequentially searches for price information in a market for a homogeneous product. Assuming a known price distribution, $F(p)$, and a unit search cost of c, a consumer optimally stops searching when he observes a price at or below the reservation price, p^*, which is determined by:

$$\int_{0}^{p^*} (p^* - p)\, dF(p) = c\,.$$

A potential drawback of these formulations is the assumption that searchers know the distribution of offers. However, Rothschild [45] shows that search behavior is qualitatively similar when $F(\cdot)$ is a prior distribution and searchers update their beliefs as they gather more information about the distribution of offers. However, the reservation price (or wage) may change during the search process as searchers update their beliefs.

While the models identified above assume that searchers act rationally, the models take the distributions of offers to be exogenous and do not take into account the incentives of firms. In contrast, Diamond [44] presents a dynamic model in which consumers visit each firm once but where firms set prices optimally, knowing the consumers' (exogenous) strategy. Diamond shows that prices converge to the monopoly price. This result, known as the *Diamond paradox*, also obtains in a static model where identical consumers search optimally and identical firms set prices to maximize their expected profits.

The Diamond paradox raised fundamental questions about the ability of search theory to rationalize the ubiquitous price dispersion documented by Stigler [12] and others. While this paradox is based on the assumption that firms have identical costs, Rothschild (1973) observes that even if firms have heterogeneous costs, a profit maximizing firm will never set a price above the reservation price. Such a response to a given reservation price would presumably induce consumers to lower their reservation price, which would induce firms to lower prices further. In short, Rothschild observes that, taking into account incentives of consumers *and* firms, it is not apparent that search costs will produce *equilibrium* price dispersion.

Reinganum [46] presents a simple but elegant model that resolves both the Diamond paradox and Rothschild's criticism. She models an environment where identical consumers with

38 See, for example, Mortensen and Pissarides (1994), Burdett and Mortensen (1998), and Mortensen (1970).

downward sloping iso-elastic demands engage in optimal sequential search[39] and where many profit-maximizing firms sell identical products but produce at different unit costs (drawn from a continuum of possible costs).[40] Since consumers have identical iso-elastic demands and the same search costs, all consumers have the same reservation price. Cost variation ensures that different firms have different monopoly prices. Reinganum shows that, in equilibrium, firms with monopoly prices above the reservation price have an incentive to price at the reservation price, thus leading to a mass of firms charging this price. But this truncation in the distribution of monopoly prices does not change the reservation price because the expected benefit of search is equal to the cost of another search at the reservation price. Thus, price dispersion can arise as a result of search costs with fully optimizing consumers and firms. It is now widely recognized that equilibrium price dispersion arises in a variety of search environments.[41]

Stigler's [12] original model of search assumes that consumers utilize a fixed sample search strategy; that is, they commit to sample N stores and then purchase at the lowest observed price. This is in contrast to sequential search, where a consumer continually determines whether to search again based on the expected benefits and costs of doing so. Morgan and Manning [47] show that there are instances in which fixed-sample size search is optimal and instances in which sequential search is optimal. The primary advantage of fixed-sample search is that it permits relatively rapid information acquisition. For example, if it takes one week for employers to respond to a job applicant, a searcher who sends applications to five potential employers receives five responses in one week, whereas sequential search would produce only one response in a week. The primary advantage of sequential search is that it avoids overinvestment in information acquisition. In general, optimal search combines features of these two search strategies.

THE PRICING OF INFORMATION

Consider a monopolist that operates a clearinghouse for the transmission and acquisition of information (e.g., a price comparison site). To maximize its profit, should the monopolist charge high fees to transmit information and low fees to access the information, or *vice versa*? Baye and Morgan [48] answer this question in the context of a two-sided market for price information.[42]

In their model, $N > 1$ geographically separated towns are each serviced by a local firm, and search costs dissuade optimizing consumers in one town from searching at stores in other towns. A third player – the information gatekeeper – operates a virtual marketplace. A firm advertising its price at the gatekeeper's website gains access to consumers in distant towns; consumers visiting

---

[39] MacMinn (1980) shows that equilibrium price dispersion also arises when consumers employ a fixed-sample search strategy as considered in Stigler [12].

[40] Stahl (1989) provides a corresponding analysis in a model with a small number of firms.

[41] Dana (1994) identifies such an equilibrium when consumers learn about price distributions during search. Arbatskaya (2007) establishes equilibrium in an environment where there is a natural order in which consumers sample firms.

[42] For related pricing issues in two-sided markets, see, e.g., Rochet and Tirole (2003, 2006), Caillaud and Jullien (2003), Armstrong (2006), and Armstrong and Wright (2007). Varian (2000) and Sundararajan (2004) provide interesting insights into the pricing of information goods such as computer software, books, CDs, DVDs, as well as digital media.

the site gain access to the list of advertised prices and benefit if they find a price lower than the price charged by their local firm. The gatekeeper recognizes that information is valuable and charges access fees to consumers and firms that use its site. Baye and Morgan show that the gatekeeper maximizes its profit by charging consumers low (possibly zero) access fees to attract the largest possible set of "eyeballs" to its site, and charging higher fees to the firms that advertise on its site.[43]

ADVERTISING AND PRICE SIGNALING

Nelson (1970) distinguishes between a "search good" (e.g., a shirt), whose quality a consumer can determine prior to purchase and an "experience good" (e.g., a can of tuna fish), whose quality cannot be determined until after purchase. Nelson [49] observes that advertisements for search goods often provide direct product information, while advertisements for experience goods typically do not. Nonetheless, expenditures on promoting the brand of an experience good may signal to consumers that it is, indeed, of high quality.

Milgrom and Roberts [50] analyze this issue formally by allowing firms to potentially signal their qualities through both advertising and price. Consistent with Nelson, their model shows that seemingly wasteful expenditures on advertising can serve an important signaling function. Since customers who learn that an experience good is of high quality are likely to make repeat purchases, an initial sale of the good is more valuable to a firm that sells a high-quality product. Consequently, such a firm is willing to spend more on advertising to attract an initial sale than is a firm that sells a low-quality product, which will experience fewer repeat sales. Typically, the seller of a high-quality product also will distinguish itself by setting a relatively high price. If the high price reduces quantity demanded sufficiently, a seller of a less expensive, low-quality product will earn a larger profit by selling a large quantity at a low price than by matching the high price set by the seller of a high-quality product.[44]

COMMUNICATION

Crawford and Sobel [51] analyze the interaction between a sender (S) who has private information about the state (θ) and a receiver (R) who is the decision-maker. The sender sends a costless, unverifiable message about the state (i.e., "cheap talk") to the receiver, who then makes a decision that affects the utility of both players. For example, the utilities of the sender and receiver might be $U^S = -\left(x - (\theta + b)\right)^2$ and $U^R = -(x - \theta)^2$, respectively. Here, x is the receiver's decision, θ is a uniformly distributed random variable, and $b > 0$ is a parameter that allows the preferences of the sender and the receiver to differ.

---

[43] These fees result in equilibrium price dispersion at the gatekeeper's site, despite the absence of heterogeneity in consumers' search costs or firms' production costs.

[44] Wolinsky (1983) considers price signaling as a moral hazard problem. Bagwell and Riordan (1991) assume that consumers learn quality over time so that the equilibrium price signal becomes less distortionary. Judd and Riordan (1994) allow consumers to have private information about their valuation of a product's quality, and show that price signaling may be credible even without cost differences.

The key issue is the degree of information that can be credibly communicated in this environment. Clearly, there exists a "babbling" equilibrium in which the receiver ignores any message and bases his decision only on his imperfect *ex ante* knowledge of θ. Crawford and Sobel show that there also exist equilibria in which the receiver acts on informative messages from the sender. However, all such equilibria involve information loss, as the sender partitions the state space into intervals and simply reports the interval in which θ lies.

Battaglini [52] shows that full information transmission can arise in an extended framework with a multi-dimensional state variable and multiple senders.[45] Farrell and Gibbons (1989) show that simultaneous communication to multiple receivers can either facilitate or hinder communication. Dewatripont and Tirole [53] analyze information transmission in a setting where the effectiveness of communication depends on the sender's effort to communicate clearly and on the receiver's effort to understand the information he receives.

MODELING INDUSTRY OUTCOMES

Formal models of oligopolistic rivalry often predict that industry outcomes vary substantially according to whether firms set prices (Bertrand) or quantities (Cournot).[46] Klemperer and Meyer [54] construct a model where firms choose supply schedules, which specify the quantity that each firm is willing to supply at different prices.[47] In contrast to environments where firms set either quantities or prices, these strategies give firms flexibility in responding to exogenous uncertainty in demand. In contrast to the case of certain demand, which can support multiple equilibria and associated coordination problems, the set of equilibria is substantially reduced (and under some conditions, unique) when there is demand uncertainty.

The Bertrand paradox indicates that, in the absence of any uncertainty, price competition among homogeneous product firms leads to marginal cost pricing and zero economic profit. Spulber [55] shows that this result critically depends on the absence of uncertainty. He considers an oligopoly price-setting environment where firms are asymmetrically informed about costs. Each firm's unit cost is an independent draw from a non-degenerate continuous distribution, and each firm knows its own cost but not rivals' costs. Even though the firm setting the lowest price captures the entire market demand, equilibrium entails each firm charging a price above its unit cost and earning positive expected profits. This is the case because firms are essentially competing in a first-price auction with independent valuations. Since a firm pricing at its unit cost always earns exactly zero profit, and since each firm perceives a positive probability that all other firms have higher unit costs, it is profitable to submit a bid (charge a price) in excess of its marginal cost.

---

45 With a one-dimensional state variable and multiple senders, Krishna and Morgan (2001) find that communication can sometimes be fully revealing. For a multi-dimensional state variable with a single sender, Chakraborty and Harbaugh (2007, 2010) find that coarse communication is possible when no communication would be possible in a single dimension, and that communication can become arbitrarily informative as the dimensionality of the state variable increases.

46 Klemperer and Meyer (1986) show that uncertainty may resolve the fundamental differences that arise when price rather than quantity is the strategic variable.

47 Bernheim and Whinston (1986) examine corresponding "menu auctions" in the absence of uncertainty.

There are, of course, a variety of other avenues through which information affects industry outcomes. Examples include herding and information cascades (Banerjee, 1992; Bikhchandani et al., 1992), bubbles (Blanchard, 1979), and consumer myopia (Gabaix and Laibson, 2006). Additionally, information affects industry outcomes through its impact on firms' incentives to bundle products (e.g., McAfee et al., 1989) and their ability to implement other non-linear pricing schemes.

COLLUSION

Stigler [56] provides an informal discussion of the challenges that cartels face in maintaining prices above competitive levels. These challenges include agreeing on a collusive price structure (firms generally charge a myriad of prices, and a particular structure may benefit some firms more than others) and entry (cartel success can encourage new firms to enter the market).

Stigler's [56] primary focus, however, is on the challenge of policing a cartel in the presence of imperfect information. Each firm has an incentive to engage in clandestine price reductions to attract customers from other cartel members. Stigler presents a simple model in which a firm's incentive to implement such a price reduction depends on the number of rivals the firm faces, the number of customers each rival serves, and the probability of repeat purchases. An additional problem is the difficulty of detecting price cuts which, as noted below, is essential in punishing defectors and thus maintaining collusive prices.[48]

The folk theorem for repeated games (see, e.g., Fudenberg and Maskin [59]) implies that collusive outcomes can be sustained as the Nash equilibrium of a repeated game if the game is repeated indefinitely and if players are sufficiently patient. Collusion can be sustained under these conditions due to the credible threat of severe future punishment for defection. If there is always a tomorrow and players are sufficiently patient, the loss from future punishments will exceed the current gains from defection, and so collusion can be sustained. Stigler [56] anticipated that, for such mechanisms to work, players must have a means of detecting deviations.[49]

Green and Porter [58] point out that in the presence of stochastic demand, firms will be unable to determine with certainty whether a low market price reflects a defection from the cartel agreement or a low demand realization.[50] The authors show that collusion is nonetheless feasible in a homogeneous product Cournot environment. In equilibrium each firm produces the collusive output as long as the market price exceeds a specified "trigger price." Whenever demand fluctuations force the market price below this threshold, firms revert to the Cournot equilibrium for a specified period of time and then collude once again. These reversionary "competitive" periods ensure that firms do not have an incentive to defect from the cartel.

---

[48] Gal-Or (1985), Raith (1996), and Shapiro (1986) examine the related issue of oligopolists' incentives to share information.

[49] Abreu et al. (1990) formally analyze the impact of imperfect monitoring on collusive outcomes.

[50] See Rotemberg and Saloner (1986) and Haltiwanger and Harrington (1991) for related models with uncertain profits, and Baye and Jansen (1996) for an analysis of the impact of interest rate uncertainty (stochastic discounting) on collusive outcomes.

A key assumption underlying both this finding and the folk theorem is that the game is repeated indefinitely. Standard backwards unraveling arguments imply that it is more challenging to construct models of a collusive equilibrium when players interact a known, finite number of times.[51] However, several authors have shown that informational asymmetries can permit firms to sustain collusive outcomes in finitely repeated games. For example, Kreps et al. [57] show that when players have incomplete information about rivals' payoffs or feasible actions, collusive outcomes can arise during early rounds of the finitely repeated prisoners' dilemma. Fudenberg and Maskin [59] provide folk theorems for a large class of finitely repeated games of incomplete information, and show that when there is an arbitrarily small chance that rivals are irrational, average (per period) equilibrium payoffs in a finitely repeated game can be arbitrarily close to fully collusive outcomes.

DETERRING ENTRY THROUGH REPUTATION OR LIMIT PRICING

In dynamic settings with incomplete information, players may find it profitable to undertake costly investments in reputation formation. Kreps and Wilson [60] extend Selten's (1978) chain store paradox to a setting where the incumbent and potential entrants have incomplete information about each other's payoff functions. For example, a "tough" incumbent might earn higher immediate profit by fighting the entrant, while a "weak" incumbent might earn higher immediate profit by acquiescing. The authors show that even if the entrant perceives only a slight chance that the incumbent is "tough," a sequential equilibrium exists in which the incumbent "fights" early on. Even though it is costly, such fighting is profitable for the incumbent because it helps to convince the entrant that the incumbent is tough and so will continue to fight the entrant.

Independently, Milgrom and Roberts [62] offer an approach that yields qualitatively similar results, but is somewhat simpler because it assumes the incomplete information concerns the incumbent's options rather than payoffs. They also discuss the implications of the results for predatory pricing and product strategies.

An incumbent might also deter entry through limit pricing. A monopolist that engages in limit pricing charges a price below the monopoly price in an attempt to induce entrants to stay out of the market. In order to affect the actions of entrants, the pre-entry price must be linked to expected post-entry prices. Milgrom and Roberts [61] analyze a model in which a monopolist and a potential entrant have incomplete information regarding one another's costs. Post-entry profits are lower for the entrant when the monopolist has lower costs. Consequently, the monopolist can reduce the entrant's expected post-entry profit by setting a lower initial (limit) price. Somewhat surprisingly, limit pricing can arise in both separating and pooling equilibria.

[51] Collusion may arise in finitely repeated games when the terminal period is random or when the stage game has multiple Nash equilibria. See Benoit and Krishna (1985).

REGULATION

Regulators often must design and implement policy in settings where they know less about prevailing industry conditions than do the firms they regulate. Consequently, the models of adverse selection discussed above (e.g., Baron and Myerson [63], Freixas et al. [64], and Lewis and Sappington [30]) have direct implications for the design of regulatory policy in both static and dynamic settings. Similarly, the aforementioned models of exclusive dealing, predatory pricing, collusion, and advertising (e.g., Martimort [65], Milgrom and Roberts [62], Stigler [56], and Nelson [49]) can inform the design of market regulation through antitrust and consumer protection policy.

INNOVATION

A large body of literature[52] examines the incentives of firms to undertake costly research and development to create new or improved products or to create cost-reducing technologies. These incentives are shown to vary with such factors as the prevailing market structure,[53] the length and breadth of patents,[54] and the nature of innovation.[55] Innovation and patent race games are related[56] to distinct literatures on tournaments,[57] rent-seeking contests,[58] the war-of-attrition, the all-pay auction, and contests with rank-order spillovers.[59] Within the literature on innovation, R&D has been modeled as a tournament (e.g., Taylor, 1995), as a contest (e.g., Che and Gale, 2003), and as a race (e.g., Mortensen, 1982; Jensen and Showalter, 2004).

Dasgupta and Stiglitz [66] consider a model of process innovation in which the probability that a discovery is made in date *t* follows a Poisson process in which the Poisson parameter is a function of the firm's effort. In their model, R&D is essentially a race to be first. Their overarching conclusion is that the impact of market structure on the efficiency of R&D in markets is generally ambiguous.

Reinganum [67] considers a model where an incumbent and challenger are engaged in a cost-reducing innovation game, and shows that results vary according to whether innovation is deterministic or stochastic. When innovation is deterministic, the incumbent invests more than the challenger. When innovation is stochastic, the challenger invests more than the incumbent.

---

52 Kamien and Schwartz (1975) and Lemley and Shapiro (2005) provide useful overviews of the literature.

53 See, Loury (1979), for example.

54 See Gilbert and Shapiro (1990), for example.

55 This includes whether there is learning through parallel research (Nelson, 1961), and more broadly, how research produces useful innovations (Evenson and Kislev, 1976).

56 Baye and Hoppe (2003) establish the strategic equivalence of a wide class of rent-seeking contests, innovation games, and patent race games.

57 See, for example, Laffont and Tirole [29], Lewis and Sappington [30], Green and Stokey (1983), and Fullerton and McAfee (1999).

58 See Nitzan (1994) for a survey of the rent-seeking literature.

59 See Baye et al. (1993, 2005, 2012) and Krishna and Morgan (1997), for example.

Therefore, whether preemptive patenting[60] is likely to help entrench an incumbent monopolist varies according to the nature of the innovation in question.

COMMITMENT

Brander and Lewis [68] examine the incentive effects of debt on industry output in a two-stage oligopoly game. Firms first choose their financial structure (degree of debt financing) and subsequently choose outputs. Each firm's profit is affected not only by these decisions, but also by a stochastic element (e.g., a demand shock) that is independently distributed across firms. A unilateral increase in debt expands the set of states in which a firm becomes bankrupt. This alters the firm's expected profit in the second stage, since limited liability ensures that it earns the same payoff in all states involving bankruptcy. The authors show that if marginal profit is increasing in the random shock, a firm that unilaterally uses debt financing produces more aggressively in the second-stage and increases its expected profit. Intuitively, the firm benefits because debt financing commits it to behave like a Stackelberg leader in second-stage game. Other firms have similar incentives, however. Equilibrium therefore entails all firms using debt financing and earning lower expected profits (due to the greater industry output and lower prices that prevail) compared to the environment where debt financing is not available.

Fershtman and Judd [69] provide related insights in a contracting environment where firm owners employ contracts (C_i) that are linear in profits (π_i) and sales (S_i) to motivate managers: $C_i = \alpha_i \pi_i + (1 - \alpha_i) S_i$. In the first stage, before uncertain market demand is realized, owners independently choose contract parameters (α_i). In the second stage, firms engage in either homogeneous product Cournot or differentiated product Bertrand competition. Under Cournot competition, equilibrium contracts set $\alpha_i < 1$, thereby committing managers to behave more aggressively in the product market. In contrast, equilibrium contracts entail $\alpha_i > 1$ under Bertrand competition, and owners induce managers to price less aggressively. In contrast to the quantity-setting case, commitment through contracts increases the equilibrium profits of firms above the levels achieved when managers are instructed to maximize expected profit.[61]

Bagwell [70] shows that commitment effects depend critically on the assumption that the second mover can perfectly observe the first-mover's decision. He shows that whenever there is even the slightest noise or imperfection in observing the first mover's decision, the set of sequential-move equilibria in the standard two-stage Stackelberg game collapses to the simultaneous-move equilibrium, and thus commitment effects are lost.

Authors have shown that commitment effects are restored when the first-mover enjoys private information (Maggi, 1999) or when mixed-strategies are allowed (Van Damme and Hurkens, 1997). Likewise, commitment effects can arise in agency or delegation games when unobserved contracts change the game being played among agents or delegates (e.g., Katz (1991), Fershtman and Kalai (1997), and Kockensen and Ok (2004)).

---

[60] See Gilbert and Newbery (1982).

[61] Sklivas (1987) obtains similar results in a model with no demand or cost uncertainty.

References

Dilip Abreu, David Pearce, and Ennio Stacchetti, 'Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,' *Econometrica*, September 1990, 58(5), 1041-1063.

Philippe Aghion, Mathias Dewatripont, and Patrick Rey, 'Renegotiation Design with Unverifiable Information,' *Econometrica*, March 1994, 62(2), 257-282.

Philippe Aghion and Richard Holden, 'Incomplete Contracts and the Theory of the Firm: What Have We Learned over the Past 25 Years?' *Journal of Economic Literature*, Spring 2011, 25(2), 181-197.

Franklin Allen and Douglas Gale, 'Measurement Distortion and Missing Contingencies in Optimal Contracts,' *Economic Theory*, January 1992, 2(1), 1-26.

Ricardo Alonso and Niko Matouschek, 'Optimal Delegation,' *Review of Economic Studies*, January 2008, 75(1), 259-293.

Maria Arbatskaya, 'Ordered Search,' *RAND Journal of Economics*, Spring 2007, 38(1), 119-126.

Mark Armstrong, 'Competition in Two-Sided Markets,' *RAND Journal of Economics*, Autumn 2006, 37(3), 668-691.

Mark Armstrong and Jean-Charles Rochet, 'Multi-Dimensional Screening: A User's Guide,' *European Economic Review*, April 1999, 43(4), 959-979.

Mark Armstrong and David E. M. Sappington, 'Recent Developments in the Theory of Regulation,' in Mark Armstrong and Robert H. Porter (eds.), *The Handbook of Industrial Organization*, *Volume 3*. (Elsevier Science Publishers, 2007), 1557-1700.

Mark Armstrong and Julian Wright, 'Two-sided Markets, Competitive Bottlenecks, and Exclusive Contracts,' *Economic Theory*, August 2007, 32(2), 353-380.

Susan Athey and Glenn Ellison, 'Position Auctions with Consumer Search,' *Quarterly Journal of Economics*, August 2011, 126(3), 1213-1270.

Kyle Bagwell and Michael H. Riordan, 'High and Declining Prices Signal Product Quality,' *American Economic Review*, March 1991, 81(1), 224-239.

Abhijit V. Banerjee, 'A Simple Model of Herd Behavior,' *Quarterly Journal of Economics*, August 1992, 107(3), 797-817.

David P. Baron and David Besanko, 'Regulation and Information in a Continuing Relationship,' *Information Economics and Policy*, 1984a, 1(3), 267-302.

David P. Baron and David Besanko, 'Regulation, Asymmetric Information, and Auditing,' *RAND Journal of Economics*, Winter 1984b, 15(4), 447-470.

Suren Basov, *Multidimensional Screening*. (Berlin: Springer-Verlag, 2005).

Michael R. Baye and Dennis W. Jansen, 'Repeated Games with Stochastic Discounting,' *Economica*, November 1996, 63(252), 531-541.

Michael R. Baye and Heidrun C. Hoppe, 'The Strategic Equivalence of Rent-Seeking, Innovation, and Patent-Race Games,' *Games and Economic Behavior*, August 2003, 44(2), 217-226.

Michael R. Baye, Dan Kovenock, and Casper G. de Vries, 'Rigging the Lobbying Process: An Application of the All-Pay Auction,' *American Economic Review*, March 1993, 83(1), 289-294.

Michael R. Baye, Dan Kovenock, and Casper G. de Vries, 'Comparative Analysis of Litigation Systems: An Auction-Theoretic Approach,' *Economic Journal*, July 2005, 115(505), 583-601.

Michael R. Baye, Dan Kovenock, and Casper G. de Vries, 'Contests with Rank-Order Spillovers,' *Economic Theory*, October 2012, 51(2), 315-350.

Michael R. Baye, John Morgan, and Patrick Scholten, 'Information, Search, and Price Dispersion,' in Terrence Hendershott (ed.), *Handbook on Economics and Information Systems*. (Amsterdam, 2006), 327-375.

Paul Beaudry, 'Why an Informed Principal may Leave Rents to an Agent,' *International Economic Review*, November 1994, 35(4), 821-832.

Jean-Pierre Benoit and Vijay Krishna, 'Finitely Repeated Games,' *Econometrica*, July 1985, 53(4), 905-922.

B. Douglas Bernheim and Michael D. Whinston, 'Menu Auctions, Resource Allocation, and Economic Influence,' *Quarterly Journal of Economics*, February 1986, 101(1), 1-31.

Helmut Bester and Daniel Krähmer, 'Delegation and Incentives,' *RAND Journal of Economics*, Autumn 2008, 39(3), 664-682.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch, 'A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,' *Journal of Political Economy*, October 1992, 100(5), 992-1026.

David Blackwell, 'Comparison of Experiments,' *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. (Berkeley: University of California Press, 1951), 93-102.

Olivier Jean Blanchard, 'Speculative Bubbles, Crashes and Rational Expectations,' *Economics Letters*, 1979, 3(4), 387-389.

Patrick Bolton and Mathias Dewatripont, *Contract Theory*. (Cambridge, MA: MIT Press, 2005).

Philip Bond and Armando Gomes, ‘Multitask Principal-Agent Problems: Optimal Contracts, Fragility, and Effort Misallocation,’ *Journal of Economic Theory*, January 2009, 144(1), 175-211.

Kenneth Burdett and Dale T. Mortensen, ‘Wage Differentials, Employer Size, and Unemployment,’ *International Economic Review*, May 1998, 39(2), 257-273.

Bernard Caillaud and Bruno Jullien, ‘Chicken and Egg: Competition among Intermediation Service Providers,’ *RAND Journal of Economics*, Summer 2003, 34(2), 309-328.

Archishman Chakraborty and Rick Harbaugh, ‘Comparative Cheap Talk,’ *Journal of Economic Theory*, January 2007, 132(1), 70-94.

Archishman Chakraborty and Rick Harbaugh, ‘Persuasion by Cheap Talk,’ *American Economic Review*, December 2010, 100(5), 2361-2382.

Yeon-Koo Che and Ian Gale, ‘Standard Auctions with Financially Constrained Bidders,’ *Review of Economic Studies*, January 1998, 65(1), 1-21.

Yeon-Koo Che and Ian Gale, ‘Optimal Design of Research Contests,’ *American Economic Review*, June 2003, 93(3), 646-671.

Yeon-Koo Che and Seung-Weon Yoo, ‘Optimal Incentives for Teams,’ *American Economic Review*, June 2001, 91(3), 525-541.

Kong-Pin Chen, ‘Sabotage in Promotion Tournaments,’ *Journal of Law, Economics, and Organization*, April 2003, 19(1), 119-140.

Edward H. Clarke, ‘Multipart Pricing of Public Goods,’ *Public Choice*, Fall 1971, 11(1), 17-33.

John R. Conlon, ‘Two New Conditions Supporting the First-Order Approach to Multi-signal Principal-Agent Problems,’ *Econometrica*, January 2009, 77(1), 249-278.

Jacques Crémer and Fahad Khalil, ‘Gathering Information before Signing a Contract,’ *American Economic Review*, June 1992, 82(3), 566-578.

Jacques Crémer, Fahad Khalil, and Jean-Charles Rochet, ‘Strategic Information Gathering before a Contract is Offered,’ *Journal of Economic Theory*, July 1998a, 81(1), 163-200.

Jacques Crémer, Fahad Khalil, and Jean-Charles Rochet, ‘Contracts and Productive Information Gathering,’ *Games and Economic Behavior*, November 1998b, 25(2), 174-193.

James D. Dana Jr., ‘Learning in an Equilibrium Search Model,’ *International Economic Review*, August 1994, 35(3), 745-771.

Mathias Dewatripont , ‘Commitment through Renegotiation-Proof Contracts with Third Parties,’ *Review of Economic Studies*, July 1988, 55(3), 377-389.

Mathias Dewatripont, ‘Renegotiation and Information Revelation over Time: The Case of Optimal Labor Contracts,’ *Quarterly Journal of Economics*, August 1989, 104(3), 589-619.

Mathias Dewatripont and Eric Maskin, ‘Contractual Contingencies and Renegotiation,’ *RAND Journal of Economics*, Winter 1995, 26(4), 704-719.

Peter Diamond and Emmanuel Saez, ‘The Case for a Progressive Tax: From Basic Research to Policy Recommendations,’ *Journal of Economic Perspectives*, Fall 2011, 25(4), 165-190.

Uwe Dulleck and Rudolf Kerschbamer, ‘On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods,’ *Journal of Economic Literature*, March 2006, 44(1), 5-42.

Ronald A. Dye, ‘Costly Contract Contingencies,’ *International Economic Review*, February 1985, 26(1), 233-250.

Ronald A. Dye, ‘Optimal Monitoring Policies in Agencies,’ *RAND Journal of Economics*, Autumn 1986, 17(3), 339-350.

Robert E. Evenson and Yoav Kislev, ‘A Stochastic Model of Applied Research,’ *Journal of Political Economy*, April 1976, 84(2), 265-281.

Joseph Farrell and Robert Gibbons, ‘Cheap Talk with Two Audiences,’ *American Economic Review*, December 1989, 79(5), 1214-1223.

Antoine Faure-Grimaud, Jean-Jacques Laffont, and David Martimort, ‘Collusion, Delegation and Supervision with Soft Information,’ *Review of Economic Studies*, April 2003, 70(2), 253-279.

Chaim Fershtman and Ehud Kalai, ‘Unobserved Delegation,’ *International Economic Review*, 38(4), November 1997, 763-774.

Richard L. Fullerton and R. Preston McAfee, ‘Auctioning Entry into Tournaments,’ *Journal of Political Economy*, June 1999, 107(3), 573-605.

Xavier Gabaix and David Laibson, ‘Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,’ *Quarterly Journal of Economics*, May 2006, 121(2), 505-540.

Esther Gal-Or, ‘Information Sharing in Oligopoly,’ *Econometrica*, March 1985, 53(2), 329-343.

Richard J. Gilbert and David M. G. Newbery, ‘Preemptive Patenting and the Persistence of Monopoly,’ *American Economic Review*, June 1982, 72(3), 514-26.

Richard Gilbert and Carl Shapiro, 'Optimal Patent Length and Breadth,' *RAND Journal of Economics*, Spring 1990, 21(1), 106-112.

Jerry R. Green and Nancy L. Stokey, 'A Comparison of Tournaments and Contracts,' *Journal of Political Economy*, June 1983, 91(3), 349-364.

Sanford J. Grossman and Oliver D. Hart, 'An Analysis of the Principal-Agent Problem,' *Econometrica,* January 1983, 51(1), 7-45.

Theodore Groves, 'The Allocation of Resources under Uncertainty: The Informational and Incentive Roles of Prices and Demands in a Team,' *Technical Report No. 1*, Center for Research in Management Science, University of California, Berkeley, August 1969.

Theodore Groves, 'Incentives in Teams,' *Econometrica*, July 1973, 41(4), 617-631.

Theodore Groves and John O. Ledyard, 'Optimal Allocation of Public Goods: A Solution to the "Free Rider" Problem,' *Econometrica*, May 1977, 45(4), 783-809.

Roger Guesnerie and Jean-Jacques Laffont, 'A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm,' *Journal of Public Economics*, December 1984, 25(3), 329-369.

John Haltiwanger and Joseph E. Harrington Jr., 'The impact of Cyclical Demand Movements on Collusive Behavior,' *RAND Journal of Economics*, Spring 1991, 22(1), 89-106.

Benjamin E. Hermalin and Michael L. Katz, 'Moral Hazard and Verifiability: The Effects of Renegotiation in Agency,' *Econometrica*, November 1991, 59(6), 1735-1753.

Bengt Holmstrom, 'On the Theory of Delegation,' in Marcel Boyer and Richard Kihlstrom (eds.), *Bayesian Models in Economic Theory*. (New York: North-Holland, 1984), 115-141.

Hideshi Itoh 'Incentives to Help in Multi-Agent Situations,' *Econometrica*, May 1991, 59(3), 611-636.

Richard Jensen and Dean Showalter, 'Strategic Debt and Patent Races,' *International Journal of Industrial Organization*, September 2004, 22(7), 887-915.

Ian Jewitt, 'Justifying the First-Order Approach to Principal-Agent Problems,' *Econometrica*, September 1988, 56(5), 1177-1190.

Kenneth L. Judd and Michael H. Riordan, 'Price and Quality in a New Product Monopoly,' *Review of Economic Studies*, October 1994, 61(4), 773-789.

Bruno Jullien, 'Participation Constraints in Adverse Selection Models,' *Journal of Economic Theory*, 93(1), July 2000, 1-47.

Morton I. Kamien and Nancy L. Schwartz, ‘Market Structure and Innovation: A Survey,’ *Journal of Economic Literature*, March 1975, 13(1), 1-37.

Michael L. Katz, ‘Game-Playing Agents: Unobservable Contracts as Precommitments,’ *RAND Journal of Economics*, Autumn 1991, 22(3), 307-328.

Fahad Khalil, ‘Auditing Without Commitment,’ *RAND Journal of Economics*, Winter 1997, 28(4), 629-640.

Paul Klemperer, ‘What Really Matters in Auction Design,’ *Journal of Economic Perspectives*, Winter 2002, 16(1), 169-189.

Paul Klemperer and Margaret Meyer, ‘Price Competition vs. Quantity Competition: The Role of Uncertainty,’ *RAND Journal of Economics*, Winter 1986, 17(4), 618-638.

Levent Koçkesen and Efe A. Ok, ‘Strategic Delegation by Unobservable Incentive Contracts,’ *Review of Economic Studies*, April 2004, 71(2), 397-424.

Frédéric Koessler and David Martimort, ‘Optimal Delegation with Multi-Dimensional Decisions,’ *Journal of Economic Theory*, September 2012, 147(5), 1850-1881.

Vijay Krishna, Giuseppe Lopomo, and Curtis R. Taylor, ‘Stairway to Heaven or Highway to Hell: Liquidity, Sweat Equity, and the Uncertain Path to Ownership,’ *RAND Journal of Economics*, Spring 2013, 44(1), 104-127.

Vijay Krishna and John Morgan, ‘An Analysis of the War of Attrition and the All-Pay Auction,’ *Journal of Economic Theory*, February 1997, 72(2), 343-362.

Vijay Krishna and John Morgan, ‘A Model of Expertise,’ *Quarterly Journal of Economics*, May 2001, 116(2), 747-775.

Jean-Jacques Laffont and David Martimort, ‘Collusion under Asymmetric Information,’ *Econometrica*, July 1997, 65(4), 875-911.

Jean-Jacques Laffont and David Martimort, ‘Collusion and Delegation,’ *RAND Journal of Economics*, Summer 1998, 29(2), 280-305.

Jean-Jacques Laffont and David Martimort, ‘Mechanism Design under Collusion and Correlation’, *Econometrica*, March 2000, 68(2), 309-342.

Jean-Jacques Laffont and Jean Tirole, ‘The Dynamics of Incentive Contracts,’ *Econometrica*, September 1988, 56(5), 1153-1175.

Jean-Jacques Laffont and Jean Tirole, *A Theory of Incentives in Procurement and Regulation*. (MIT Press, 1993).

Mark A. Lemley and Carl Shapiro, ‘Probabilistic Patents,’ *Journal of Economic Perspectives*, Spring 2005, 19(2), 75-98.

David Levhari, Jacob Paroush, and Bezalel Peleg, ‘Efficiency Analysis for Multivariate Distributions,’ *Review of Economic Studies*, January 1975, 42(1), 87-91.

Dan Levin and James L. Smith, ‘Equilibrium in Auctions with Entry,’ *American Economic Review*, June 1994, 84(3), 585-599.

Tracy R. Lewis and David E. M. Sappington, ‘Regulating a Monopolist with Unknown Demand,’ *American Economic Review*, December 1988, 78(5), 986-998.

Tracy R. Lewis and David E. M. Sappington, ‘Information Management in Incentive Problems,’ *Journal of Political Economy*, August 1997, 105(4), 796-821.

Steven A. Lippman and John McCall, ‘The Economics of Job Search: A Survey,’ *Economic Inquiry*, June 1976, 14(2), 155-189.

Glenn C. Loury, ‘Market Structure and Innovation,’ *Quarterly Journal of Economics*, August 1979, 93(3), 395-410.

Richard D. MacMinn, ‘Search and Market Equilibrium,’ *Journal of Political Economy*, April 1980, 88(2), 308-327.

Giovanni Maggi, ‘The Value of Commitment with Imperfect Observability and Private Information,’ *RAND Journal of Economics*, Winter 1999, 30(4), 555-574.

Giovanni Maggi and Andrés Rodriguez-Clare, ‘On Countervailing Incentives,’ *Journal of Economic Theory*, June 1995, 66(1), 238-263.Eric Maskin and John Riley, ‘Monopoly with Incomplete Information,’ *RAND Journal of Economics*, Summer 1984, 15(2), 171-196.

Eric Maskin and Jean Tirole, ‘The Principal-Agent Relationship with an Informed Principal: The Case of Private Values,’ *Econometrica*, March 1990, 58(2), 379-409.

Eric Maskin and Jean Tirole, ‘The Principal-Agent Relationship with an Informed Principal, II: Common Values,’ *Econometrica*, January 1992, 60(1), 1-42.

R. Preston McAfee and John McMillan, ‘Auctions and Bidding,’ *Journal of Economic Literature*, June 1987a, 25(2), 699-738.

R. Preston McAfee and John McMillan, ‘Auctions with a Stochastic Number of Bidders,’ *Journal of Economic Theory*, October 1987b, 43(1), 1-19.

R. Preston McAfee and John McMillan, ‘Auctions with Entry,’ *Economics Letters*, 1987c, 23(4), 343-347.

R. Preston McAfee and John McMillan, 'Optimal Contracts for Teams,' *International Economic Review*, August 1991, 32(3), 561-577.

R. Preston McAfee and John McMillan, 'Bidding rings,' *American Economic Review*, June 1992, 82(3), 579-599.

R. Preston McAfee, John McMillan, and Michael D. Whinston, 'Multiproduct Monopoly, Commodity Bundling, and Correlation of Values,' *Quarterly Journal of Economics*, May 1989, 104(2), 371-383.

Paul Milgrom, 'Auctions and Bidding: A Primer,' *Journal of Economic Perspectives*, Summer 1989, 3(3), 3-22.

James Mirrlees, 'The Theory of Moral Hazard and Unobservable Behavior: Part I,' *Review of Economic Studies*, January 1999, 66(1), 3-21.

Dilip Mookherjee, 'Decentralization, Hierarchies, and Incentives: A Mechanism Design Perspective,' *Journal of Economic Literature*, June 2006, 44(2), 367-390.Dale T. Mortensen, 'Job Search, the Duration of Unemployment, and the Phillips Curve,' *American Economic Review*, December 1970, 60(5), 847-862.

Dale T. Mortensen, 'Property Rights and Efficiency in Mating, Racing, and Related Games,' *American Economic Review*, December 1982, 72(5), 968-979.

Dale T. Mortensen and Christopher A. Pissarides, 'Job Creation and Job Destruction in the Theory of Unemployment,' *Review of Economic Studies*, July 1994, 61(3), 397-415.

John F Muth, 'Rational Expectations and the Theory of Price Movements,' *Econometrica*, July 1961, 29(3), 315-335.

Phillip Nelson, 'Information and Consumer Behavior,' *Journal of Political Economy*, March/ April 1970, 78(2), 311-29.

Richard R. Nelson, 'Uncertainty, Learning, and the Economics of Parallel Research and Development Efforts,' *Review of Economics and Statistics*, November 1961, 43(4), 351-364.

Shmuel Nitzan, 'Modelling Rent-Seeking Contests,' *European Journal of Political Economy*, May 1994, 10(1), 41-60.

Motty Perry and Philip J. Reny, 'On the Failure of the Linkage Principle in Multi-unit Auctions,' *Econometrica*, July 1999, 67(4), 895-900.

Joaquin Poblete and Daniel Spulber, 'The Form of Incentive Contracts: Agency with Moral Hazard, Risk Neutrality, and Limited Liability,' *RAND Journal of Economics*, Summer 2012, 43(2), 215-234.

Michael Raith, 'A General Model of Information Sharing in Oligopoly,' *Journal of Economic Theory*, October 1996, 71(1), 260-288.

Eric Rasmusen, 'Moral Hazard in Risk-Averse Teams,' *RAND Journal of Economics*, Autumn 1987, 18(3), 428-435.

Alfred Rényi, 'On Measures of Entropy and Information,' *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*. (Berkeley: University of California Press, 1961), 547-561.

John G. Riley, 'Silver Signals: Twenty-Five Years of Screening and Signaling,' *Journal of Economic Literature*, June 2001, 39(2), 432-478.Jean-Charles Rochet and Lars A. Stole, 'The Economics of Multidimensional Screening,' in Mathias Dewatripont, Lars P. Hansen, and Stephen J. Turnovsky (eds.), *Advances in Economics and Econometrics*, Volume 1. (Cambridge, UK: The Press Syndicate of the University of Cambridge, 2003), 115-150.

Jean-Charles Rochet and Jean Tirole, 'Platform Competition in Two-Sided Markets,' *Journal of the European Economic Association*, June 2003, 1(4), 990-1029.

Jean-Charles Rochet and Jean Tirole, 'Two-Sided Markets: A Progress Report,' *RAND Journal of Economics*, Autumn 2006, 37(3), 645-667.

Richard Rogerson, Robert Shimer, and Randall Wright, 'Search-Theoretic Models of the Labor Market: A Survey,' *Journal of Economic Literature*, December 2005, 43(4), 959-988.

Julio J. Rotemberg and Garth Saloner, 'A Supergame-Theoretic Model of Price Wars during Booms,' *American Economic Review*, June 1986, 76(3), 390-407.

Alvin E. Roth, 'Two-Sided Matching with Incomplete Information about Others' Preferences,' *Games and Economic Behavior*, June 1989, 1(2), 191-209.

Michael Rothschild, 'Models of Market Organization with Imperfect Information: A Survey,' *Journal of Political Economy*, November-December 1973, 81(6), 1283-1308.

Thomas J. Sargent and Neil Wallace, '"Rational" Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule,' *Journal of Political Economy*, April 1975, 83(2), 241-254.

Jesus K. Seade, 'On the Shape of Optimal Tax Schedules,' *Journal of Public Economics*, April 1977, 7(2), 203-235.

Reinhard Selten, 'The Chain Store Paradox,' *Theory and Decision*, April 1978, 9(2), 127-159.

Sergei Severinov, 'An Efficient Solution to the Informed Principal Problem,' *Journal of Economic Theory*, 141(1), July 2008, 114-133.

Carl Shapiro, 'Exchange of Cost Information in Oligopoly,' *Review of Economic Studies*, July 1986, 53(3), 433-446.

Steven Shavell, 'Risk Sharing and Incentives in the Principal and Agent Relationship,' *Bell Journal of Economics*, Spring 1979, 10(1), 55-73.

Herbert A. Simon, *The Sciences of the Artificial*. (Cambridge, MA: MIT Press, 1981).

Bernard Sinclair-Desgagné, 'The First-Order Approach to Multi-Signal Principal-Agent Problems,' *Econometrica*, March 1994, 62(2), 459-465.

Steven D. Sklivas, 'The Strategic Choice of Managerial Incentives,' *RAND Journal of Economics*, Autumn 1987, 18(3), 452-458.

Kathryn E. Spier, 'Incomplete Contracts and Signalling,' *RAND Journal of Economics*, Autumn 1992, 23(3), 432-443.

Dale O. Stahl, 'Oligopolistic Pricing with Sequential Consumer Search,' *American Economic Review*, September 1989, 79(4), 700-712.

Arun Sundararajan, 'Nonlinear Pricing of Information Goods,' *Management Science*, December 2004, 50(12), 1660-1673.

Dezsö Szalay, 'Contracts with Endogenous Information,' *Games and Economic Behavior*, March 2009, 65(2), 586-625.

Curtis R. Taylor, 'Digging for Golden Carrots: An Analysis of Research Tournaments' *American Economic Review*, September 1995, 85(4), 872-890.

Robert M. Townsend, 'Optimal Contracts and Competitive Markets with Costly State Verification,' *Journal of Economic Theory*, October 1979, 21(2), 265-293.

Eric van Damme and Sjaak Hurkens, 'Games with Imperfectly Observable Commitment,' *Games and Economic Behavior*, October 1997, 21(1-2), 282-308.

Hal R. Varian, 'Buying, Sharing and Renting Information Goods,' *Journal of Industrial Economics*, December 2000, 48(4), 473-488.

Hal R. Varian, 'Position Auctions,' *International Journal of Industrial Organization*, December 2007, 25(6), 1163-1178.

John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Second Edition. (Princeton: Princeton University Press, 1947).Michael Waldman, 'Job Assignments, Signalling, and Efficiency,' *RAND Journal of Economics*, Summer 1984, 15(2), 255-267.

Michael Waldman, 'Classic Promotion Tournaments versus Market-Based Tournaments,' *International Journal of Industrial Organization*, May 2013, 31(3), 198-210.

Joel Watson, 'Alternating-Offer Bargaining with Two-Sided Incomplete Information,' *Review of Economic Studies*, July 1998, 65(3), 573-594.

Oliver Williamson, *Markets and Hierarchies: Analysis and Antitrust Implications*. (New York: Free Press, 1975).

Asher Wolinsky, 'Prices as Signals of Product Quality,' *Review of Economic Studies*, October 1983, 50(4), 647-58