

UNIVERSITY OF CALIFORNIA, BERKELEY

Department of Economics

Berkeley, California 94720

Working Paper 8742

ADVERSE SELECTION WITHOUT HIDDEN INFORMATION

Paul Milgrom

June 2, 1987

Key words: Adverse selection, hidden information, free rider problem.

Abstract

The usual distinction between "hidden action" and "hidden information" problems is criticized. The criticism is founded partly on purely logical grounds, partly on a review showing how the received literature can be better understood when the usual distinction is suppressed, and partly on the development of new static and dynamic models of adverse selection without hidden information.

ACKNOWLEDGMENT

Sections 1 and 2 of this paper are based on teaching notes from a graduate course entitled Advanced Microeconomic Theory taught at Yale University in 1984-5. I am grateful to John Roberts and to my students for their comments and suggestions. I am also pleased to acknowledge the financial support of the National Science Foundation and the Guggenheim Foundation.

Paul Milgrom

Yale University

University of California, Berkeley

ADVERSE SELECTION WITHOUT HIDDEN INFORMATION

1. Introduction

Consider the problem of an asset owner who rents his asset to another party and wishes to motivate the renter to care for it properly. One possibility is that the efforts supplied by the renter make deterioration of the asset less likely. Another is that the owner incurs lower maintenance costs if the renter reports, say, a leaky valve when he first notices the leak, because then any needed repair can be accomplished before additional damage occurs. It has been customary in the literature of economic theory to distinguish the former possibility, in which the problem is to motivate the renter to take an appropriate hidden action (his effort), from the latter, in which the problem is to motivate him to take an appropriate observable action (his report) on the basis of hidden information.¹ Different methods have usually been applied to the analysis of these two kinds of models. Our purpose is to argue that this

¹Arrow [1985], Fudenberg and Tirole [1986], Grossman and Hart [1983], and Myerson [1985] are among those who make this distinction.

customary distinction is logically unsound and that the received literatures on principal-agent problems, free-rider problems, signalling, adverse selection, and dynamic incentive contracts can all be better understood when this distinction is rejected.

The logical flaw in treating hidden information and hidden action as distinct problems is an elementary one: When an individual makes a report or takes any other observed choice on the basis of hidden information, the strategy that he uses (the function from his information to his choice) is unobserved and therefore constitutes a hidden action. Hidden information is a particular case of hidden action, so any general proposition that applies to hidden action models can be specialized for application to hidden information problems.

This logical criticism is not by itself very damaging, for two reasons. First, some hidden action models employ assumptions that make it impossible to specialize the model to obtain a hidden information case. For example, it is often assumed that the hidden action is represented by a real number -- the "hidden effort" -- and that output varies monotonically with effort. Hidden information models in which an observable effort is taken as a function of unobservable information cannot generally be reduced to the choice of a one dimensional parameter while preserving the

required monotonicity. Second, because hidden information is a special case of hidden action, the logical criticism leaves open the possibility that the principal conclusions obtained in hidden information problems actually do require the hidden information assumption.

The remainder of this paper is devoted to establishing that the principal conclusions in a wide range of important hidden information and hidden effort models are essentially identical. We proceed using a series of examples. The first example investigates the possibilities and pitfalls in extending a hidden action principal-agent model to incorporate the hidden information case. The other examples consist primarily of pairs of models, one of the hidden effort sort and one of the hidden information sort, for which the results can be compared to see in what ways they correspond. The two pairs studied in this Section are models of the free rider problem and models of signalling to deter entry or expansion by a competing producer. For these two pairs, we are able to draw models from the existing literature. Sections 2 and 3 study situations where individuals with private information decide whether to enter into, renew, or terminate a contract or trading relationship (that is, adverse selection problems) in which a hidden one-dimensional effort choice replaces the usual exogenously specified hidden information. Section 2 presents our analysis of

a "lemons" problem resembling that studied by Akerlof [1970], but without hidden information. Section 3 examines the role of adverse selection at contract renewal dates in a multiperiod contracting problem.

We begin with the particular hidden action principal-agent model analyzed by Grossman and Hart [1983]. Their formal mathematical analysis includes two main cases: one in which the agent's cardinal utility function over income (I) and action (a) takes the additively separable form $u(I) - v(a)$ and one in which it takes the multiplicatively separable form $u(I)v(a)$. The agent's action affects the distribution of some observed outcome x ; the principal's profits depend on x ; and the employment contract specifies the agent's compensation as a function of x .

Now consider the corresponding hidden information problem in which the agent observes privately an exogenous random variable θ and then makes a publicly observed choice a .² Let α designate the agent's choice strategy, so that $a = \alpha(\theta)$. Suppose that the principal's profits and the agent's compensation depend on the choice a . Further suppose that the agent's preferences take the additively separable form $u(I) - w(a, \theta)$ and define $v(\alpha) = E[w(\alpha(\theta), \theta)]$.

²Sappington [1983] presents a related hidden information model, in which the agent's limited financial resources rather than his risk aversion motivate the principal's participation in the enterprise.

For any random variable I , we have:

$$(1) \quad E[u(I) - w(\alpha(\theta), \theta)] = E[u(I) - v(\alpha)] .$$

Consequently, if we identify a as the outcome and α as the hidden action, this hidden information model is seen to be a special case of the Grossman-Hart hidden action model, so their various conclusions apply to it. No similar construction is possible for the multiplicatively separable case, making it impossible to apply the multiplicatively separable version of the model to hidden information problems.

Consider next the "public goods free rider problem" as studied by Green and Laffont [1979]. This problem is one of hidden information: Individuals have private information about their personal valuation functions $u_i(\cdot)$ of a public good g . Preferences are represented as $u_i(g) + t_i$, where t_i is the net transfer received or, if negative, the net taxes paid. Let $C(g)$ be the cost of providing the good. Suppose that each consumer is asked to report a valuation function \hat{u}_i and that g is then chosen to maximize the reported social surplus $\sum \hat{u}_i(g) - C(g)$. Then there exists a tax and transfer scheme (the Vickrey-Groves-Clarke scheme) that makes it a dominant strategy for consumers to report their preferences honestly, that is, to set \hat{u}_i equal to u_i . If the scheme is used and consumers play their dominant strategies, then the public decision g will maximize actual social surplus. However, this tax and transfer scheme

cannot be constructed to ensure that the budget balances: that is, $\sum t_i + C(g)$ cannot generally be set equal to zero. A feasible tax policy that does not throw away money is a requirement for first-best efficiency, so the Vickrey-Groves-Clarke scheme, despite ensuring the efficient choice of g , does not lead to fully efficient outcomes. D'Aspremont and Gerard-Varet [1979] have exhibited conditions under which both of the desired conditions -- a balanced budget and an efficient choice of g -- are simultaneously possible. However, their result requires that the notion of equilibrium be weakened from dominant strategies to Bayes Nash equilibrium.

Let us compare this hidden information free rider model and analysis to Holmstrom's [1982] treatment of a hidden action free rider problem. Suppose the total output of a joint venture or partnership takes the form $\sum e_i + \eta$, where e_i is the effort level of partner i and η is a random "noise" term. Let the partner i 's utility for effort and net transfers be $u_i(e_i) + t_i$, where each u_i is smooth, $u_i' < 0$, $u_i'' < 0$, and $u_i'(0) > -1$ (the last assumption assures that a positive level of effort is efficient). Now, if the total output x is to be shared among the partners according to smooth sharing rules $s_i(x)$ with $\sum s_i(x) = x$ then, as Holmstrom showed, the first-best effort levels ($u_i'(e_i^*) = -1$) cannot constitute a Nash equilibrium: At any equilibrium, some partner will find it optimal to contribute less than

the first-best level of effort because the private marginal returns from his efforts are less than the marginal returns to the partnership. However, if the "balanced budget" requirement that total payments must add up to total output is dropped, then the compensation scheme $s_i(x) = x + k_i$ (where the k_i 's are arbitrary constants that sum to zero) makes it a dominant strategy for each agent to undertake the optimal level of effort, since each receives the full marginal benefit of his contributions. The same scheme makes it a dominant strategy for partners always to contribute first best effort levels in a variant of Holmstrom's model where the u_i functions are unknown but the efforts e_i are observed. This variant is also a special case of the hidden information model reviewed in the preceding paragraph, where the "public decision" g is the list of the partners' effort levels.

Is there a general model that encompasses the existing hidden information and unobserved effort free-rider models? What is the analogue to the results of D'Aspremont and Gerard-Varet for the unobserved effort and general cases? The artificial distinctions that have been maintained between hidden information and hidden action models has kept questions such as these off the research agenda.

Our final example in this Section juxtaposes two models of firms that signal to limit the entry or expansion of a

competitor. The first is a model of Milgrom and Roberts [1982], in which an incumbent monopolist facing potential competition has hidden information about its production costs. Milgrom and Roberts find that the incumbent often produces a quantity in excess of the full information monopoly quantity in order to signal to the potential entrant that its costs are low and so to reduce the probability of entry. The second is a model of Riordan [1985], in which Cournot oligopolists that compete over two periods in a market with uncertain demand make unobserved quantity choices. There is no hidden information in this set-up, but both firms do observe the price that prevails in the first-period and use it to draw inferences about the underlying demand. Riordan finds that each firm produces more than the full information Cournot quantity in the first period in order to reduce its competitors' estimate of second period demand, hoping to curtail its second period output. In both models, firms increase output because they expect that action to reduce a competitor's next period output. Both models should be classified as signalling models, because both lead to the standard signalling conclusion ("overinvestment" in the signal) for the standard reason (to influence another actor's choices by manipulating his beliefs). Other signalling models without hidden informa-

tion have been developed by Holmstrom [1982a] and Fudenberg and Tirole [1986].

2. A Market for "Lemons"

Consider a situation in which there are many (in fact, a continuum) of used car owners, whose payoff from owning a car of quality q is q , and equally many (potential) buyers whose payoff from buying a car of quality q at price p is $B(q) - p$, where $B' > 0$. A buyer who does not buy gets a payoff of zero; an owner who sells for a price p gets payoff p . We assume that $B(q) > q$ for all q , so that there are potential gains from trade. Each owner chooses the quality q of his car by the level of maintenance expense he incurs. We assume that the cost of maintaining a car to quality $q \geq 0$ is $C(q) = q^2/2$. Thus, an owner who intends to keep his car will solve $\text{Max}_q q - q^2/2$, setting $q = 1$ and receiving a payoff of $1/2$. Since this option is always available to an owner, the owners' expected payoffs in every equilibrium model we study will always be at least $1/2$.

There are several ways to analyze the likely outcome in the market. The first, following along the lines of Akerlof's [1970] analysis, defines an equilibrium price as one at which the market clears when the potential buyers correctly anticipate the quality of the vehicles being offered. If $B(0) \geq 1/2$, then there is an equilibrium in

which $p = B(0)$, the quality of all cars are set to 0, and all cars are traded. This equilibrium is inefficient, because the level of maintenance is below the first-best level (because $B'(0) > C'(0) = 0$), but the market does clear with all cars traded. As in Akerlof's original analysis, only low quality cars are traded at equilibrium.

If $B(0) < 1/2$, there is no Akerlof equilibrium with positive trading volume. For, if there were, then the owners who offer their cars would all set $q = 0$, and then competition among the buyers would lead to $p = B(0) < 1/2$. However, as we have seen, a seller who sets $q = 1$ and keeps the car gets a payoff of $1/2$, so no cars would be offered for sale at the price p . One might wish to argue that there is an equilibrium with $p = B(0)$ and all owners keeping their cars, since the "average" quality of the cars offered can be arbitrarily set to zero in this case, but this path leads to difficult questions of interpretation about what the price means and where it comes from.

A more fruitful approach is to specify a mechanism by which prices are formed and trade occurs and to analyze it using the methods of game theory. We first suppose that prices are quoted by the car owners. Since each owner knows the quality of his car, it is possible that the price-quote may signal something about the quality of the car being sold. After hearing the prices, the buyers decide whether

to try to buy and, if so, at which one of the offered prices. If there is an excess supply of cars at any named price, the sellers at that price are assumed to be rationed (each with equal probability). If there is excess demand, the buyers at that price are rationed. The rules of the game allow the possibility that trades may take place at several prices and that there may be rationing of buyers at some prices and of sellers at others. Whether these things can happen at equilibrium is a matter to be established by analysis of the model.

Let $s(p)$ be the probability that a seller is rationed at price p at equilibrium and let $b(p)$ be the probability that a buyer is so rationed. Repeating the arguments made above for the non-existence of an Akerlof equilibrium shows that, when $B(0) < 1/2$, there is no Nash equilibrium with a positive volume of trade and no rationing. However, there are many equilibria with rationing. Indeed, for any quality $q \leq 1$ with

$$(2) \quad q^2/2 + (1-q)B(q) \geq 1/2$$

and provided $B(0) < 1/2$, there is an equilibrium in which all owners set their qualities at q , ask the price $p = B(q)$, and are rationed with probability $s(q) = q$. The buyers' strategies call for them to purchase with certainty anything priced at or below $B(0)$, to randomize so that a car offered at price $p = B(q)$ is sold with probability $1-s(q)$, and never

to purchase a car with a price above $B(0)$ and different from $B(q)$. This is a "signalling" equilibrium (but without hidden information) in which buyers "believe" that prices other than $B(q)$ signal quality zero, and act accordingly.

VERIFICATION: Let us check that this is, in fact, an equilibrium. An owner faced with the others' strategies has three options. He can set a price at which nobody will buy, set quality accordingly, and earn an expected payoff of $1/2$. Or, he can set a price at which he is not rationed (the highest such price is $B(0)$), set quality to zero, and earn at most $B(0) < 1/2$. Finally, he can set a price $p = B(q)$, plan to sell if he is not rationed, and choose a quality Q privately to solve:

$$(3) \quad \begin{aligned} &\text{Maximize } s(p)Q + [1 - s(p)]p - Q^2/2 . \\ &Q \geq 0 \end{aligned}$$

The unique maximum of this strictly concave problem is achieved by $Q = s(p) = q$. In view of (2), the resulting payoff is at least $1/2$, so this is the best of the owner's options. Given the owners' strategies, the buyers have no strategies yielding positive payoffs, so the specified zero payoff strategy is a best response for them. This verifies that the specified strategies form an equilibrium.

For some specifications of B , there exist equilibria at which several different qualities are supplied at equilibrium and with a different price charged for each level of

quality. We illustrate this possibility with an extreme example in which a whole interval of prices prevail simultaneously in the market at equilibrium. Let $B(q) = (1+q)/2$. Then, we claim, for any distribution of quality levels on the interval $[0,1]$, there is an equilibrium in which the owners randomize their choices of q according to that distribution and set corresponding prices $p = B(q)$. Buyers randomize so that at equilibrium a seller's probability of being rationed at price p is $s(p) = B^{-1}(p) = q$.

VERIFICATION: To verify that this is a Nash equilibrium, note from (3) that for any price choice p , the owner's corresponding best quality choice is $s(p)$. For each quality level, a calculation shows that owner's profits are precisely $1/2$. Since an owner earns the same maximum expected payoff from each price and also from planning to keep his car, his strategy of randomizing among these alternatives is a best response for him to the others' strategies. A buyer has no strategy that earns a positive profit, so the buying strategy that generates the specified rationing probabilities $s(p)$ is a best response to the others' strategies. The equilibrium is thus verified.

When an owner decides to offer a car for sale at some price p , he determines the quality by solving (3). The resulting quality is always less than unity - the quality that would be maintained if there were no possibility of

trade. Since this holds regardless of $B(q)$, it is evident that there can be an arbitrary undersupply of quality. If more than one quality is offered, then one sees from (3) that the highest quality goods at the separating equilibrium are those least likely to be traded, just as in Akerlof's analysis. Because actual quality is degraded by the possibility of trade, the gains from trade at equilibrium quality levels is a poor indicator of the welfare benefits of trade. Indeed, in the separating equilibrium described above, all the cars are offered for sale and the total surplus is positive, but there is no ex ante gain to anyone from the institution of trade: At equilibrium, the expected payoff to each buyer is zero and that to each seller is $1/2$, just as they would be if no trade were possible.

Next, consider the game in which the prices are named by the buyers, who have no private information to signal. For technical reasons, we shall suppose that any buyer's offer of an unexpected price is directed randomly to a single owner. The resulting model resembles that analyzed by Rothschild and Stiglitz [1976], and its analysis leads to similar conclusions. One of the hallmarks that distinguished the analysis of Rothschild and Stiglitz from that of Akerlof was the possibility that there might be no equilibrium. We verify below that if $B(0) < 1/2$, then our corres-

ponding hidden action game has no sequential equilibrium (as defined by Kreps and Wilson [1982]).

VERIFICATION: There is no equilibrium at which no trade occurs: If there were, all owners would set $q = 1$ and some buyer would then find it profitable to offer a price between 1 and $B(1)$ (because the offeree would gladly accept the offer). Suppose there were an equilibrium at which some trade does occur and let \underline{q} be the lowest quality level at which trade occurs at that equilibrium. Then \underline{q} solves (3) for some \underline{p} , and is indeed the only value of Q that solves (3) (because the maximand is strictly concave). The Bertrand form of competition among the buyers ensures that the buyers earn zero profits at each quality level, so that $\underline{p} = B(\underline{q})$. Since $\underline{q} < 1$ and the owner's profit is at least $1/2$, it follows that $\underline{p} > 1/2$ and hence that $\underline{q} > 0$ (because $B(0) < 1/2$) and so that the owner is rationed at price \underline{p} . For ϵ small, an offer of $\underline{p} - \epsilon$ would be accepted with positive probability by the offeree, since the offeree may be rationed at a price of \underline{p} . Hence, such an offer earns a positive expected profit for the buyer, while the purported equilibrium strategy earns a zero profit -- a contradiction. Thus, there is no equilibrium with trade and none without, so the claim is verified.

3. Multiperiod Incentive Models

Some recent research in information and incentives has emphasized the problem of incentives in long-term relationships when long-term commitments cannot be made. For example, Fudenberg, Holmstrom, and Milgrom [1987] have investigated the extent of the necessary inefficiencies caused by contract renegotiations in a long-term relationship. They find that if, at the dates of renegotiation or renewal of a contract (when each party has an option to terminate the relationship), the parties' preferences over future contracts are common knowledge, then the possibility of renegotiation entails no necessary inefficiencies. Indeed, a sequence of short-term contracts then performs as well as a comprehensive long-term contract. The common knowledge condition in the analysis of Fudenberg, Holmstrom and Milgrom implies that, at each recontracting date, there can be no problem of adverse selection, whether due to hidden past efforts or hidden information or some mixture of the two. Then, the recontracting process at each renewal date is efficient, so there is no loss from renegotiation.³

When the agent in a principal-agent relationship takes a hidden action that affects his preferences among future contracts, the adverse selection problem at recontracting

³The significance of this conclusion for economic organization theory has been elaborated by Milgrom and Roberts [1987].

dates can be a source of inefficiency. One such example has been given by Fudenberg, Holmstrom, and Milgrom [1987]. The one presented below is a variant of a hidden information model introduced by Laffont and Tirole [1985], and serves to establish once again the similarity between the implications of models with and without hidden information.

The model of Laffont and Tirole represents a situation in which a principal contracts with an agent in two successive periods. The principal's superior bargaining power is modeled by the assumption that he can offer contracts to the agent on a take-it-or-leave-it basis in each period. However, the principal is unable to commit himself to a long term contract.

In the model of Laffont and Tirole, the agent has a cost function $C(\cdot; \theta)$ and knows the value of the (real) parameter θ , which however is unknown to the principal. It is assumed that the partial derivative $C_\theta < 0$ and that θ is distributed according to some distribution function F with positive density f . Without loss of generality, F can be taken to be the uniform distribution on $(0,1)$.

The initial contract requires that the agent make some report $\hat{\theta}$ about his cost function. As a function of the report, the principal prescribes some action $a = \alpha(\hat{\theta})$ and pays some compensation $s(\hat{\theta})$. The agent's profit is $\pi^1 = s(\hat{\theta}) - C(\alpha(\hat{\theta}); \theta)$ if he accepts the contract; he can alterna-

tively quit and receive a profit of zero. The pair (α, s) is the first period contract. In the second period, the principal, knowing $\hat{\theta}$, asks for another report $\bar{\theta}$, and specifies another action $b = \beta(\bar{\theta}, \hat{\theta})$ and compensation $t(\bar{\theta}, \hat{\theta})$. The pair $(\beta(\cdot, \hat{\theta}), t(\cdot, \hat{\theta}))$ is the second period contract. The agent's second period profit is $\pi^2 = t(\bar{\theta}, \hat{\theta}) - C(\beta(\bar{\theta}, \hat{\theta}), \theta)$ if he accepts both the first and second period contract; otherwise it is zero. We may take the principal's payoff if the agent works to be $\alpha(\hat{\theta}) - s(\hat{\theta})$ in the first period and $\beta(\bar{\theta}, \hat{\theta}) - t(\bar{\theta}, \hat{\theta})$ in the second; the principal receives zero for any period in which the agent does not work.

The model is analyzed as a two stage principal-agent game, in which the principal moves first at each stage by offering a contract, which the agent accepts or rejects. The agent then makes a report and an action is determined as the contract requires. As is customary in principal-agent games, it is assumed that if the agent is indifferent among moves at some point in the game tree, he chooses a move that is most preferred by the principal. A major result of the Laffont-Tirole analysis is that there is no sequential equilibrium in which the agent always accepts the contracts offered and $\hat{\theta} \equiv \theta$, that is, it is never the case that the agent has an incentive to work and report all of his information truthfully. Consequently, at equilibrium, the principal does not have precise beliefs about the agent's

productive capacity at the beginning of period two: By this we mean that the principal's beliefs about θ at the beginning of the second period given the report $\hat{\theta}$ are not represented by a probability distribution that concentrates all its mass on a single point.

We shall show below that, in a hidden effort version of the Laffont-Tirole model in which (at equilibrium) the agent chooses $\theta > 0$ and works in both periods, the principal's beliefs at the start of the second period can never be precise. In particular, there is no pure strategy equilibrium of the model, since any pure choice of θ would lead the principal to have precise beliefs. Models in which θ is chosen constitute a significant extension of the original model; for example, they capture the case where θ represents the investment an agent makes in capital that is specialized for the current principal.

Thus, suppose the agent chooses θ at cost $K(\theta)$ after the initial contract is signed. Assume that $K(0) = 0$ and $K'(\theta) > 0$ for all θ . We shall argue by contradiction.

Assume that there is an equilibrium at which, following some contract offer of the principal, the agent either chooses some level of θ with certainty or plays a mixed strategy and sometimes chooses θ and reports it so that the principal has precise beliefs. At equilibrium, the agent's strategy is a best response to the principal's; in parti-

cular, the agent cannot benefit by keeping his reports and contract acceptance decisions unchanged and raising θ slightly. Thus, writing the agent's total expected profits from the two periods following the initial contract offer as $-K(\theta) + \pi^1(\theta, \hat{\theta}) + E[\pi^2(\theta, \hat{\theta}, \bar{\theta})]$, we must have:

$$(4) \quad \pi^1_{\theta}(\theta, \hat{\theta}) + E[\pi^2_{\theta}(\theta, \hat{\theta}, \bar{\theta})] - K'(\theta) \leq 0,$$

where the expectation represents the possibility that the principal will randomize over β . Since $\pi^2_{\theta} = -C_{\theta}(\beta(\hat{\theta}, \bar{\theta}), \theta) > 0$, it follows that $\pi^1_{\theta}(\theta, \hat{\theta}) - K'(\theta) < 0$.

Now, as Laffont and Tirole have shown, $\pi^2 = 0$ at equilibrium when the principal has precise beliefs at the start of the second period. (Otherwise, the principal could benefit by deviating at the start of the second period, calling for the same action but reducing the agent's compensation.) Consequently, another best response for the agent is to choose θ and $\hat{\theta}$ as he would at equilibrium but then to quit at the beginning of the second period. In particular, the agent cannot do better by reducing θ while still reporting $\hat{\theta}$ and working just one period: $\pi^1_{\theta}(\theta, \hat{\theta}) - K'(\theta) \geq 0$. But that is inconsistent with the conclusion of the last paragraph, establishing the required contradiction.

Laffont and Tirole also showed that the sequence of two contracts is strictly worse for the principal than a two period contract would be, if commitment to such a contract were possible. That conclusion applies here as well.

Indeed, the best two period contract requires that the agent always choose the (generically unique) θ^* that solves

$$\text{Max}_{\theta} \text{Max}_{a,b} -K(\theta) + a - C(a;\theta) + b - C(b;\theta).$$

Our preceding analysis establishes that no such pure choice can be part of an equilibrium when the principal must offer a sequence of short term contracts.

4. Concluding Remarks

We have established that many of the propositions derived in single period and multiperiod hidden information models apply equally when there is a hidden effort choice, and reversely. In terms of our opening example of an asset owner and a renter, it doesn't matter much whether the incentive problem is one of motivating the renter to supply effort or timely and honest reporting. That fact is comforting when one seeks to apply the principles of incentive theory to complex situations, and it is helpful when one tries to understand the relations among the diverse models that populate the modern incentive theory literature.

REFERENCES

- Akerlof, George [1970], "The Market for Lemons: 'Quality Uncertainty and the Market Mechanism,'" Quarterly Journal of Economics 84: 488-500.
- Arrow, Kenneth [1985], "The Economics of Agency" in Principals and Agents: The Structure of Business, John Pratt and Richard Zeckhauser (eds), Boston: Harvard Business School Press.
- D'Aspremont, Claude and Louis Gerard-Varet [1979], "Incentives and Incomplete Information," Journal of Public Economics, 11: 25-45.
- Fudenberg, Drew, Bengt Holmstrom and Paul Milgrom [1987], "Short-Term Contracts and Long-Term Agency Relationships," unpublished.
- and Jean Tirole [1986], "A Signal-Jamming Theory of Predation," Rand Journal of Economics, 17: 366-76.
- Grossman, Sanford and Oliver Hart [1983], "An Analysis of the Principal-Agent Problem," Econometrica 51: 7-45.
- Green, Jerry and Jean-Jacques Laffont [1979], Incentives in Public Decision Making, Amsterdam: North-Holland.
- Holmstrom, Bengt [1982], "Moral Hazard in Teams," Bell Journal of Economics, 13: 324-40.
- [1982a], "Managerial Incentive Problems: A Dynamic Perspective," in Essays in Economics and Management in Honor of Lars Wahlbeck, Helsinki: Swedish School of Economics.
- Kreps, David and Robert Wilson [1982], "Sequential Equilibria," Econometrica, 50: 863-94.
- Laffont, Jean-Jacques and Jean Tirole [1985], "The Dynamics of Incentive Contracts," unpublished working paper.
- Milgrom, Paul and John Roberts [1982], "Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis," Econometrica, 50: 443-59.

- [1987], "Bargaining and Influence Costs and the Organization of Economic Activity," Stanford GSB Working Paper.
- Myerson, Roger [1985], "Bayesian Equilibrium and Incentive Compatibility: An Introduction," in Social Goals and Social Organization, Leonid Hurwicz, David Schmeidler and Hugo Sonnenschein (eds), Cambridge University Press, 229-259.
- Riordan, Michael [1985], "Imperfect Information and Dynamic Conjectural Variations," Rand Journal of Economics, 16: 41-50.
- Rothschild, Michael and Joseph Stiglitz [1976], "Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Imperfect Information," Quarterly Journal of Economics, 90: 629-50.
- Sappington, David [1983], "Limited Liability Contracts Between Principal and Agent," Journal of Economic Theory, 29: 1-21.