

# Relying on the information of interested parties

Paul Milgrom\*

and

John Roberts\*\*

*We investigate the conventional wisdom that competition among interested parties attempting to influence a decisionmaker by providing verifiable information elicits all relevant information. We find that, if the decisionmaker is strategically sophisticated and well informed about the relevant variables and about the preferences of the interested party or parties, competition may be unnecessary to achieve this result. If the decisionmaker is unsophisticated or not well informed, competition is not generally sufficient. If the interested parties' interests are sufficiently opposed, however, or if the decisionmaker is seeking to advance the parties' welfare, then competition can reduce or even eliminate the decisionmaker's need for prior knowledge about the relevant variables and for strategic sophistication. In other settings only the combination of competition among information providers and a sophisticated skepticism is sufficient to allow effective decisionmaking.*

“. . . [T]he only way in which a human can make some approach to knowing the whole of a subject is by hearing what can be said about it by persons of every variety of opinion, and studying all modes in which it can be looked at by every character of mind.”

—John Stuart Mill, *On Liberty*

“So long as dissent is not suppressed, there will always be some who will query the ideas ruling their contemporaries and put new ideas to the test of argument and propaganda.”

—Friedrich A. Hayek, *The Road to Serfdom*

“What we usually call ‘scientific knowledge’ is, as a rule . . . information regarding the various competing hypotheses and the way they have stood to various tests.”

—Karl R. Popper, *The Open Society and Its Enemies*

---

\* Yale University.

\*\* Stanford University.

This research was supported by National Science Foundation Grants IST-8411595 and IST-8420606 to Yale University and SES-8308723 to Stanford University. Part of this work was completed while Milgrom was a Fellow of the Institute for Advanced Studies at the Hebrew University of Jerusalem.

We thank the participants at the Stanford Summer Accounting Workshop, the Northwestern University Summer Workshop on Strategic Behavior and Competition, the Cowles Foundation Seminar, the Theory Seminar at Princeton, the Stanford University Conference on Adaptive Institutions, and the Conference on Games and the Theory of Firm Behavior at the University of Western Ontario for their suggestions and comments. Bernard Desgagne provided diligent and able research assistance. Stan Besen and the anonymous referees helped us to clarify the exposition.

## 1. Introduction

■ A common problem faced by decisionmakers is the need to rely on suggestions and information provided by individuals who are affected by their decisions. Although interested parties may try to manipulate the decisionmaker's choice by concealing or distorting information, their efforts do not always succeed. An archetypical example of an institution that is designed to prevent concealment and distortion is the adversary system used to resolve legal disputes. The chief perceived virtue of an adversary system is that, since any relevant piece of information favors one disputant or the other, one can rely on the disputants themselves to report all relevant information. More generally, it has been argued that "free and open discussion" or "competition in the marketplace of ideas" will result in the truth's becoming known and appropriate decisions' being made in a variety of political, scientific, legal, regulatory, and market contexts.

In this article we examine the validity and scope of this argument by studying the problem of a decisionmaker who must rely on one or more interested parties to provide information about possible decisions and their consequences. We identify conditions under which a decisionmaker can reach a good decision despite possibly severe limitations both on the decisionmaker's prior information and ability to draw sophisticated inferences and on the capacity of the interested parties to communicate what they know.

There are two general kinds of strategies that may be available to a decisionmaker to overcome some of these limitations. First, although lacking information about the specific situation, a decisionmaker may nevertheless be sophisticated about interpreting any information reported to him by recognizing that self-interest tinges the reports. We shall find that a sophisticated *skepticism* can be an important weapon in the decisionmaker's arsenal. Second, a decisionmaker—even one who is too unsophisticated to implement an effective skeptical strategy—may still be able to extract useful information by inducing well-informed parties with competing interests to compete in providing information. Sometimes, a combination of these two techniques is required for optimal decisionmaking.

For our formal study of the decisionmaker's problem, we introduce a class of "persuasion games." In these games the decisionmaker and the interested party or parties interact only once, so that issues of reputations do not arise.<sup>1</sup> We assume that the interested parties can withhold information, but that the decisionmaker can freely verify anything that is reported to him. For example, a seller might verifiably report that his product meets or exceeds a standard that is out-of-date or just barely met, but at the same time he might simply fail to mention a more relevant or stringent standard that the product does not meet.

The assumption that reported information can be freely verified is important for our analysis, although one could substitute the assumption that there are penalties for perjury, false advertising, or warranty violations that are sufficiently sure and heavy that false reporting never pays. When information is not verifiable, the reliability of any report depends in part on the degree of consonance between the objectives of the decisionmaker and those of the interested party or parties. The case with one interested party is equivalent to a problem of delegation, and has been studied by Holmström (1977) and Crawford and Sobel (1982), among others.

We begin our formal analysis in Section 2 by studying the problem of a decisionmaker who relies on a single interested party to report information. Initially we investigate how a sophisticated decisionmaker can use skepticism as a weapon to extract maximum information. Our first results here extend ones previously reported by Milgrom (1981) and Grossman (1981): If the interested party has known monotone preferences over the decisionmaker's choice set (e.g., a seller wants to sell as much as possible, an electric utility company prefers

---

<sup>1</sup> See Sobel (1985) for an analysis of reputation building in information provision.

less restrictive emissions standards) and has information that bears on the decisionmaker's preferences, and if the decisionmaker knows what information to seek, then (i) the decisionmaker's unique equilibrium strategy is to *assume the worst*, that is, to make the inference that, consistent with his information, leads to the least favorable decision for the interested party and (ii) the equilibrium decision is the *full-information decision*—the decision that would have been reached if the decisionmaker had perfect access to the interested party's information.

Often an interested party will be uncertain about whether the decisionmaker is *sophisticated*, that is, capable of game-theoretic reasoning. There are many kinds of unsophisticated behavior. For example, an unsophisticated buyer may be naively credulous, that is, he may interpret any information reported to him as if it were the complete report of some disinterested observer<sup>2</sup> and act accordingly. When the interested party is unsure whether the decisionmaker is sophisticated or naively credulous, his equilibrium reporting strategy is the strategy that elicits the most favorable decision from a credulous decisionmaker.<sup>3</sup> For example, a rational salesman will treat every buyer as if he were naively credulous. This treatment, however, does not harm the sophisticated decisionmaker: at equilibrium, he uses the skeptical "assume-the-worst" strategy, interprets the salesman's report correctly, and reaches the full-information decision. Similarly, the presence in the market of sophisticated buyers does not benefit an unsophisticated buyer; he hears the same report and makes the same decision as he would if the seller knew him to be naively credulous.

It is interesting to contrast these results with the results of search-theoretic models in which consumers search for the lowest price. There, the presence of knowledgeable consumers, or ones with low search costs, benefits less knowledgeable consumers by shifting downward the distribution of prices (Wilde, 1977). Similarly, consumers with high search costs harm other consumers by shifting the distribution of prices upwards. In our model there is no such effect, because the seller finds that the same strategy is effective against both sophisticated and credulous consumers.

The assumptions needed to justify the conclusions of Section 2 are many. First, one needs to assume that the interested party's preferences are known. For example, this assumption would be satisfied if the interested party is a seller who is offering a single product and wants to sell as much as possible, or if he is offering one unit of any of several products and wants to sell the most expensive one. On the other hand, if mark-ups vary across several products in a manner unknown to the buyer, then the result does not apply—the seller may benefit by withholding information. Similarly, if the decision involves, say, setting product safety standards, the interested parties' preferences are likely to be unknown and complex, and the result again does not apply.

Second, the decisionmaker must know the factors about which the interested party has information to detect situations in which information is being withheld. For example, if a used car salesman has information about recent repairs to a car but does not report it, the buyer may not know that information has been withheld.

Third, the decisionmaker must be sophisticated enough to draw the appropriate inference when information is withheld. When any of these first three conditions fails, the decisionmaker will be unable to implement the strategy of extreme skepticism, and will suffer a loss of utility as a result.

Finally, for the full information decision to be an appropriate welfare standard, the decisionmaker must be able to draw the proper inferences and reach the right decision when all relevant information has been made available to him.

---

<sup>2</sup> The way a naively credulous decisionmaker forms his beliefs is made precise in Section 2.

<sup>3</sup> A similar result is shown to hold for a wide class of unsophisticated behaviors, provided that the decisionmaker responds positively to "favorable information." The result that the interested party treats a decisionmaker as if he were unsophisticated also holds for more complex environments, in which various kinds of unsophisticated behavior are possible.

In Section 3 we begin to study how competition among interested parties in providing information may substitute for the many restrictions listed above, so that an unsophisticated decisionmaker, with little or no idea of the set of available options, of the issues bearing on the decision, or of the preferences of the interested parties, might overcome all these handicaps to reach a good decision. Our central result is that if the full-information decision is Pareto-undominated among the interested parties—more precisely, if in every situation and for every proposed decision there is an interested party who is well informed, who has an opportunity to report, and who prefers the full-information decision to the proposed decision—then only the full information decision can be reached at equilibrium. For if any other decision were proposed, some interested party would find it advantageous to propose the full-information decision and provide enough information to support it.

Two sorts of applications emerge from this simple proposition. The first arises in situations where the decisionmaker seeks to maximize a Bergson-Samuelson social welfare function that is an increasing function of the interested parties' utilities. Some regulatory or legislative situations might be appropriately viewed in this light. The second concerns situations where the preferences of the interested parties are generally opposed, as they would be in many purchasing decisions or legal contests. In these applications competition among informed interested parties allows even an unsophisticated decisionmaker who fails to recognize the strategic incentives of the interested parties to reach the full-information decision.

The strategy to be employed by the decisionmaker in Section 3 does not make great demands on his rationality. It does not require the decisionmaker to know the possible states of information, the space of available decisions, or the preferences of the interested parties—or even to have beliefs about these things. Nor does the decisionmaker have to make any sophisticated inferences to unravel possible strategic dissembling by interested parties. The demands placed on the abilities of the interested parties to convey information and on the decisionmaker to process the information he receives are relatively severe, however.

In Section 4 we explore a class of problems in which the interested parties are unable to transmit all of their information; they are constrained to suggesting a limited range of alternatives, to suggesting dimensions to be considered in ranking the alternatives, and to reporting how the alternatives rank on these dimensions. We assume that the decisionmaker faces severe handicaps in dealing with the interested parties: He does not know their preferences, the set of alternatives, the particular attributes of any alternatives that may be suggested, or even the relevant dimensions on which alternatives might be ranked. We do assume, however, that the decisionmaker is sophisticated and that the interested parties are all equally well informed. We retain the assumption of Section 3 that the full-information decision is Pareto-undominated among the interested parties. Then, at the equilibrium of the persuasion game, competition among the interested parties leads them to reveal the dimensions or issues that are relevant to the decision, and the decisionmaker's skepticism extracts enough additional information to ensure that the full-information decision is reached.

The condition that there does not exist another decision that is unanimously weakly preferred to the full-information decision is quite useful for understanding when the full-information decision might not be reached. For example, the condition fails in a selling situation where the buyer can choose only one product to buy and can also buy nothing, thereby avoiding the whole product class. Then, as Posner (1977, pp. 136–137) has noted, if product safety is an issue, “the manufacturer who advertises that his product is safer than his competitors' product [*sic*] runs the risk of planting fears where none existed. This . . . may reduce the advertiser's sales by more than his claim of relative safety increases them.”

It is interesting to contrast the main conclusion of Section 2, which requires that the decisionmaker know the interested party's preferences, with that of Section 4, which has no such requirement. To understand this fully, one must understand something of the

nature of skeptical strategies. For general persuasion games, we say that a strategy exhibits *skepticism* if it calls for holding pessimistic expectations about alternatives that are (or are thought to be) favored by some well-informed interested party and about which little has been reported. When there is but a single interested party, skepticism involves forming expectations that downgrade his *most favored* alternatives—a strategy that requires information about his preferences. With several competing interested parties who do not unanimously prefer any alternative to the full-information decision, skepticism consists of forming pessimistic expectations about *every* alternative. This is effective because *somebody* favors the full-information decision over any other alternative; that party can be relied upon to suggest it and provide the information needed to justify it.

Our theory is a close cousin to the extensively developed theory of mechanism design. We review some similarities and differences between the theories in Section 5, and present our conclusion in Section 6.

## 2. One seller and a sophisticated buyer

■ We consider here a game with two players: for concreteness, we take them to be a seller and a buyer, although the model fits other situations as well. The single seller provides verifiable information about product quality to the buyer, who then decides how much to purchase. The seller wants to maximize sales; the buyer's objective is to maximize the expected utility of his consumption, which depends upon the quality of the product being offered.

The extensive form of this persuasion game is as follows. First, Nature selects a point  $x$ , representing the seller's information, from a finite set  $X$ . The buyer believes that the probability that any particular  $x$  has been chosen is  $P(x) > 0$ .<sup>4</sup> The seller observes  $x$  and makes an assertion  $A$  to the buyer;  $A$  is a subset of  $X$ . We restrict the seller to make true assertions, that is, we require that  $x \in A$ . The buyer observes  $A$  and then selects a quantity  $q \in \mathbb{R}_+$ . The payoffs are  $u(x, q)$  to the buyer and  $v(x, q)$  to the seller, where  $v$  is a function increasing in  $q$ . In particular, the monotonicity of  $v$  in  $q$  implies that the buyer knows the seller's ordinal preferences. Finally, we assume that for each  $x$  there is a unique  $q^*(x)$  that maximizes the buyer's utility.

The normal form of the game can be derived from the extensive form in the usual way. A *reporting strategy* for the seller is a function  $r$  from  $X$  to the subsets of  $X$  such that  $x \in r(x)$ ;  $r$  specifies what assertion the seller will make as a function of his information  $x$ . A *buying strategy*  $b$  for the purchaser is function  $b$  mapping subsets of  $X$  to purchase decisions in  $\mathbb{R}_+$ . When Nature chooses  $x$  and the strategies are  $r$  and  $b$ , the seller's payoff is  $v(x, b(r(x)))$  and the purchaser's is  $u(x, b(r(x)))$ .

To solve the game we use the concept of sequential equilibrium introduced by Kreps and Wilson (1982). For this game, a sequential equilibrium is described by a triple  $(r, b, p)$ , where  $r$  and  $b$  are the reporting and buying strategies, respectively, and  $p$  specifies what the buyer believes when the seller makes a report. Thus,  $p(x|A)$  is the probability that the buyer assigns to the information state  $x$  when the seller reports  $A$ . The triple is a sequential equilibrium in pure strategies if it satisfies the following four conditions:

- (i) *Seller maximization.*  $r$  is the seller's best response to  $b$ , that is,  $r(x)$  is the assertion  $A$  that maximizes  $v(x, b(A))$  subject to  $x \in A$ .
- (ii) *Buyer maximization.* For all  $A$ ,  $b(A)$  is the best purchase for the buyer, given his beliefs, that is, it is the quantity  $q$  that maximizes  $\sum u(x, q)p(x|A)$ .
- (iii) *Rational buyer expectations.* If  $A = r^x(z)$  for some  $z$ , then  $p(x|A)$  is  $P(x)/P(r^{-1}(A))$  for  $x$  in  $r^{-1}(A)$  and is zero otherwise.

<sup>4</sup> The seller may or may not know  $P$  or he may have imperfect information about it; that part of the specification of the game does not affect the solution.

(iv) *Consistent beliefs.*  $p(x|A) = 0$  for all  $x$  not in  $A$ .

For the same game, a Nash equilibrium can be described as a triple satisfying conditions (i), (iii), (iv), and a weakened form of (ii) as follows:<sup>5</sup>

(ii) *Nash buyer maximization.* For all  $A$  in the range of  $r$ ,  $b(A)$  is the best purchase for the buyer, given his beliefs, that is, it is the quantity  $q$  that maximizes  $\sum_x u(x, q)p(x|A)$ .

The difference between the Nash and sequential equilibrium concepts lies in their notions of what rational behavior is. Sequential equilibrium requires what Kreps and Wilson call “sequential rationality”—the buyer always maximizes when it is his turn, no matter what has previously happened. Nash equilibrium requires the buyer to maximize only if there have been no “surprises.”

This game has many Nash equilibria, some of which are quite implausible. For example, one Nash equilibrium  $(r, b)$  consists of the pair of strategies  $r(x) = X$  for all  $x$  and  $b(A) = q^*$  for all  $A$ , where  $q^*$  is the quantity that maximizes  $\sum_x P(x)u(x, q)$ . At this Nash equilibrium, the buyer is stubbornly determined (contrary to his own interests) to ignore any information offered by the seller, and the seller, believing that what he says is irrelevant, offers no information. But if there is any  $x$  such that the buyer’s best choice, given  $x$ , is to buy more than  $q^*$ , one should expect that the seller will try to communicate that information and that a rational buyer will pay heed. Henceforth, we shall focus primarily on the sequential equilibrium solution concept; we use the unmodified term “equilibrium” to refer to it.

How will a sophisticated buyer behave in this persuasion game? Let us call a pair  $(b, p)$  satisfying (ii) and (iv) a “posture” for the buyer. A *naively credulous* posture is one in which the buyer takes the seller’s report at face value and simply puts  $p(x|A) = P(x)/P(A)$  for  $x \in A$ . A *skeptical* posture  $(\bar{b}, \bar{p})$  is one such that, for every report  $A$ ,  $(\bar{b}(A), \bar{p})$  solves:

$$\text{Minimize } Q_{Q,p}$$

subject to

$$Q \text{ maximizes } \sum_x u(x, q)p(x|A)$$

$$p(\cdot | \cdot) \text{ a conditional probability}$$

$$p(x|A) = 0 \quad \text{for all } x \notin A.$$

Every posture requires that the buyer form beliefs consistent with his information and maximize accordingly. A *skeptical* posture minimizes (over all postures) the quantity the buyer will purchase. In the game we have described, there is always an equilibrium in which the buyer adopts a skeptical posture and the seller reports everything he knows. Indeed, somewhat more is true.

*Proposition 1.* If the buyer’s posture  $(b, p)$  is a skeptical posture and the full-information decision is always reached, i.e.,  $b(r(x)) = q^*(x)$  for all  $x$ , then  $(r, b, p)$  is an equilibrium.

It is straightforward to verify that any triple  $(r, b, p)$  satisfying the conditions of Proposition 1 is in fact an equilibrium. In particular, if  $(b, p)$  is a skeptical posture and  $r(x) = \{x\}$ , the hypotheses of the proposition are satisfied, so an equilibrium does exist.

With the assumptions made so far, there may be other equilibria in which the buyer

---

<sup>5</sup> The condition (iv) is actually irrelevant in the definition of a Nash equilibrium, since it is implied by (iii) for sets  $A$  in the range of  $r$  (Milgrom and Roberts, 1982). The present description emphasizes that the major difference between the equilibrium concepts lies in the fact that Nash equilibrium does not require maximizing behavior at every node in the game tree.

does not adopt a skeptical posture. For example, if the buyer's preferences would lead him to purchase more than any full-information quantity when he has no information, there is an equilibrium in which the seller is always silent ( $r(x) \equiv X$ ). But the set of equilibria shrinks to just those described in Proposition 1 when the buyer's utility function  $u$  is strictly concave and continuously differentiable in  $q$ . The argument goes as follows.

At a pure strategy equilibrium, for each state  $x$  the buyer will never buy less than the full-information, utility-maximizing quantity  $q^*(x)$ . Otherwise, the seller could not be optimizing; he could do better by telling the *whole truth* (reporting  $r(x) = \{x\}$ ). The concavity of the buyer's utility function and the uniqueness of the optimum,  $q^*(x)$ , ensure that if the buyer always buys at least the full-information quantity  $q^*(x)$  (as he must at any pure-strategy equilibrium) and sometimes buys an excessive amount, then he could do better by reducing his purchases slightly. Hence, at any pure-strategy equilibrium of the game we have described, the buyer always buys precisely the full-information quantity. (A formal proof, which allows for mixed-strategy equilibria as well, appears in the Appendix.)

*Proposition 2.* If the buyer's utility function is strictly concave and continuously differentiable in  $q$ , then at every equilibrium the buyer adopts a skeptical posture  $(b, p)$  and always purchases the full-information quantity:  $b(r(x)) = q^*(x)$  for all  $x$ .

Notice that the seller's equilibrium strategy is not unique, even though the buyer's response to it is. All that can be said about the seller's strategy is that when his information is  $x$ , he reports some  $A$  whose elements  $y$  are all (at least) *as favorable as*  $x$ , in the sense that  $q^*(y) \geq q^*(x)$ . Notice, too, that the equilibrium strategies of both players are independent of the prior distribution  $P$ , so the demands of Bayesian rationality and common knowledge priors are less extreme in our model than in most Bayesian game models.

An interesting variation of the game arises when the seller suspects that the buyer may be too unsophisticated to adopt a skeptical posture. Suppose that the seller believes that with some probability  $\pi$  ( $0 < \pi \leq 1$ ), the buyer will adopt some other posture, such as the naively credulous posture. We restrict attention to postures for which the buyer's purchases increase as the seller's report becomes more favorable. More precisely, a posture is *responsive to favorable information* if whenever  $q^*(x) \leq q^*(y)$  for all  $y \in A$ , and  $A \setminus \{x\}$  is nonempty,  $b(A \setminus \{x\}) \geq b(A)$ .

As before, we describe an equilibrium by a triple  $(r, b, p)$ , where  $r$  is the seller's reporting strategy and  $(b, p)$  is the posture of the sophisticated buyer. The strategy of the unsophisticated buyer is not specified as part of the equilibrium, since it is given exogenously.

*Proposition 3.* In the variant game suppose  $u$  is strictly concave and continuously differentiable in  $q$  and the unsophisticated buyer's posture is responsive to favorable information. Then the triple  $(r, b, p)$  is an equilibrium if and only if  $(b, p)$  is a skeptical posture for the buyer and  $r$  is a strategy that maximizes sales to the unsophisticated buyer. At equilibrium, the sophisticated buyer purchases the full-information quantity:  $b(r(x)) = q^*(x)$ .

Propositions 2 and 3 are proved in the Appendix.

Thus, at equilibrium, the seller acts as if the buyer were certain to be unsophisticated, and makes a larger sale if the buyer is, in fact, unsophisticated. Given our assumption that even an unsophisticated buyer is responsive to favorable information, the sales-maximizing report rules out all information states that are "less favorable" than the truth. At equilibrium, a sophisticated buyer adopts a skeptical posture and correctly infers that the actual information state  $x$  is the least favorable one consistent with the seller's report.

### 3. Competition among interested parties

■ In this section we relax substantially our assumptions about the sophistication and the prior information of the decisionmaker, but we introduce multiple interested parties who

compete in providing any information upon which the decision will be based. The question at issue is under what circumstances competition among providers of information can help to protect unsophisticated and ill-informed decisionmakers from the self-interested dissembling of information providers.

As examples, when does competition among sellers reveal actual product qualities? When does lobbying by interest groups help regulators and legislators reach better decisions? Does competition between divisions for corporate resources generally assist in making correct investment and capital budgeting decisions? As we shall see, competition may help both the credulous and the sophisticated decisionmaker by reducing the amount of prior information that they need about their sets of options, the relevant aspects of each option, and the interested parties' preferences, and by reducing the strategic sophistication that they need to interpret the messages they receive and to reach the full-information decision.

This persuasion game is structured as follows. First, Nature chooses a point  $x$  from the finite set  $X$  according to the distribution  $P$ . Then, each of  $N$  interested parties observes  $x$ .<sup>6</sup> Simultaneously, each suggests a set  $D_i$  of possible decisions  $d$ , chosen from a finite set  $\Delta$ , and asserts a true proposition (that is, a set  $A$  such that  $x \in A$ ). Let  $D = \cup D_i$  be the set of decisions suggested.

We model the decisionmaker as a naive automaton—not as a player in the game. The automaton takes the conjunction (intersection)  $I$  of the assertions and selects a mixed decision (that is, a probability distribution over  $D$ ) to maximize the objective function  $E[u(x, d)|I]$ .<sup>7</sup> The probabilities used by the automaton for the expected utility calculation are such that  $P(x|I)$  is zero for information states  $x$  not in  $I$ . For simplicity, we assume that the maximizing decision is unique for each  $I$  and  $D$  so that the chosen mixed decision will, in fact, be simply one of the suggested decisions  $d \in D$ . Recalling the notation of the last section, we call this decision  $b(I, D)$ . Let  $f(x) = b(\{x\}, \Delta)$  designate the full-information decision.<sup>8</sup>

The payoffs to the interested parties are denoted  $v_i(x, b(I, D))$ . Since these payoffs depend on  $x$ , they need not be known *a priori* to the decisionmaker, but they must in effect be verifiably reportable to him.

*Proposition 4.* Suppose that for every  $x$  and every decision  $d$  in  $\Delta \setminus \{f(x)\}$  there is some interested party who prefers the full-information decision  $f(x)$  to  $d$ . Then at every pure-strategy Nash equilibrium, the full-information decision is taken. Moreover, if there is no mixed decision with support in  $\Delta \setminus \{f(x)\}$  that is weakly preferred to  $f(x)$  by every interested party, then at every Nash equilibrium the decision  $f(x)$  is taken.

This proposition is supported by a simple argument. If there were an equilibrium with any decision other than  $f(x)$  being taken, then some interested party would prefer the full-information decision to the equilibrium outcome. That party could therefore do better by suggesting the full-information decision and reporting  $\{x\}$ , thereby contradicting the assumption of equilibrium. A similar argument establishes the second part of the proposition.

For ease of reference, we shall sometimes refer to the condition of Proposition 4 as the

<sup>6</sup> This assumption is similar to, but is stronger than, the assumption of the conventional argument that each party has access to all the information that favors its side. We could formalize the conventional argument by restricting the model so that each piece of information favors one side or the other, regardless of whatever other information may be reported, and then by allowing each party to know at least all the information that favors its side. We believe, however, that it is significant for some applications (for example, where the interested parties are competing sellers) that such restrictive assumptions about the nature of information are not needed to generate strong conclusions about the effects of competition.

<sup>7</sup> The restriction to truthful reporting ensures that the various reports are consistent in that  $I$  contains at least the true state  $x$ .

<sup>8</sup> This assumption implies that at any pure-strategy equilibrium of the persuasion game, some particular pure decision will result. Then, pure-strategy equilibria can be analyzed by using only the ordinal preferences of the interested parties.



assumption that the full-information decision is Pareto-optimal. Actually, however, the assumption is a bit stronger than that, since it requires that no other decision be even Pareto-indifferent to it, that is, that no other decision can be so good in the eyes of every interested party.

Proposition 4 has three easy and useful corollaries. First, consider the case of a decision made by a regulatory body that seeks to advance the welfare of the various constituencies affected by the decision. Suppose that each constituency has interests that are aggregated and represented (honestly) by a lobbyist, and that each lobbyist knows all the relevant information,  $x$ . Then the regulator's payoff is  $u(v_1, \dots, v_n)$ , where  $u$  is an increasing function and  $v_i$  is the utility of constituency  $i$  corresponding to the decision taken. We assume, as above, that there is always a unique maximizing decision for the decisionmaker. Plainly, any full-information decision is a Pareto-optimal one.

*Corollary 1.* At every pure-strategy Nash equilibrium of this "persuade the regulator" game the equilibrium decision is the full-information decision.

Notice that, in particular, the regulator can rely on the lobbyist to suggest the full-information decision.

A second variation arises when the interests of the parties are *strongly opposed*, that is, for every  $x$  and every pair of (possibly mixed) decisions  $d$  and  $d'$  there are interested parties  $i$  and  $j$  such that  $v_i(x, d) > v_i(x, d')$  and  $v_j(x, d') > v_j(x, d)$ . For example, the decisionmaker may be deciding how to allocate a given volume of purchases at predetermined prices among a group of suppliers. The archetypal example of competition in information provision—the adversary system in legal disputes—also would often involve strongly opposed interests in this sense.

*Corollary 2.* At every pure-strategy Nash equilibrium of the persuasion game with strongly opposed interests, the equilibrium decision is the full-information decision.

A third variation, which perhaps better models the selling game, arises when the sellers not only provide information, but also quote prices for their goods. More precisely, suppose each seller has a product to offer whose cost of production is  $c_i(x)$ . If a sale is concluded at price  $p$ , the seller's payoff is  $p - c_i(x)$ . The buyer's utility from purchasing product  $i$  at price  $p$  is  $u(x, i) - p$ . Moving simultaneously, sellers name prices  $p_i$  for their products and make reports  $r_i(x)$ . The buyer then selects one of the products to purchase. Finally, all players receive their payoffs from this "persuasion and pricing" game.

Since all moves are simultaneous, nothing is lost if we choose to think of the sellers as first all setting prices and then, without knowledge of the prices set by others, all making reports. Now, at any *pure-strategy* Nash equilibrium of this game the sellers at the second stage will all act as if they knew the prices that their competitors had set, and Corollary 2 applies. Consequently, given the prices, the buyer will make the full-information decision. Anticipating that result, the sellers will set prices at the first stage as if playing a price-setting oligopoly game with a single fully informed buyer. Thus, we have the following result.

*Corollary 3.* At every pure-strategy Nash equilibrium of the persuasion and pricing game, the equilibrium choice and price are the same as in the corresponding full-information price-setting game.

#### *Remarks.*

(1) Corollaries 2 and 3 would change if we allowed the buyer to purchase nothing and to obtain some utility  $u^*(x)$ . In that case the arguments made above lead to the conclusion that at every pure-strategy Nash equilibrium, in any state  $x$  where the outcome of the full-information price-setting game involves buying from some seller, the equilibrium choice and price are as in the full-information game. An unsophisticated buyer, however, may

sometimes be fooled into making a purchase when he should not. For example, competition among cigarette producers will not lead them to reveal that cigarette smoking may shorten the smoker's life.

(2) It is fair to say that the decisionmaker in this model may have "no idea" what the range of alternatives is, and may have little idea about the possible states of the world. The prior probability distribution  $P$  on information states  $X$  plays no role at all, since the game is formally one of complete information. All the *players*—the interested parties—know the state  $x$  precisely.

(3) Proposition 4 and its corollaries are stated for Nash equilibria, rather than for sequential equilibria as used in the last section. The difference arises because we have specified in the structure of the game how the decisionmaker uses the information provided to him, whereas, in the previous section, the information use assumption was introduced through the equilibrium concept. Here, the Nash and sequential equilibria coincide.

#### 4. Competition and sophistication

■ The results of the last section show that, when the interested parties are all fully informed and able to report all they know and when the full information decision is Pareto optimal for them, competition in suggesting alternatives and providing information can obviate the need for the decisionmaker to be well informed and sophisticated. In this section we relax the assumption that all interested parties can report all they know, but we reintroduce sophistication on the part of the decisionmaker in drawing inferences. We also spell out the dimensions of uncertainty by adding a special structure to the information-state space.

Our model is designed to represent a situation in which not only is the decisionmaker ignorant of the set of possible alternatives, the facts necessary to evaluate the alternatives, and the preferences of the interested parties, but also he does not know the relevant dimensions on which each alternative should be evaluated. For example, the consumer who buys a new forced-air furnace may remember to ask about prices, maintenance costs, and standard fuel-efficiency ratings, but may forget to ask about how quietly it operates or how well it will function with the existing ductwork. Similarly (to recall a historical example), a Department of Defense analyst reviewing an Air Force proposal for a Rapid Deployment Force may not ask whether the huge, newly proposed troop and equipment carrying plane (the C5-A Galaxy Transport) will be wide enough to accommodate the *next* planned generation of tanks (Weintraub, 1980). We do assume, however, that, given a set of alternatives to evaluate and a set of relevant attributes, the decisionmaker can assess the information available about the attributes and can anticipate the strategic dissembling of the information providers.

We also assume that there is too much "relevant" information for any interested party to report it all and that the interested parties cannot verifiably report their own preferences. Accordingly, we limit each interested party to suggesting one or more alternatives, naming some set of relevant attributes or dimensions, and providing information about the standing of his suggested alternatives on each indicated dimension.<sup>9</sup> The parties may not report information about one another's preferences, although the decisionmaker might be able to infer something about them from the suggested alternatives.

The extensive form of this persuasion game is described as follows. First, Nature determines the set of relevant attributes  $Z$ , which is some finite subset of the set of possibly relevant attributes  $\mathcal{Z}$ . Nature also determines grades  $x_{dm}$  of each possible decision  $d \in \Delta$  on each relevant attribute  $m \in Z$  and a parameter  $y$  which may affect the interested parties' preferences (but not those of the decisionmaker). The number and the identities of the

---

<sup>9</sup> This rules out the use of comparative advertising by a seller. When the buyer is too unsophisticated to make his own comparisons, this arrangement can lead to a reduction in his welfare.

elements in  $Z$  are random, so that the decisionmaker must rely on the interested parties to identify the relevant attributes. Each grade  $x_{dm}$  is selected from some finite set  $X$ . In sum, Nature chooses a triple  $\omega = (Z, x, y)$  according to some probability distribution  $P$ .

Each interested party observes  $\omega$ . The parties then simultaneously make assertions to the decisionmaker. Interested party  $i$  suggests a set of alternatives  $D_i$  from some feasible set of suggestions  $\mathbf{D}$  with the property that for each  $d \in \Delta$  there is some  $D \in \mathbf{D}$  such that  $d \in D$ . Thus, any alternative can be suggested, but not necessarily in isolation. Some alternatives, once suggested, may necessarily call to mind a whole set of variants. Interested party  $i$  also reports a subset  $A_i$  of  $Z$ , interpretable as the relevant attributes that he chooses to identify, and, for each suggested alternative  $d_{ij} \in D_i$ , sets  $A_{ijm}$  which represent verifiable assertions about the grade of alternative  $d_{ij}$  on each attribute  $m \in A_i$ .

The decisionmaker collects the suggested alternatives and hears all the reports. We represent the information contained in the reports by the letter  $I$ . He then selects a decision  $d$  from the set of suggested alternatives  $\mathbf{D}$ . The decisionmaker's payoff is  $u(Z, x_d, d)$ , where  $x_d$  is the list of actual attributes for the decision  $d$  actually taken. The  $i$ th interested party's payoff is  $v_i(\omega, d)$ .

A sequential equilibrium is now defined very much as in Section 2. Noting that for any fixed strategy of the decisionmaker there is an induced game among the interested parties, we may describe a sequential equilibrium of the overall game as an  $(n + 2)$ -tuple  $(r_1, \dots, r_N, b, p)$  such that: (i) given the decisionmaker's strategy  $b$ , the reporting strategies  $r_i$  of the interested parties form a Nash equilibrium of the induced game; (ii) the decision strategy  $b$  is optimal, given the decisionmaker's beliefs; (iii) these beliefs satisfy a rational expectations condition; and (iv) the beliefs are consistent with the decisionmaker's information.

Observe that the analysis of Section 3 does not apply directly to this setting. The reason is that an interested party cannot report fully about all the alternatives to convince the decisionmaker that a particular alternative is best. The decisionmaker must do some of the work on his own and discard alternatives whose advocates do not justify them adequately.

On the basis of the analysis in Section 2, one might think that to adopt an effective skeptical posture the decisionmaker needs to know the interested parties' preferences (that is, the value of  $y$ ). In the presence of sufficiently intense competition among the interested parties, however, a skeptical posture entails skepticism towards *the alternatives themselves*; the decisionmaker need not consider exactly how the interested parties' preferences color their reports. More precisely, a skeptical posture entails the decisionmaker's believing that  $Z$  is the union of the  $A_i$ 's and adopting beliefs about the attribute ranks of each suggested decision  $d$  that, while consistent with  $Z$ 's being the union of the  $A_i$ 's and with the information provided, result in the lowest possible expected utility for alternative  $d$ . With this posture each interested party is required to prove the merits of his suggestion: any attribute of any suggested alternative not proved to rank high will be regarded as if it were proved to rank low. With such beliefs and the corresponding optimizing choices by the decisionmaker, if the full-information decision is Pareto-undominated (as previously defined), then it will be in someone's interests to suggest it and to provide supporting information.

*Proposition 5.* Suppose that for no  $\omega$  and no  $d \in \Delta$  is it true that  $d$  is preferred or indifferent to the full-information decision at  $\omega$  by all interested parties. Then there exists a sequential equilibrium at which the decisionmaker adopts the skeptical strategy described above. At every such equilibrium, the decision reached is the full-information decision.

The existence claim is supported by having each interested party suggest the full-information decision, report the full-attribute set, and provide accurate information about the suggested decision. The characterization of all equilibria involving the skeptical strategy follows by noting that once the skeptical posture is adopted, the argument associated with Proposition 4 applies.

In this game there may exist other equilibria as well at which the decisionmaker does not behave skeptically and the full-information decision is not reached. These other equilibria would, of course, be destroyed if we modified the game to allow the opposing parties to rebut each other's alternatives. In a situation like this with multiple equilibria, it is to the decisionmaker's advantage to select (if he can) the one that favors him, for example by announcing his intention to play the skeptical strategy. That equilibrium seems to be a focal point, since there are many specifications of the information-state space for which it is the only equilibrium.<sup>10</sup> Thus, it is reasonable for us to focus on skeptical behavior as a descriptive account of the behavior of rational decisionmakers facing the kind of uncertainties considered here.

## 5. Comparison with mechanism design

■ The questions we have studied in this article are related to ones that have been studied in the burgeoning economic literature on "mechanism design."<sup>11</sup> In its standard formulation the mechanism designer's problem is to select rules for an institution that advances his objectives by exploiting the private information of one or more individuals or by motivating the individuals to take prescribed actions, or both. Despite the similarity of the problems studied, there are several major differences between the models we have used here and the kinds used in the mechanism design literature.

First, we have focused our attention on general purpose institutions, ones that can be used in a variety of different decision environments and can even be implemented by a decisionmaker with little idea of what the environment is. In mechanism design theory the recommended mechanism is often a function of such fine details of the environment as the exact form of the various agents' prior beliefs.<sup>12</sup> Finely tuned mechanisms may be of limited use to a decisionmaker who knows little about what the relevant environment is, and, indeed, the institutions we actually observe do not typically use the detailed information that is assumed to be common knowledge in theories of mechanism design. Here we have shown that, when reported information can be verified, a decisionmaker can sometimes do quite well with a general purpose mechanism. Moreover, in the games we have studied, the behavior on the part of the interested parties called for by the equilibrium is quite straightforward, so that the assumption of equilibrium does not seem strained. In contrast, equilibrium behavior in theories of mechanism design is often very complicated. Note, too, that it should generally be easier to test theories of general purpose mechanisms, since their predictions do not depend on the unobservable beliefs of the mechanism designer.<sup>13</sup>

Second, we have assumed that the decisionmaker has no power to restrict the kinds of reports that the interested parties can give (other than to ensure that they are consistent with the true state), and that the decisionmaker neither ignores information nor takes any

---

<sup>10</sup> For one example, suppose that each interested party is a potential supplier of a homogeneous product, which not all the sellers have available. The decisionmaker must decide from whom to buy, and how much. The only relevant attributes are who carries the product, at what prices, and the quality of the product. Then the logic of Proposition 2 implies that the skeptical strategy is the only equilibrium strategy for the buyer.

<sup>11</sup> The subject of optimally designing economic institutions or "mechanisms" was introduced by Hurwicz (1960, 1972). Among the many applications of Bayesian mechanism theory are ones to the design of income tax schemes (Mirrlees, 1971), incentive contracts (Harris and Raviv, 1979), and regulatory procedures (Baron and Myerson, 1982), to name just a few.

<sup>12</sup> The sensitivity of optimal schemes to fine details of the environment is particularly evident in the optimal auction work of Matthews (1983), Maskin and Riley (1984), and especially Cremer and McLean (1985).

<sup>13</sup> Assuming that the decisionmaker is sophisticated, our theory generates several testable hypotheses. These include: (i) that the decisionmaker adopts a skeptical strategy; (ii) that a salesman will not differentiate his treatment of customers who are probably sophisticated and those who are probably not and, since that information is valueless to him, will not seek it; and (iii) that with sufficiently opposed interests or in appropriate regulatory settings, efficient decisions will be reached.

decision that is not optimal, given his beliefs. Mechanism design approaches normally assume that the decisionmaker can set the rules of the game to restrict the options of interested parties and to commit himself to making decisions that will not be in his interests *ex post*. For most of the cited applications, it is difficult to justify such an asymmetric treatment of the decisionmaker and the interested parties: Why are not all the parties equally able to commit? It is by no means certain that commitment can often be achieved in the situations we wish to study. One can sometimes break commitments by asserting that the underlying conditions on which the commitment was premised have not been met. When enforcement costs are high, one can simply renege on a so-called commitment. As a result, the commitments required in the mechanism design approach may not be credible. We have shown that the decisionmaker can sometimes do as well without the use either of commitments or of restrictions on the interested parties as he could do with these devices. Consequently, our theory applies even in those situations where the ability of the decisionmaker to control the rules and to achieve commitment is limited or uncertain.

Third, we allow the decisionmaker a far greater range of uncertainty than is common in the mechanism design literature. Our decisionmaker may not know the alternatives that are available or even those that might possibly be available, and may be forced to rely on interested parties for suggestions. Such uncertainties make it impossible to formulate prior beliefs about the set of alternatives, and so rule out the use of Bayesian decision theory. Uncertainty of this kind is an important aspect of reality, and it is a significant finding that competition among interested parties sometimes alleviates this uncertainty. Identifying mechanisms that work well in the face of such thoroughgoing uncertainty lies wholly outside the realm of the traditional approach to mechanism design, since that theory requires Bayesian priors on everything to define an objective function for the optimizing process.

Finally, the models used in this article deal with verifiable information, in the same spirit as earlier work by Milgrom (1981) and Grossman (1981). Research in mechanism design has often dealt with information about variables, like personal taste, for which direct verification may be impossible. In such cases the decisionmaker must provide incentives to the interested parties for reporting their information truthfully. Both perfect verifiability and perfect privacy are highly limited as models of reality, yet both shed some light on the important intermediate cases in which some, but not all, information can be verified.

## 6. Conclusion

■ We have used game theory to examine the logic of the argument that when all interested parties have access to complete and verifiable information, competition among them in attempting to influence a decision leads to the emergence of “truth” or, more precisely, of all relevant decision alternatives and information. Some parts of our analysis apply to the case of a buyer’s being courted by many sellers; other parts apply to governmental proceedings in which all interested parties are represented. Our analysis has obvious relevance for persuasive situations within firms, as well as for legal contests, legislative battles, regulatory hearings, etc. It indicates that, at least in some situations, skepticism on the part of the decisionmaker and/or competition among interested parties can result in the emergence of all the relevant information and the selection of an optimal decision.

The scope of the conclusion that competition leads to the revelation of truth is in some respects wider and in others narrower than would appear to be commonly thought. It is not always true in competitive situations that each piece of information favors one of the interested parties. Interested parties then may not know which piece favors them, and so they may unwittingly withhold even favorable information. On the other hand, even if the parties do not have access to all information, or if they cannot report all that they know, rational skepticism by a decisionmaker can lead to a full-information decision by inducing one party to reveal information that is damaging to its interests. The party reveals this

information for fear that withholding it will lead to an *even more unfavorable supposition* by the skeptical decisionmaker.

One of the most striking results of our analysis has been the prominent role played by skeptical strategies. Evidently, sophisticated decisionmakers must often adopt skeptical strategies.

Informally, skepticism means that one assumes the interested parties' reports are tinged by self-interest. Within the models we have studied, sophisticatedly skeptical behavior consists of the decisionmaker's systematically downgrading the alternatives that he suspects may be favored by a well-informed interested party. When the parties' preferences are known, a skeptical strategy is relatively easy to describe and implement. When they are unknown, but the interested parties' interests are strongly enough in conflict, the decisionmaker can safely assume that every relevant decision is preferred by someone. Then, skepticism consists of downgrading *every* suggested alternative. We expect subsequent research to show that sophisticated decisionmakers in more general settings use a strategy of "calibrated skepticism" by downgrading most those alternatives that are expected to find articulate, well-informed advocates and by relying on the advocates to provide the information to justify their favored decisions.

## Appendix

■ The proofs of Propositions 2 and 3 follow.

*Proof of Proposition 2.* Since the buyer's utility function is strictly concave in  $q$ , so is his expected utility, given any information. Thus, it will never be optimal for the buyer to adopt a mixed strategy. To accommodate mixed strategies on the part of the seller, we allow that  $r(x)$  may be a random variable. Then, for an equilibrium the seller-maximization condition (i) must hold with probability one. In particular, this implies that at equilibrium  $b(r(x))$  is a constant, even though  $r(x)$  may be a random variable.

In view of Proposition 1, it is only necessary to check that every equilibrium triple  $(r, b, p)$  has the specified form.

The seller-maximization condition of equilibrium requires that for every  $x$ , the seller must weakly prefer reporting  $r(x)$  to reporting  $\{x\}$ , that is,  $b(r(x)) \geq b(\{x\})$  for all  $x$ . If this inequality is strict for some  $x'$ , then the concavity of  $u$  implies that

$$E[\partial u(x, b(r(x')) + q)/\partial q | p(\cdot | r(x'))] < 0, \quad (A1)$$

and the buyer could do better by reducing his purchases in response to the report  $r(x')$ . This contradicts the buyer-maximization condition of equilibrium. Hence,  $b(r(x)) = b(\{x\})$  for all  $x$ . It remains to show that the buyer adopts a skeptical posture  $(b, p)$ , that is, for any report  $A$ ,  $b(A) = \min \{b(\{x\}) | x \in A\}$ .

The results that  $b(r(x)) = b(\{x\})$ , together with the seller-maximization equilibrium condition, imply that for all  $A$  and all  $x \in A$ ,  $b(\{x\}) \geq b(A)$  (otherwise the seller does better to report  $A$  when the state is  $x$ ). Therefore,

$$b(A) \leq \min \{b(\{x\}) | x \in A\}.$$

Suppose that there is some  $A$  for which this inequality is strict. Then, using the consistent-beliefs equilibrium condition and the strict pseudo concavity of the buyer's preferences, one obtains

$$E[\partial u(x, b(A) + q)/\partial q | p(\cdot | A)] > 0, \quad (A2)$$

which contradicts the buyer-maximization condition. *Q.E.D.*

*Proof of Proposition 3.* Let  $(r, b, p)$  be a purported equilibrium and let  $r^*$  be a strategy that maximizes sales to a naively credulous buyer.

First, we observe that, since the unsophisticated buyer is "responsive" (as defined in the body of the article),  $r^*$  must have the property that for all  $x' \in r^*(x)$ ,  $b(\{x\}) \leq b(\{x'\})$ . (Otherwise, reporting  $r^*(x) \setminus \{x'\}$  would result in more sales to the unsophisticated buyer, thereby contradicting the definition of  $r^*$ .)

Using the just-proved property, the strict concavity of buyer preferences, and the consistent-beliefs condition, we find that for all  $Q < b(\{x\})$ ,

$$E[\partial u(x, Q + q)/\partial q | p(\cdot | r^*(x))] > 0. \quad (A3)$$

Hence,  $b(r^*(x)) \geq b(\{x\})$  for all  $x$ . Now, once we show that  $b(r(x)) = b(\{x\})$  for all  $x$  for any equilibrium strategy  $r$ , we shall have established that  $r$  sells no more to sophisticated buyers than  $r^*$ . Therefore,  $r$  cannot be a best

response unless, like  $r^*$ , it also maximizes sales to unsophisticated buyers, and it will follow that any equilibrium  $r$  must maximize sales to unsophisticated buyers.

If for all  $x$ ,  $b(r(x)) \geq b(\{x\})$  and there is strict inequality for some  $x'$ , then by the consistent-beliefs condition and the strict concavity of buyer preferences, (A1) holds and thus violates the buyer-maximization condition. This leaves two possibilities: either  $b(r(x)) < b(\{x\})$  for some information state  $x$ , or  $b(r(x)) = b(\{x\})$  for all  $x$ . Suppose first that  $b(r(x)) < b(\{x\})$  for some information state  $x$ . For that  $x$ , reporting  $r^*(x)$  increases sales to sophisticated buyers compared with  $r(x)$  and maximizes sales to naive buyers. This contradicts the seller-maximization equilibrium condition. Hence, this case cannot arise at equilibrium, and we conclude that  $b(r(x)) = b(\{x\})$ , which completes the proof. *Q.E.D.*

## References

- BARON, D. AND MYERSON, R. "Regulating a Monopolist with Unknown Costs." *Econometrica*, Vol. 50 (1982), pp. 911-930.
- CRAWFORD, V. AND SOBEL, J. "Strategic Information Transmission." *Econometrica*, Vol. 50 (1982), pp. 1431-1451.
- CREMER, J. AND MCLEAN, R. "Full Extraction of Surplus in Bayesian and Dominant Strategy Auctions." University of Pennsylvania, CARESS Working Paper #85-17, 1985.
- GROSSMAN, S. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics*, Vol. 24 (1981), pp. 461-483.
- HARRIS, M. AND RAVIV, A. "Optimal Incentive Contracts with Imperfect Information." *Journal of Economic Theory*, Vol. 20 (1979), pp. 231-259.
- HAYEK, F. A. *The Road to Serfdom*. Chicago: University of Chicago Press, 1972.
- HOLMSTRÖM, B. "On Incentives and Control in Organizations." Ph.D. Dissertation, Stanford University, 1977.
- HURWICZ, L. "Optimality and Informational Efficiency in Resource Allocation Processes" in K.J. Arrow, S. Karlin, and P. Suppes, eds., *Mathematical Methods in the Social Sciences*, Stanford: Stanford University Press, pp. 27-46.
- . "On Informationally Decentralized Systems" in C.B. McGuire and R. Radner, eds., *Decision and Organization*, New York: American Elsevier Publishing Company, 1972.
- KREPS, D. AND WILSON, R., "Sequential Equilibria." *Econometrica*, Vol. 50 (1982), pp. 863-894.
- MASKIN, E. AND RILEY, J. "Optimal Auctions with Risk-Averse Buyers." *Econometrica*, Vol. 52 (1984), pp. 1473-1518.
- MATTHEWS, S. "Selling to Risk-Averse Buyers with Unobservable Tastes." *Journal of Economic Theory*, Vol. 30 (1983), pp. 370-400.
- MILGROM, P. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics*, Vol. 12 (1981), pp. 380-391.
- AND ROBERTS, J. "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis." *Econometrica*, Vol. 50 (1982), pp. 443-459.
- MILL, J.S. *On Liberty*. Indianapolis: Bobbs-Merrill, 1977.
- MIRRLIES, J. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies*, Vol. 38 (1971), pp. 175-208.
- POPPER, K.R. *The Open Society and Its Enemies, Vol. 1*. Princeton: Princeton University Press, 1966.
- POSNER, R. *Economic Analysis of Law*, 2nd ed. Boston: Little, Brown and Company, 1977.
- SOBEL, J. "A Theory of Credibility." *Review of Economic Studies*, Vol. 52 (1985), pp. 557-574.
- WEINTRAUB, B. "Warning Crucial to Air and Sea Transport Abilities." *New York Times* (September 26, 1980).
- WILDE, L. "Labor Market Equilibrium under Nonsequential Search." *Journal of Economic Theory*, Vol. 16 (1977), pp. 373-393.