

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 749

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than acknowledgment that a writer had access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

RELYING ON THE INFORMATION OF INTERESTED PARTIES

by

Paul Milgrom

and

John Roberts

May, 1985

RELYING ON THE INFORMATION OF INTERESTED PARTIES

July, 1984
(revised May, 1985)

Paul Milgrom*
Yale University
135 Prospect, New Haven, CT. 06520

John Roberts**
Graduate School of Business
Stanford University, Stanford CA. 94305

*Department of Economics, Cowles Foundation, and School of Organization and Management, Yale University.

**Graduate School of Business, Stanford University.

Milgrom's work was supported by National Science Foundation Grant IST-8411595 to Yale University and by the Institute for Advanced Studies of the Hebrew University of Jerusalem. Roberts was supported by National Science Foundation Grant SES-8308723 to Stanford University. We thank the participants at the Stanford Summer Accounting Workshop, the Northwestern University Summer Workshop on Strategic Behavior and Competition, the Cowles Foundation Seminar at Yale, the Theory Seminar at Princeton, and the Stanford University Conference on Adaptive Institutions for their suggestions and comments.

ABSTRACT

We investigate the conventional wisdom that competition among interested parties attempting to influence a decisionmaker by providing verifiable information brings out all the 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, while if the decisionmaker is unsophisticated or not well informed, competition is not generally sufficient. However, if the interested parties' interests are sufficiently opposed, 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.

Journal of Economic Literature Classification Number

Primary: 026

Secondary: 612, 613

Keywords: Persuasion Games, Revelation Games, Lobbying, Regulation, Strategic Information Transmission, Adversary System, Asymmetric Information.

RELYING ON THE INFORMATION OF INTERESTED PARTIES

by Paul Milgrom and John Roberts

I. INTRODUCTION

A common problem faced by decisionmakers is the need to rely on the suggestions and information provided by individuals who are interested in the decision. Although interested parties may try to manipulate the decision by concealing and distorting information, their efforts do not always succeed. An archtypical example of an institution that copes successfully with the problem of interested parties is the adversary system used to resolve legal disputes, for which it is often argued -- to put it crudely but simply -- 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 becoming known and appropriate decisions being made in the context of a wide variety of political, scientific, regulatory, and economic institutions.

In this paper, we examine the validity and scope of that argument by studying the problem of a decisionmaker who must rely on interested parties to provide information about possible decisions and their consequences. The focus of our work is to identify those conditions under which skepticism on the part of the decisionmaker about the claims of interested parties and competition among the interested parties in providing information permits the full information optimal decision to be reached, despite

possible severe limitations on the decisionmaker's prior information, on his ability to draw sophisticated inferences, and on the capacity of the interested parties to communicate the information they have.

To study 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. (See Sobel [1983] for an analysis of reputation building in information provision.) We assume that the interested parties will use guile if they can to attempt to dupe the decisionmaker, but that the decisionmaker can freely verify any information that is provided. For example, a seller might verifiably report that his product meets or exceeds standard "XYZ" when that standard is out-of-date or is just barely met, but at the same time might simply fail to mention a more relevant standard that the product does not meet.

The assumption that information can be freely verified is important for our analysis, though 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 Holmstrom [1977] and Crawford and Sobel [1982], among others.

In section II, we investigate how a sophisticated decision-maker might use skepticism as a weapon to extract maximum information from a single interested party. The 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 (a seller wants to sell as much as possible, an electric utility company prefers 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 one of extreme skepticism ("assume the worst") and (ii) the equilibrium decision is the "full information decision" -- the same decision as would have been reached if the decisionmaker had perfect access to the interested party's information.

Next, we study the case where the interested party is unsure about whether the decisionmaker is sophisticatedly skeptical or naively credulous; a credulous decisionmaker is one who interprets the information reported to him without regard to the reporting party's interests, as though it were the result of an objective experiment. Even though the interested party is unsure of the decisionmaker's rationality, his equilibrium reporting strategy is uniquely determined as the strategy that elicits the most favorable decision from a naive decisionmaker. Thus, for example, the rational salesman treats every buyer as if he were naive. The resulting masking of information does not harm the sophisticated decisionmaker: At equilibrium, he interprets the

reports correctly and reaches the full information optimal decision.

The hypotheses needed in section II to generate those happy conclusions are many. First, one needs to assume that the set of decisions is one-dimensional and that the interested party's preferences are known and monotone on that dimension. For example, this assumption is satisfied if the interested party is a seller who is offering a single product and wants to sell as much of it as possible, or 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 the several brands offered 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 in order 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 three conditions fail, the decisionmaker will be unable to implement the strategy of extreme skepticism, and he will suffer a loss of utility as a result. Finally, the decisionmaker must also be sophisticated enough to draw the proper inferences and reach the right decision when all relevant information has been

made available; otherwise, the full-information decision is not an appropriate welfare standard.

In section III, 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, the issues bearing on the decision, or the preferences of the interested parties might overcome all these handicaps to reach a good decision. Our main result is that if in every situation and for every proposed decision, there is some interested party who is well informed, has an opportunity to report, and 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 specific applications emerge from this line of reasoning. The first concerns 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 sufficiently opposed, as they would be in many purchasing decisions or legal contests. In each of these areas of application, competition between the informed interested parties allows even a naive decisionmaker who fails to recognize the strategic incentives of

the interested parties to reach the full information optimal decision.

These conclusions of section III do not formally require the decisionmaker to know the possible states of information, the space of available decisions, or the preferences of the interested parties. Nor does the decisionmaker have to make any sophisticated inferences to unravel possible strategic dissembling by interested parties. However, relatively severe demands are placed on the abilities of the interested parties to convey information and on the decisionmaker to process the information he does receive.

In section IV, we explore a class of problems in which the interested parties are unable to transmit all of their information. As well, the decisionmaker does not know even the relevant dimensions on which alternatives could be ranked, let alone the actual rankings. We do assume, however, that the decisionmaker is sophisticated and that the interested parties are all equally well informed. We find that, at equilibrium, competition among the interested parties reveals the dimensions or issues that are relevant to the decision. Then skeptical, sophisticated reasoning by the decisionmaker extracts the remaining information, so that the full information decision is reached.

Our theory is a close cousin to the extensively developed theory of mechanism design. Some similarities and differences between the theories are reviewed in section V followed by our conclusion in section VI.

II. 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 buy. The seller wants to sell as much as possible; 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 $P(x) > 0$. 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. The seller observes x , and then selects a true assertion to make to the buyer, which is a subset A of X such that $x \in A$. The buyer observes the assertion A and then selects a quantity $q \in R_+$. The payoffs are $u(x,q)$ to the buyer and $v(x,q)$ to the seller, where v is 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 a function b mapping subsets of X to purchase decisions in 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 $u(x, b(r(x)))$.

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 $\int P(x)u(x, q)$. At this equilibrium, the buyer is stubbornly determined 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 respond to the information that is revealed. Nash equilibrium allows the players to be implausibly pig-headed, so we shall want to use an equilibrium concept that does not permit a rational buyer to ignore any useful, verifiable information he receives.

To force the buyer to listen when spoken to, we employ the concept of a sequential equilibrium introduced by Kreps and Wilson [1982]. In 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 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: $b(A)$ 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 $A = r(z)$ for some z then $p(x|A)$ is $F(x)/P(r^{-1}(A))$ for x in $r^{-1}(A)$ and is zero otherwise, and

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

Conditions (i), (ii), and (iii) can be used as an alternative to the standard definition of a Nash equilibrium: The equilibrium strategies permitted by the two definitions are identical (cf. Milgrom and Roberts [1982]). It is the addition of condition (iv) that distinguishes the sequential equilibria from the broader class of Nash equilibria by requiring the buyer to "listen" when verifiable information is provided. In the developments that follow, we use the unmodified word "equilibrium" to refer to a sequential equilibrium, while employing the term "Nash equilibrium" when condition (iv) is not relevant (see Proposition 4, its corollaries, and Proposition 5).

Let us call a pair (b,p) satisfying (ii) and (iv) a "posture" for the buyer. A *naive* 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 is one that satisfies:

$$b(A) = \text{Min}_{Q,p} Q$$

subject to

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

$p(\cdot|A)$ a conditional probability

$$p(x|A) = 0 \text{ for all } x \text{ not in } A.$$

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 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 never makes an informative report ($r(x)=X$). However, 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 $q^*(x)$. For 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 optima $q^*(x)$ ensures 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 too much, 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, is given 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 π ($0 < \pi < 1$), the buyer will adopt some other, more credulous, posture, such as the naive posture. This other posture is exogenously specified as part of the game. We require only that the credulous posture be responsive to favorable information. 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 . 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 credulous 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 credulous. Given our assumption that even a credulous buyer is responsive to favorable information, this implies that the seller's report rules out all information states which 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.

Thus, there are no externalities worked among the buyers: The possible presence of a naive buyer does not cause the seller to effectively withhold information from the sophisticated buyer, nor does the possible existence of a sophisticated buyer force the seller to alter his reporting strategy in a way that helps or harms the naive buyer. It always pays the seller to assume that his customer is naive and it always pays the rational buyer to be skeptical.

III. COMPETITION AMONG INTERESTED PARTIES

In this section, we relax substantially our assumptions about the sophistication and the prior information of the decisionmaker but 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 strategically unsophisticated, 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 naive 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 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 . (This captures -- but is stronger than -- the assumption of the conventional argument that each party has access to all information that favors its side.) Simultaneously, each suggests a set D_i of possible decisions d , chosen from a finite set Δ , 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.

The decisionmaker is modeled 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 $E[u(x,d)|I]$. (The restriction to truthful reporting ensures that the various reports are consistent in that I contains at least the true state.) The probabilities used by the automaton for the expected utility calculation are some probabilities 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.

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 $\{x\}$, 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 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 as 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 which 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 decision-

maker. 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 buy. 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, sellers choose 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 added an option allowing the buyer to purchase nothing and 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. However, a naive buyer (in the sense of Proposition 3) 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.

IV. 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 decision-maker 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 quietness of operation or how well the furnace will function with the existing ductwork. Similarly (to recall an historical example), a Department of Defense analyst reviewing an Air Force proposal for a Rapid Deployment Force may forget to ask whether the huge, newly proposed troop and equipment carrying plane (the C5-A) will be wide enough to accommodate the next planned generation of tanks. However, we do assume that given a set of alternatives to evaluate and a set of relevant attributes, the decisionmaker can evaluate 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, and provi-

ding information about the standing of his suggested alternatives on each indicated attribute.

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 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 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 ω . (Note that this means that they all have access to the same information.) 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 D with the property that for each $d \in D$ there is some $D \in D$ such that $d \in D$. (This models the possibility that when a party suggests one alternative, he cannot avoid calling another to mind.) 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 that alternative 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 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 II. 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 III does not apply directly to this setting. The reason is that an interested party cannot report fully about all the alternatives in order to convince the decisionmaker that a particular alternative is best. The decisionmaker must do some of the work on his own, discarding alternatives whose advocates do not justify them adequately.

Normally, in order to adopt an effective skeptical posture, it is necessary for the decisionmaker to know the interested parties' preferences (that is, the value of y). However, in the presence of sufficiently intense competition among the interested parties, it may be enough for the decisionmaker to be skeptical *about the alternatives themselves*, without concerning himself about how the interested parties' preferences color their reports. In our present model, a skeptical posture toward alternatives entails the decisionmaker believing that Z is the union of the A_i 's and then adopting beliefs about the attribute ranks of each suggested decision d which, while consistent with Z being the union of the A_i 's and with the information provided, gives the lowest possible expected utility of choosing 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 optimal (as in the previous section), then it will be in someone's interests to suggest it and provide supporting information.

Proposition 5. Suppose that for no ω and no $d \in \Delta$ is it true that d is preferred or indifferent to the full information decision at ω 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. (For one example, suppose 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 various products. Then the logic of Proposition 2 implies that the skeptical strategy is the only equilibrium strategy for the buyer.) So 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.

V. COMPARISON WITH MECHANISM DESIGN

The questions we have studied in this note are related to ones that have been studied in the burgeoning economic literature on "mechanism design." 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 utilized 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 agent's prior beliefs. 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 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 con-

trast, 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.

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 cannot commit himself to ignore information or to take any decision that is not optimal, given his information. 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 taking decisions which will not be in his interests *ex post*. Commitment, however, is a subtle matter. In reality 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. 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, so our theory applies even in some situations where the ability of the decisionmaker to control the rules and achieve commitment is in doubt.

Third, we allow the decisionmaker a far greater range of uncertainty than is common in the mechanism design literature. Our decisionmaker may not know what alternatives are available to him and may be forced to rely on interested parties for suggestions. He may not know what the relevant dimensions of a decision are. Indeed, he may not even be a Bayesian with subjective probabilities to describe his beliefs about these things. Uncertainty of this kind is an important aspect of reality, and it is a significant finding that competition among interested parties sometimes alleviates this kind of 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 in sight in order to define an objective function for the optimizing process.

Finally, the models used in this paper deal with verifiable information, in the same spirit as earlier work by Milgrom [1981] and Grossman [1981]. Research in mechanism design has mostly dealt with information about variables like taste, for which no direct verification is possible. Then, the decisionmaker must provide incentives to the interested parties to report their information truthfully. Both formulations are highly limited as models of reality, and both shed some light on the important intermediate case in which some information can be verified while other information must be coaxed from the interested parties.

VI. 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 "ideas" and information. Some parts of our analysis apply to the case of a buyer being courted by many sellers; other parts apply to hearings 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 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.

The nature of skepticism is illuminated to some small degree by our study. Informally, skepticism means that one assumes the interested parties' reports are tinged by self-interest. In our first model, skeptical behavior requires that the decisionmaker know the preferences of the interested party well, since a skeptical interpretation of a report in that model depends on what objective one thinks the interested party is pursuing. In our later models, we found that the decisionmaker sometimes has an effective strategy that entails skepticism about the alternatives themselves, without regard to the individual parties' preferences. Rational skepticism is a fascinating and subtle matter about which much remains to be learned.

REFERENCES

- Vincent Crawford and Joel Sobel [1982], "Strategic Information Transmission", *Econometrica* 50, 1431-1451.
- Sanford Grossman [1981], "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 461-483.
- Bengt Holmstrom [1977], *On Incentives and Control in Organizations* Ph.D. dissertation, Stanford University.
- David Kreps and Robert Wilson [1982], "Sequential Equilibria," *Econometrica* 50, 863-894.
- Paul Milgrom [1981], "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 380-391.
- Paul Milgrom and John Roberts [1982], "Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis," *Econometrica* 50, 443-459.
- Joel Sobel [1983], "A Theory of Credibility," mimeo, University of California - San Diego.

APPENDIX
PROOFS OF PROPOSITIONS 2 AND 3

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

$$(A1) \quad E\{u(x, b(r(x')) + q) / \lambda q + p(r(x'))\} < 0,$$

and the buyer could do better by reducing his purchases in response to the report $r(x')$, contradicting 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 , $b(A) = \min \{b(\{x\}) \mid x \in A\}$.

The result that $b(r(x)) = b(\{x\})$, together with the Seller Maximization equilibrium condition, implies 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,

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

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

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

which contradicts the Buyer Maximization condition. QED

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 paper), r^* must have the property that for all $x' \in r^*(x)$, $b(\{x\}) \leq b(\{x'\})$. (Otherwise, reporting $r^*(x) \setminus \{x\}$ would sell more to the unsophisticated buyer, 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 $Q < b(\{x\})$,

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

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 will have established that r sells no more to sophisticated buyers than r^* . So, 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 Consistent Beliefs and the strict concavity of buyer preferences, (A1) holds, violating Buyer Maximization. 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 to $r(x)$, and maximizes sales to naive buyers, contradicting 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. QED