12

Differential Information and Strategic Behavior in Economic Environments: A General Equilibrium Approach

Andrew Postlewaite
David Schmeidler

1. Introduction

The last three decades have seen the development of what is probably the most general model in economic theory—the Arrow-Debreu model of general equilibrium theory. Within this model it has been possible to prove the existence of equilibrium under very general conditions and to prove that the equilibrium allocations are Pareto efficient. Over the last decade, however, there have been a plethora of papers dealing with (sometimes) less general models in which either the existence or the efficiency of equilibria is called into question. Akerlof, in his now famous "Lemons" paper (1970), showed that in a model in which buyers and sellers have asymmetric information about the quality of a product, it is possible that the equilibrium outcome is inefficient.

Rothschild and Stiglitz (1976) examined a model of insurance contracts in which the potential purchasers were heterogeneous and had better information about their likely accident rates than did the sellers of the insurance. They defined an equilibrium, and showed that for large sets of environments no equilibrium could exist.

Kreps (1977) made a single change in the Arrow-Debreu model, namely, that one agent would know the state of nature and a second would not. He showed by example that there might not be an equilibrium when the "uninformed" agent used the information that might be contained in the price.

The common element of these papers is differential information; different agents have different information about the environment. Our purpose in this chapter is to present a general equilibrium model that makes explicit the information or beliefs that an agent has as part of his primitive characteristics. The model we present is a reinterpretation of Harsanyi's model of incomplete information games. Once we have presented the model, we will be able to apply the Harsanyi-Nash existence theorem to this model. Our purpose is essentially expository. Many models of asymmetric information that have been studied are subsumed in this framework.

This research was partially supported by NSF Grant No. SES-8026086. A number of people carefully read and commented on earlier drafts. We would particularly like to thank an anonymous referee, Beth Allen, Peter Hammond, Eric Maskin, David Sappington, and Jenny Wissink. The present chapter is introductory work to our joint paper "Implementation in differential economies," which appeared in *Journal of Economic Theory* 39 (1986):14–33.

In Section 2, before introducing incomplete information, we describe the notion of Nash equilibrium for a strategic outcome function in a general equilibrium environment. In Section 3 we extend this framework to include differential information. The differences from Harsanyi's approach and the motivation for these differences are explained. In Section 4 the revelation principle is defined, examined, and critiqued. We analyze the problems of welfare evaluation in differential information environments in Section 5.

2. The Complete Information Model

Formally, we consider pure exchange economies with a finite number of commodities and a finite number of agents. The consumption set for each agent is the nonnegative orthant of l-dimensional Euclidean space, R^l_+. The set of agents will be denoted by T. For notational simplification we identify T with $\{1,\ldots,n\}$. Each agent is characterized by a utility function, $u: R^l_+ \to R$, and an initial endowment, $w \in R^l_+$. Each utility function is assumed to be continuous, concave, and strictly increasing in each coordinate. Furthermore, whenever expected utility must be considered, we assume that the agents are characterized by von Neumann-Morgenstern utility functions. The set of all such utility functions is denoted by U. A finite pure exchange economy, e, is then a list $(w_t,u_t)_{t\in T}$, where $(w_t,u_t) \in R^l_+ \times U$.

The set of all economies will be denoted by $\tilde{E}$. Given an economy $e \in \tilde{E}$, an allocation for e is a list $(x_t)_{t\in T}$, where $x_t \in R^l_+$ and $\Sigma_{t\in T}(x_t - w_t) = 0$. As stated in the introduction, we are interested in the reallocation of initial endowments. We will model the rules or institutions that govern exchange using the concept of a strategic outcome function. Imagine a set $\tilde{S}$ that contains all conceivable acts that could be undertaken by some agent. An agent may be restricted by his initial endowment from undertaking certain acts in $\tilde{S}$. To represent this idea formally, we define a correspondence $S: R^l_+ \to \tilde{S}$, where $S(w)$ are the acts available to an agent whose initial endowment is w. To avoid notational complication, we define the strategic outcome function for a given n and a list $(w_t)_{t\in T}$ of initial endowments (and of course the given correspondence S above). The domain of this strategic outcome function f is $\underset{\sim}{S} \equiv \chi^n_{t=1} S_t$, where $S_t = S(w_t)$ $t = 1,\ldots,n$. The range of f is the set of allocations compatible with $(w_t)_{t\in T}$. We have explicitly assumed that two agents with the same initial endowments have the same acts available to them. Notice that we can introduce an asymmetry between agents in the definition of f rather than in its domain. We will elaborate on this point after the range of f is defined.

We call a strategic outcome function *applicable* to an economy e if it is parameterized by the initial endowments of e. If f is applicable to an economy e, then it is applicable to any other economy with the same number of agents and the same list of initial endowments. We will call a pair e and f compatible if f is applicable to e. We denote by $\underline{s}$ the list $(s_t)_{t\in T}$ and by $f_t(\underline{s})$ the bundle agent t receives in the allocation $f(\underline{s})$. Notice that $f_t(\underline{s})$ does not depend on the particular economy e so long as f is applicable to e. Specifically, $f_t(\cdot)$ is independent of the utility functions of the agents in the economy. The asymmetry mentioned above can be modeled by letting $f_t(\underline{s})$ be different from $f_{t'}(\underline{s})$ even if $w_t = w_{t'}$ and $s_t = s_{t'}$.

The dependence of the outcome on the acts $\underline{s}$ and not on the utility functions is meant to convey the partial decentralization of the exchange rules or institutions. We use the term *partial decentralization* since f is parameterized by $(w_t)_{t\in T}$. (Refer to Schmeidler [1980] and Hurwicz, Maskin, and Postlewaite [1982] for a more detailed treatment of this distinction.)

Suppose a set T of n agents with characteristics $(w_t,u_t)_{t\in T}$ is faced with institutional rules (including available acts) represented by the strategic outcome function f. Our basic goal is to describe the allocations that result from the institutional rules in this economy. Since e and f compatible formally define a game to be played by the n agents, we can restate the question of which allocation arises in the game as: which solution concept is appropriate for the game? The Nash equilibrium of this game, which is the focus of this chapter, embodies the notion of incentive compatibility. That is, each agent is assumed to choose acts governed by his own self-interest, where an agent's interests are represented solely by his utility function. Notice that an agent's interest in the allocations that others receive could be easily introduced in this model through the agent's utility function without altering the solution concept.

Formally, a list of strategies $\underline{s}^* = (s_t^*)_{t\in T} \in \underline{S}$ is a *Nash equilibrium* for (e,f) if $u_t(f_t(\underline{s}^*)) \geq u_t(f_t(\underline{s}^*|s_t))$ for all $s_t \in S_t$ and $t = 1,\ldots,n$. The resulting allocation $f(\underline{s}^*)$ is called the *Nash allocation*. We will not attempt to fully justify the use of Nash equilibrium as an appropriate solution concept for our problem. Rather, we will concentrate our attention on a specific criticism of its use for this problem. At first sight it would seem that for an agent to choose a best response to the actions of the others, he must know or be able to predict their choices. Since we wish to think of the game as being played without communication, justifications of Nash equilibria on the grounds that they can be interpreted as self-enforcing agreements are not applicable. An alternative justification has an agent computing the Nash equilibria and playing his part while assuming the other agents do likewise. We disregard, for a moment, the problem of multiplicity of Nash equilibria. There is an implicit assumption in this second justification, namely, that each agent knows the game he is playing. That is, each agent knows e as well as f. This assumption is the basis of the particular criticism that has motivated this chapter. Harsanyi (1967–68) introduced a notion of incomplete information games specifically to address this issue. We expand on this notion in what follows.

3. Incomplete Information and Harsanyi-Nash Equilibria

A. The Formal Model

Incomplete information in allocation models includes uncertainty about agents' characteristics, the characteristics of the commodities, and the rules of the game. A complete description of an economy should resolve these uncertainties. Following Savage (1954), a state of the world (state of nature) will include such a description. If the state of the world is known and the agents' choices of acts are known, each agent will know the outcome and his corresponding utility.

In contrast to the single-person decision problem of Savage, we consider n people. Thus, another aspect of incomplete information must be taken into account. Specifically, a complete description of a state of the world must contain the information not only for resolving the uncertainties mentioned but also for determining the extent to which each agent knows the state of the world. There must be some states of the world that are indistinguishable from others if there is incomplete information. The degree to which states are indistinguishable will affect agents' behavior and, therefore, must be part of the description of a state of the world. Including the possibility of agents' being able to differentiate among states allows us to analyze asymmetric information among agents. This is the essence of incomplete information beyond that treated by statistical decision theory.

We now explain how one can formally describe agents' information about agents' characteristics, commodities, and allocation rules. Let ω' be a complete description of agents' characteristics, commodities, and allocation rules, and let Ω' be the set of all possible ω'. For each agent $t \in T$, let $X_t\colon \Omega' \to R$ be a real-valued function. $X_t(\omega')$ is the "signal" observed by agent t if ω' occurs. Agent t can distinguish between ω_1' and ω_2' if $X_t(\omega_1') \neq X_t(\omega_2')$. A state of the world, ω, is $(\omega',(X_t(\omega'))_{t\in T})$ and Ω is the set of all such points. An economic agent is assumed to observe his signal before choosing an action. Mimicking Savage, we assume that each agent has a probability distribution (a prior) on Ω'. This, in conjunction with his signal, yields a posterior probability distribution on Ω'. We assume throughout this chapter that Ω' is finite. Given the prior on Ω', P_t', and the random variable X_t, agent t will have for each $\omega' \in \Omega'$ a conditional probability distribution on Ω', $P_t'(\cdot|X_t(\omega'))$.

For agent t, any $\omega \in \Omega$ is associated with a unique pair $(\omega', X_t(\omega')) \in \Omega' \times R$. Hence we can define $P_t\colon \Omega \times \Omega \to [0,1]$ by $P_t(\omega,\bar{\omega}) = P_t'(\omega'|X_t(\bar{\omega}'))$, where ω' and $\bar{\omega}'$ are the unique elements of Ω' associated with ω and $\bar{\omega}$, respectively. We will write $P_t(\omega,\bar{\omega})$ as $P_t(\omega|\bar{\omega})$ since we want to interpret P_t as the conditional probability that t assigns to ω when $\bar{\omega}$ is the true state.

For our purposes, the function $P_t(\cdot|\cdot)$ will describe the information available to agent t at the time of his decision. The origin or derivation is immaterial. We will take $P_t(\cdot|\cdot)$ as a primitive in the description of agent t. The generation of Ω and P_t from Ω', X_t, and P_t' is a special case of interest.

To help illustrate and motivate this structure, a common knowledge assumption on $(P_t(\cdot|\cdot))_{t\in T}$ will be helpful. If agent 1 knew $P_{t_2}(\cdot|\cdot)$ and $\omega \in \Omega$, we would also know that agent 2's probability distribution over Ω (beliefs) was $P_{t_2}(\cdot|\omega)$. But agent 1 typically does *not* know ω even when ω is the true state. $P_{t_1}(\cdot|\omega)$ is precisely our way of representing his beliefs about Ω when ω is the true state. Thus $P_{t_1}(\omega'|\omega)$ "is" the probability that agent 1 assigns to the possibility that 2's beliefs about Ω are $P_{t_2}(\cdot|\omega')$. The "is" is in quotation marks because there may exist $\omega'' \neq \omega'$ with $P_{t_1}(\omega''|\omega) > 0$ and $P_{t_2}(\cdot|\omega'') = P_{t_2}(\cdot|\omega')$. Thus the probability that agent 1 assigns to $P_{t_2}(\cdot|\omega')$ being agent 2's beliefs about Ω must be computed by summing $P_t(\omega'|\omega)$ over such ω'. In the same way we can generate agent 1's beliefs about agent 2's beliefs about agent 1's beliefs about Ω, and so on.

If $\omega \in \Omega$ contains a complete description of the state of the world, it must contain a description of the characteristics of any commodity that is not completely known to all agents. In this case, an agent's utility will depend not only on $x \in R_+^l$, but on ω as well. Thus, a utility function is $u\colon R_+^l \times \Omega \to R$. Similarly, uncertainty about the strategic outcome function can be represented by extending the domain of f from $\underset{\sim}{S} \equiv X_{t\in T}S_t$ to $\underset{\sim}{S} \times \Omega$. We illustrate these concepts using the special case in which Ω is induced by Ω' and $(X_t(\cdot))_{t\in T}$. Let $\Omega' = \Omega_1' \times \Omega_2' \times \Omega_3'$, where Ω_1' is that part of the description of a state of the world concerning the uncertainty about agents' utility functions. Specifically, $\Omega_1' \equiv X_{t\in T}\Omega_1'(t)$, where $\Omega_1'(t)$ relates to the uncertainty about agent t's utility function. A special case of interest is the one in which agent t's utility function does not depend on $\Omega_1'(\bar{t})$, $\bar{t} \neq t$. Ω_2' contains that part of the description of the state of the world which details the characteristics of commodities. Thus, u_t is generally defined on $R_+^l \times \Omega_1'(t) \times \Omega_2'$. Ω_3' is that part of the description of the state of the world concerning the uncertainty about the strategic outcome function. Thus, the strategic outcome function f is defined on $\underset{\sim}{S} \times \Omega_3'$.

This interpretation and factorization of Ω is for expository purposes. There are, of course, many other interpretations of Ω consistent with the formal presentation.

Given an economy $\{\Omega,(w_t,u_t,P_t(\cdot|\cdot))_{t\in T}\}$, a strategy for agent $t \in T$ is a mapping from Ω into S_t, the set of acts available to agent t. However, since an agent does not know precisely which $\omega \in \Omega$ has occurred when he must choose his act, the action chosen must be the same for all $\omega \in \Omega$ that are indistinguishable. That is, if $\omega_1, \omega_2 \in \Omega$ are such that $P_t(\cdot|\omega_1) = P_t(\cdot|\omega_2)$, then the action agent t chooses must be the same in states ω_1 and ω_2. Formally, a strategy for agent $t \in T$ is a mapping $\sigma_t\colon \Omega \to S(w_t)$ satisfying $\sigma_t(\omega_1) = \sigma_t(\omega_2)$ if $P_t(\cdot|\omega_1) = P_t(\cdot|\omega_2)$. If $P_t(\cdot|\omega_1) \neq P_t(\cdot|\omega_2)$, then it is reasonable to assume that $P_t(\omega_1|\omega_2) = P_t(\omega_2|\omega_1) = 0$ since in each case the agent knows that his beliefs are inconsistent with that state arising. If we assume, in addition, that $P_t(\omega|\omega) > 0$ for all $\omega \in \Omega$ (recall that Ω is finite), we see that the conditional probabilities induce a partition on Ω. If $\omega, \omega' \in \Omega$ and $P_t(\omega|\omega') > 0$ (this, of course, implies $P_t(\cdot|\omega) = P_t(\cdot|\omega')$), then ω and ω' are in the same element of the partition. We denote by Σ_t this partition. A strategy is then a function $\sigma_t\colon \Omega \to S(w_t)$ which is measurable with respect to Σ_t.

We will now define a *Harsanyi-Nash equilibrium,* also called a *Bayesian equilibrium* or an *incomplete information equilibrium* (IIE). A list of strategies $(\sigma_t^*)_{t\in T}$ is said to be an IIE if for all $h \in T$ and for all acts $s \in S(w_t)$ and for every $\omega' \in \Omega$: $\Sigma_{\omega\in\Omega}u_h(f_h((\sigma_h^*(\omega),[\sigma_t^*(\omega)]_{t\neq h}),\omega))\ P_h(\omega|\omega') \geq \Sigma_{\omega\in\Omega}u_h(f_h((s,[\sigma_t^*(\omega)]_{t\neq h}),\omega))\ P_h(\omega|\omega')$. This equilibrium concept is most plausible if the strategic outcome function f, the probability functions $(P_t(\cdot|\cdot)_{t\in T}$, the utility functions $(u_t)_{t\in T}$, and endowments $(w_t)_{t\in T}$ are common knowledge. Recall that $(u_t)_{t\in T}$ being common knowledge does not preclude uncertainty about agents' utilities since there may be parameters of the utility function that are not commonly known. The assumption that an agent's conditional probabilities, $P_t(\cdot|\omega)$, are either identical or singular is not necessary for the definition of strategies or for the definition of an IIE. Similarly, the assumption that $P(\omega|\omega) > 0$ $\omega \in \Omega$ is not necessary.

Under appropriate conditions on f and $(S_t)_{t\in T}$ (e.g., continuity of f and compactness of S_t, $t \in T$), existence of an IIE in our case of finite Ω and mixed acts, or equivalently mixed strategies, follows from an argument analogous to that of Harsanyi (1967–68) and Nash (1950, 1951).

B. Elaboration

We will now discuss the relationship between our presentation and that in Harsanyi (1967–68) where the concept of incomplete information games and their equilibria was introduced. Since Harsanyi was not specifically interested in economic applications, he did not develop an economic structure as we have presented here. In particular, he considered only the mapping of n-lists of strategies and states of nature into n-lists of utilities. Game theory typically analyzes how a specific game would or should be played. That is, it treats the game as a parameter of the problem. Welfare economics, on the other hand, recognizes that for some purposes, some parts of the description of the game are not parameters but rather are choice variables of the problem. Welfare economists are willing to compare the equilibria of *different* games according to some criterion. More generally, one part of welfare economics focuses on how we should select the game to be played from some larger class of possible games. In our framework, the space of acts *and* the strategic outcomes function relating combinations of acts to *physical* outcomes together represent the economic institutions and are the choice variables for some problems. Hurwicz (1972), it should be noted, pointed out the importance of distinguishing between the economic environment and the allocation mechanism in welfare applications of game theory.

The separation of environments and allocation mechanisms is also required when the same economic institutions are used at different times by possibly different information. In this situation we are not simply comparing the outcome of one game with that of another when we consider two different institutional structures (strategic outcome functions). Rather, there will be a set of games indexed by an institutional structure, one for each conceivable economic environment for which this institutional structure will be used. Within this context, one of the goals of welfare economics is to compare alternative institutional structures. For example, one institutional structure might be said to be preferred to another if, for every economic environment and every informational structure, it gives rise to a Pareto-superior allocation. A similar, but weaker, criterion would be to compare the average performance of two institutional structures over a range of economic environments. These comparisons obviously cannot be made unless the strategic outcome function is separated from the characteristics of the agents, in particular, the utility functions. This distinction is absent in Harsanyi's original treatment. Furthermore, this point of view leads us to a criticism of certain applications of the revelation principle, which will be discussed in Section 4.

Next we consider questions concerning the existence of equilibria and mixed acts. As we stated in the introduction, there are a number of models in which there are problems with existence of equilibria. Because we present a general model with existence of equilibria, we should reconcile this with the failure of existence in these other models. Our equilibrium concept is Nash equilibrium for incomplete information games (Bayesian equilibrium as it is sometimes called) as defined above. This equilibrium may involve an agent randomizing his acts, that is, the equilibrium may call for an agent, in some circumstances, to choose an act according to some probability distribution. This leads to several problems. Randomized choice of acts has been questioned in descriptive models on the grounds that "people do not really behave that way." To answer this criticism, it has been suggested that the concept of purification and ε-purification of mixed-strategy equilibria, which we describe below, be employed.

Aumann et al. (1983) showed that if individual agents have some "private information," then pure-strategy equilibria will exist. While their paper dealt with some very deep mathematical problems, the result we are referring to was simply that if an agent sees the outcome of a random variable that is "rich enough" and independent of everything other agents see, his act can be a function of his private information. This eliminates the need for an additional random variable. Since the purpose of this chapter is to examine environments with uncertainty, a restriction to assume enough private information to avoid randomized acts may not seem overly restrictive.

There is another more serious problem with randomized acts. One reason we are interested in the existence of equilibrium is that if for some set of parameters or environments there is no equilibrium, then there is a fundamental problem with the consistency of the model. If an equilibrium does not exist, the model cannot describe realistically an existing institution, nor can it be the basis for a proposed institution. If an equilibrium exists, it does not matter whether it is "pure" or "mixed" from this standpoint. There is a more subtle modeling problem, however. In our model, as in all models that use the Nash-equilibrium solution concept, there is an implicit assumption that the choice of strategy is irrevocable. The definition of Nash equilibrium states that no single agent will want to change his strategy unilaterally. However, we must be careful to distinguish between acts and strategies. In a game of complete information, a strategy is a choice of an act, possibly randomized. If the strategy involves a nondegenerate probability distribution over acts, it is the *strategy*, rather than the act, that an agent would not change unilaterally. In fact, for many games in which the equilibrium involves randomization, the *act* chosen is *not* a best response to the acts of others.

The subtle modeling question posed by the above discussions is whether we really believe that strategies (and consequent acts) are irrevocable. In some models, for example, the insurance model of Rothschild and Stiglitz mentioned above, the answer is no. The implicit strategy available to the firm is an insurance contract offered to an agent. If the contract is not a best response to the (now public) choice of other firms' contracts, the contract may be changed. Thus, for an equilibrium to be plausible, the acts themselves should be best responses to each other.

Formally, we could introduce a stronger equilibrium notion than IIE. We could ask that, in addition to the properties of an IIE, the acts chosen be best responses to each other for every ω in Ω. The existence of such an *ex post* equilibrium is, of course, not guaranteed by the existence theorem. In fact, there are many games in which there will *not* be an *ex post* equilibrium. The nonexistence problem of Rothschild and Stiglitz that was mentioned in the introduction is precisely this kind. In the next section we present an example in which the IIE has this *ex post* equilibrium property; however, the property is destroyed when we move to the revelation game.

4. The Revelation Principle

Given an incomplete information economy $\{\Omega,(w_t,u_t,P_t(\cdot|\cdot))_{t\in T}\}$, a strategic outcome function can be defined where the set of acts available to agent t, S_t, is $\{P_t(\cdot|\omega)|\omega \in \Omega\}$. That is, agent t announces one of the probability beliefs that could arise from some $\omega \in \Omega$. For each T-list of such announcements, a feasible allocation would be assigned. In essence, each agent is asked to reveal the information that is not already common knowledge. Such games are called revelation games. (There is a slight difference in the treatment of mixed acts which will be discussed below.) It is of interest to characterize the set of revelation games for which correct revelation is an IIE for each $\omega \in \Omega$.

REVELATION PRINCIPLE. Given an economy $\{\Omega,(w_t,u_t,P_t(\cdot|\cdot))_{t\in T}\}$, a strategic outcome function $f : \underset{\sim}{S} \to (R^l_+)^T$, and an IIE $(\sigma^*)_{t\in T}$, there exists a revelation game in which truthful revelation is an IIE. Furthermore, for this game, truthful revelation gives rise to an allocation identical to that arising from $(\sigma^*)_{t\in T}$ and f for every $\omega \in \Omega$.

Gibbard (1973) first used this principle in the context of complete information. For incomplete information this result first appeared in Rosenthal (1978) in the context of a bargaining model. Myerson (1979) stated it in full generality, and Harris and Townsend (1981) independently derived a version for economic environments, although not in the same manner as presented here. The proof of this principle is immediate. Denoting the strategic outcome function for the revelation game by f', we define $f'(P_t(\cdot|\omega)_{t\in T}) \equiv f((\sigma_t^*(\omega))_{t\in T})$ for all $\omega \in \Omega$. If for some $\omega \in \Omega$ and $t \in T$, $P_t(\cdot|\omega)$ were not a best response to the correct revelation of the other agents, then $\sigma_t^*(\omega)$ cannot be a best response to the strategies $(\sigma_{t'}^*)_{t'\neq t}$ of the others. This proof is similar to the proof in Gibbard (1973).

REMARK. If the IIE $(\sigma_t^*)_{t\in T}$ for the original strategic outcome function f assigns a mixed act for some agent t at $\omega \in \Omega$, the revelation game will assign a random outcome at some n-list of strategies of the revelation game, $[P_t(\cdot|\omega)]_{t\in T}$. Hence, the range of the strategic outcome function of the revelation game, f', may contain distributions over allocations even if the range of f did not. The translation from the original strategic outcome function, f, to the revelation strategic outcome function, f', has essentially transferred the randomness from acts to outcomes.

A problem with the revelation principle and revelation games involves the complexity of the strategies presumed to be available to the agents. There is a class of strategic outcome functions applied to pure exchange economies (see, for example, Dubey, Mas-Colell, and Shubik 1980; Postlewaite and Schmeidler 1978; or Schmeidler 1980) for which the acts available to an agent are contained in a subset of a finite-dimensional Euclidean space. Furthermore, the set of acts is independent of the number of agents in the economy. We can easily apply one of these institutions (i.e., strategic outcome functions) in an incomplete information environment. The strategy sets and the outcome function remain the same; only the individual agent's choice of an act is complicated by the incomplete information structure. Consider an n-person exchange economy with the agent's information described as follows. There is a common prior over n-types of individual characteristics. Each agent learns his own characteristics with certainty and observes other agents' characteristics with noise. Suppose that any neoclassical utility can be obtained as an observation. In this case the set of acts in the revelation game consists essentially of all n-tuples of neoclassical characteristics (assuming whatever independence is necessary). This example demonstrates that the complexity involved in the transmission or description of an act in the revelation game may be dramatically more complicated than in the "original" strategic outcome function. Although this example violates our assumption of finite Ω, it illustrates that the acts in the revelation game may be more, rather than less, complicated than in the original game.

Maskin (1977) investigated the performance possible from strategic outcome functions for social choice problems. He introduced individual strategy sets that were n-tuples of preference orderings, sometimes referred to as profiles. Maskin's positive results have been criticized on the grounds that the strategy spaces were unreasonably complex. This complexity is sometimes contrasted with the simplicity of revelation games, in which an agent "simply announces his type." (In fact, Maskin's game is itself a revelation game in which each agent's private information includes the true profile of preferences, a complete description of the economy.)

In Section 3.B, where we elaborated on the concept of an IIE, we mentioned that the property of *ex post* equilibrium may not be preserved if we substitute the revelation game for the original game. The following example illustrates the point.

EXAMPLE. $\Omega = \{\omega_1,\omega_2,\omega_3,\omega_4,\omega_5\}$, each with probability 1/5. There are two goods, x and y, and two agents, 1 and 2, whose utility functions are given below:

$$u_1(x,y;\omega_i) = \begin{cases} x + (1/2)y & i = 1,3 \\ x + 2y & i = 2,4 \\ 3x & i = 5 \end{cases} \quad u_2(x,y;\omega_i) = \begin{cases} (1/2)x + y & i = 1,2 \\ 2x + y & i = 3,4. \\ 3y & i = 5 \end{cases}$$

Initially agent 1 holds one unit of x and agent 2 holds one unit of y. The information sets for agent 1 are $F_1^1 = \{\omega_1,\omega_2\}$, $F_2^1 = \{\omega_3,\omega_4\}$, and $F_3^1 = \{\omega_5\}$. For agent 2 they are $F_1^2 = \{\omega_1,\omega_3\}$, $F_2^2 = \{\omega_2,\omega_4\}$, and $F_3^2 = \{\omega_5\}$. We will now define a strategic outcome function. The acts available to the agents are $A = \{a,b\}$. The outcome function $f: A \times A \to R_+^2$ is defined by $f_1(a,a) = (0,1)$, $f_2(a,a) = (1,0)$, and $f_1 = (1,0)$, $f_2 = (0,1)$ otherwise. The strategies $\sigma_i(F_j^i) = a$ and $\sigma_i(F_3^i) = b$ for $i,j = 1,2$ constitute a Nash equilibrium with net trades $(-1,1)$ and $(1,-1)$ to agents 1 and 2, respectively, in states ω_1, ω_2, ω_3, and ω_4 and net trades 0 in ω_5. Agent 1's conditional probability distribution is $p(\omega_1) = p(\omega_2) = 1/2$ if F_1^1 is the event, $p(\omega_3) = p(\omega_4) = 1/2$ if F_2^1 is the event, and $p(\omega_5) = 1$ if F_3^1 is the event. If agent 1 knows the equilibrium strategy of agent 2 and *also* the act chosen by agent 2, his probability beliefs do not change; the strategy and act chosen by agent 2 reveal no information to agent 1 beyond that which he already knows from his information set.

Consider now the revelation game that is induced from this equilibrium. Each agent is one of three types; thus $S_1 = \{F_1^1,F_2^1,F_3^1\}$ and $S_2 = \{F_1^2,F_2^2,F_3^2\}$. The outcome function $h: S_1 \times S_2 \to R_+^l$ is defined as $h_1(F_i^1,F_j^2) = (0,1)$, $h_2(F_i^1,F_j^2) = (1,0)$, unless $i = 3$ or $j = 3$. If $i = 3$ or $j = 3$, $h_1(F_i^1,F_j^2) = (1,0)$ and $h_2(F_i^1,F_j^2) = (0,1)$. This equilibrium, of course, gives rise to the same outcome as before for each state of nature ω_i. But, there is a difference in agent 1's probability beliefs if he knows the strategy and act of agent 2. These, together with his own information, reveal the state of nature perfectly. Thus in states 1 and 3, agent 1 would prefer to have chosen act F_3^1, since this would result in no trade, yielding higher utility than the outcome $(0,1)$ in states 1 and 3. Similarly, agent 2 would "regret" not having announced F_3^2 in the case that states 1 or 2 resulted, and, similarly, the state will be known with certainty if agent 2 knows the strategy and the act chosen by agent 1.

In summary, the original game had the property that not only were the strategies chosen by the agents best responses to each other, but so were the specific acts, regardless of the state of the world. This property is not preserved when we move to the induced revelation game. Information that may be revelant for some applications is lost.

If for modeling purposes a social planner desires *ex post* equilibrium, it may not be enough to consider only revelation games. The example illustrates a particular outcome which can result from an *ex post* equilibrium, but not of a revelation game.

Let us now turn to the problem of multiple equilibria. The meaning and plausibility of Nash equilibria are, in general, weakened if there are multiple equilibria. Multiple equilibria are not a problem specific to revelation outcome functions. In particular, if there exist multiple equilibria, a problem arises concerning how the revelation principle is used in practice. For example, many principal-agent problems are formulated as an outcome for each piece of information that the agent reveals. The principal's utility is maximized over all such schemes subject to the constraint that the agent(s) has (have) no incentive to lie. By applying the revelation principle, we are implicitly assuming that the solution to this problem is the best that can be achieved—even among nonrevelation schemes.

This process will yield an equilibrium outcome that maximizes the principal's expected utility. The mechanism that generates this best outcome may have more than one equilibrium. If there is more than one equilibrium, the "truthful revelation" equilibrium may not be compelling among all equilibria. The following example illustrates the problem.

EXAMPLE. Consider the following economy with four equally likely states of nature, five commodities, and three agents. The agents' utility functions in each state are as follows:

	ω_1	ω_2	ω_3	ω_4
Agent 1:	$x_1 + 2x_3 + x_5$	$x_1 + 2x_3 + x_5$	$2x_1 + x_3 + x_5$	$2x_1 + x_3 + x_5$
Agent 2:	$x_2 + 2x_4 + x_5$	$2x_2 + x_4 + x_5$	$x_2 + 2x_4 + x_5$	$2x_2 + x_4 + x_5$
Agent 3:	$x_3 + x_4 + 2x_5$	$x_2 + x_3 + 2x_5$	$x_1 + x_4 + 2x_5$	$x_1 + x_2 + 2x_5$.

The information structures and initial endowments for the agents are:

Agent 1:	$\{\omega_1,\omega_3\}\ \{\omega_2,\omega_4\}$	$w_1 = (0,0,0,0,1)$
Agent 2:	$\{\omega_1,\omega_2\}\ \{\omega_3,\omega_4\}$	$w_2 = (0,0,0,0,1)$
Agent 3:	$\{\omega_1,\omega_2,\omega_3,\omega_4\}$	$w_3 = (1,1,1,1,0)$.

Suppose we look for the strategic outcome function that maximizes the expected utility of agent 3 subject to the constraint that agents 1 and 2 are at least as well off in any state as they would be if they consumed their initial endowments. The state contingent allocations that maximize agent 3's utility subject to the constraints are given as follows:

	ω_1	ω_2	ω_3	ω_4
Agent 1:	(1,0,0,0,0)	(1,0,0,0,0)	(0,0,1,0,0)	(0,0,1,0,0)
Agent 2:	(0,1,0,0,0)	(0,0,0,1,0)	(0,1,0,0,0)	(0,0,0,1,0)
Agent 3:	(0,0,1,1,2)	(0,1,1,0,2)	(1,0,0,1,2)	(1,1,0,0,2)

If we consider the revelation mechanism that gives these allocations for these states, we see that "truth" is an equilibrium. If agent 2 truthfully announces whether event $\{\omega_1,\omega_2\}$ or event $\{\omega_3,\omega_4\}$ obtains, agent 1 has the ability by his announcement to "confuse" states 1 and 2 and states 3 and 4. That is, if ω_1 is the true state, agent 2 will announce $\{\omega_1,\omega_2\}$, and agent 1 effectively announces which of $\{\omega_1,\omega_2\}$ is the state by announcing $\{\omega_1,\omega_3\}$ or $\{\omega_2,\omega_4\}$. Since agent 1 receives the same allocation in states 1 and 2, and similarly gets identical bundles in states 3 and 4, he has no incentive to lie. Agent 2 can "confuse" states 1 and 3 and states 2 and 4 in a similar manner if agent 1 is announcing truthfully. But since his bundles are constant across the states, he can mix (or confuse); he also has no incentive to lie. Thus, truthful revelation is an equilibrium and gives rise to utility levels of 1, 1, and 4 to agents 1, 2, and 3, respectively, regardless of which state arises.

The point the example illustrates is that although there is no incentive for the agents to lie, there is also no disincentive for doing so. If agent 1 announces the "wrong" state, his utility is still 1, given that agent 2 announces truthfully. Moreover, if agent 2 announces the wrong state, agent 1's utility will be 2 rather than 1. Thus, although truthful revelation is an equilibrium, this strategy is weakly dominated by announcing incorrectly. Incorrect revelation, in fact, is a second equilibrium, an equilibrium that gives both agent 1 and agent 2 higher utilities than the truthful revelation equilibrium. It is plausible that the incorrect revelation equilibrium, rather than the truthful revelation equilibrium, may be the one that arises.

This example is motivated by examples in the auction literature. The fact that incorrect revelation weakly dominates truthful revelation is not robust in this example. Small changes in the utility functions for agents 1 and 2 could make unilateral deviations from truth yield payoffs worse than those arising from truthful revelation. However, examples do exist in which there is a second equilibrium which Pareto dominates truthful revelation for those who deviate from the truth that are robust to perturbations in the utility functions and endowments. In situations where one of the agents in the problem chooses a mechanism that maximally exploits the other agents, it may not be surprising that there could exist nontruthful equilibria that are better for the exploited agents. Consider, for example, problems of the "single principal—many agent" type. Many of these problems have been analyzed by maximizing the principal's expected utility over revelation games, subject to the constraint that the agents have no incentive to lie. This example points out that multiple equilibria can pose particularly difficult problems when using this approach. We note that this problem does not arise in the standard "single principal—single agent" problem as it is generally posed. Since there is only one agent who chooses strategically once the mechanism is chosen, there cannot be nontrivial multiple equilibria.

5. Welfare Evaluation in Differential Information Environments

In this section we discuss the concept of Pareto efficiency. We are interested in the question of how individual agents can compare two allocations or two institutional structures or, in general, any two entities. As is traditionally done in welfare economics, the comparison is assumed to be on an individualistic basis. Questions of equity, justice, and similar concerns are excluded. The desirability of making such comparisons stems from the presumption that society's members facing two allocations or two institutions will unanimously prefer the allocation or institution that Pareto dominates the other. Hence, we want to concentrate our attention on those allocations or institutions that are undominated.

To make the above remarks more precise, we must specify how an individual makes comparisons within our economic-information structure. If we are considering a situation involving a single environment with no uncertainty, the manner in which an individual compares alternatives (allocations) is unambiguous. When we add uncertainty, the simplest case is one in which the uncertainty is identical across agents and in which there are no contingent markets. In this case it has been customary to distinguish between *ex ante* and *ex post* efficiency (see, for example, Holmstrom and Myerson 1983 for a detailed discussion of this issue). In the case of incomplete information, the situation is much more complex. In the above case there is possibly incomplete information about an agent's characteristics. Thus, there may be uncertainty as to this agent's ranking of alternatives. The question arises whether a comparison of alternatives should be pointwise across his possible characteristics or in "expectation" across his characteristics. This latter point can be clarified in terms of the timing of the welfare evaluation with respect to the possible stages of information. Let us now illustrate and separate the various stages of decision making and (possible) acquisition of information in our model.

At this point we will make a brief digression. Recall that the strategic outcome function has as its arguments the agents' strategies, that is, acts that depend on the state of nature. Thus, in evaluating a strategy within a mechanism, an agent should realize that he will consume different bundles in different states. We are, therefore, led to consider an agent's expected utility of the contingent bundle $x_t(\omega)$. We do not imply, by this notation, that contingent markets exist. Indeed, we separate the evaluation of a contingent bundle, which is clearly meaningful, from the possibility of achieving a contingent bundle, which is a question of implementation. One attractive feature of this model is that it allows for an endogenous determination of which contingent bundles are possible. Furthermore, not only is it formally possible for agent t to calculate the expected utility of a bundle $x_t(\omega)$, even if his private information does not distinguish which ω has arisen, but such contingent allocations can arise in equilibrium. For example, consider the following differential information economy with two agents, two commodities, and two states of nature.

Agent 1 has endowment (1,0), information partition $\{\omega_1,\omega_2\}$, and utility function

$$u_1(x_1,x_2,\omega_i) = \begin{cases} x_1 \text{ if } i = 1 \\ x_2 \text{ if } i = 2. \end{cases}$$

Agent 2 has endowment (0,1), information partition $\{\omega_1\}$ $\{\omega_2\}$, and utility function

$$u_2(x_1,x_2,\omega_i) = \begin{cases} x_2 \text{ if } i = 1 \\ x_1 \text{ if } i = 2. \end{cases}$$

Consider the revelation game with an outcome function that gives final bundles of (1,0) and (0,1) to agents 1 and 2, respectively, in state one and the reverse in state two. Truth is an equilibrium that results in an allocation to agent 1 given by $x(\omega_1) = (1,0)$ and $x(\omega_2) = (0,1)$, which is contingent on the state of nature despite the fact that he cannot distinguish the states. The bulk of the principal-agent literature can be interpreted as an attempt to design mechanisms to achieve trades contingent upon states of nature that are observable only to the agent.

We now return to our main concern—welfare evaluation of uncertain bundles. In the beginning, say stage one, the agents have not yet observed their private information. They may or may not have a common prior on Ω. The economy is common knowledge. In particular, every agent knows the type both he and other agents may be, that is, the possible information structures that may arise are known. For some problems this stage may be hypothetical in the sense that there may be agents for whom their type was never unknown to them. Part of Harsanyi's contribution in this area was to introduce this (possibly hypothetical) stage to consistently model agents with different information. The strategic outcome function is common knowledge at this stage. Given an economy, an n-list of strategies, and a strategic outcome function, an outcome will arise that depends on ω. Individual t' can calculate his expected utility, $\Sigma_{\omega\in\Omega}P_{t'}(\omega)u_{t'}(x_{t'}(\omega),\omega)$, where $P_{t'}(\omega)$ is the unconditional probability beliefs of agent t'. It is reasonable to assume that the prior probabilities, $(P_t(\cdot))_{t\in T}$, are common knowledge. (Alternatively, we can think of them as being known by an ethical observer.) We do not assume that all probabilities $(P_t(\cdot))_{t\in T}$ are identical. It may be that the probabilities are based on different models of the world and the agents agree to disagree. In the case in which agents agree to disagree, there is a fundamental problem

in selecting mechanisms or outcomes. Should we use inconsistent beliefs or not? It may be that each agent, on the basis of his own beliefs, prefers alternative A to alternative B, and yet there does not exist a single distribution that would support the individuals' orderings simultaneously. Since we have nothing to add here, for the remainder of this section we restrict our attention to the case in which the priors, $(P_t(\cdot))_{t\in T}$, exist and are identical for all agents. Note that a common prior, that is, $P_t(\cdot)$ independent of t, does not make $P_t(\cdot|\cdot)$ independent of t; the agents may have different partitions of Ω.

Even in the case of a common prior, there is a question concerning the reasonableness, or applicability, of the welfare evaluation involving averaging utilities $u_t(x_t,\omega)$. In the previous section we differentiated between uncertainty pertaining to the neoclassical preferences of the agent and uncertainty about a characteristic of the commodity or the institution. In the first instance, aggregation across different neoclassical preferences is tantamount to making interpersonal comparisons of utility. Any objection to interpersonal comparison of utility applies equally to the use of Pareto efficiency in stage one. If the uncertainty concerns a commodity or the institution, however, we have fixed preferences and there is generally no problem raised by the use of expected utility. We should note that in modeling a particular problem involving uncertainty, it may be unclear whether the uncertainty pertains to the preferences or the commodity. The plausibility of the modeling of uncertainty as pertaining to preferences or to commodities cannot be determined formally, but must be decided on in the context of the problem being analyzed.

Next we consider welfare evaluation at stage two, that is, after private information has been revealed and the conditional probability beliefs of agent t, $P_t(\cdot|\omega)$, are known by agent t. At this stage agent t will prefer $x_t(\cdot)$ to $y_t(\cdot)$ if for every $\omega \in \Omega$, $\Sigma_{\omega'\in\Omega}P_t(\omega'|\omega)u_t(x_t(\omega'),\omega') \geqq \Sigma_{\omega'\in\Omega}P_t(\omega'|\omega)u_t(y_t(\omega'),\omega')$ with strong inequality for at least one $\omega \in \Omega$. If at this stage an agent does not know his neoclassical preferences with certainty, the objections mentioned above apply again. However, it is assumed here, by definition, that no uncertainty as to the neoclassical preferences remains. Without this assumption, the plausibility of an Incomplete Information Equilibrium (Harsanyi-Nash-Bayes Equilibrium) is in doubt. Having eliminated uncertainty regarding agents' own preferences, we have eliminated the problems associated with the use of expected utility in agents' evaluations of alternatives. Nevertheless, the importance of the first-stage efficiency derives from the following self-evident observation: if an n-list of strategies results in an allocation that is undominated in the first-stage efficiency sense, it is undominated in the second-stage efficiency sense. The converse, however, is not true. It is straightforward to construct examples for which the allocation is pointwise Pareto efficient (stage two) but dominated by other allocations when the expectation is taken before private information is revealed. This should not be surprising, of course. Taking expected utility at stage one allows for the possibility of "trading off" utility, given one type of private information, which may be revealed for utility, given another type of private information. If stage one is not hypothetical and there is a point at which agents can make decisions prior to the revelation of their private information, first-stage efficiency may well be the proper concept for comparisons.

In the case in which first-stage comparisons are not meaningful, we can consider evaluation at stage two, after individual agents have their private information. There is, however, a basic problem with the Pareto-efficiency concept resulting from this stage-two welfare evaluation. Heuristically, the differential information entering into the individuals' evaluations of random outcomes is inconsistent with the notion of unanimity, which we expect from Pareto efficiency. Intuitively we expect Pareto efficiency to be characterized by an unwillingness of agents to recontract. The following example illustrates the problem.

EXAMPLE. Consider a differential information economy with two commodities, three equally probable states $\{\omega_1,\omega_2,\omega_3\}$, and two agents, 1 and 2. The information structures are $\{\omega_1,\omega_2\}$ $\{\omega_3\}$, for agent 1 and $\{\omega_1\}$ $\{\omega_2,\omega_3\}$ for agent 2; the initial endowments, independent of the state of nature, are (2,0) and (0,2), respectively.

The agents' utility functions $u_i(x,y;\ \omega_j)$ $i = 1, 2$ $j = 1, 2, 3$ are given as

	ω_1	ω_2	ω_3
$u_1(x,y;\ \omega)$:	$x + 3y$	$2x + y$	$4x + y$
$u_2(x,y;\ \omega)$:	$x + 4y$	$x + 2y$	$3x + y$.

The initial endowment is easily seen to be the unique symmetric Pareto-efficient allocation (in the stage-one sense), which, of course, is Pareto efficient for each ω_i. Consider the situation if ω_2 arises. The conditional beliefs of agent 1 will be $p(\omega_1) = p(\omega_2) = 1/2$. Similarly, the beliefs of agent 2 will be $p(\omega_2) = p(\omega_3) = 1/2$. The expected utility of the initial endowment given the private information, conditioned on ω_2, is 3 for each agent. The expected utility of (0,2) to agent 1 given his information is 4, and, similarly, (2,0) yields expected utility 4 to agent 2 given his information. Clearly, each agent would prefer trade to no trade here.

However, this individual willingness to trade may not lead to realized trade. Notice that agent 1's willingness to trade stems from his inability to distinguish ω_1 from ω_2. The trade is beneficial in ω_1, detrimental in ω_2, and beneficial in expected value. But agent 2 will not be willing to trade in ω_1. Thus, agent 2's willingness to trade conveys additional information for agent 1, information which when used makes agent 1 unwilling to trade.

We will now define a concept of Pareto efficiency when stage-one comparisons are not applicable. Given a partition $\tilde{B} = (B_1,\ldots,B_k)$ of Ω, an allocation $(y_t(\omega))_{t\in T}$ is said to $\tilde{B}$-dominate $(x_t(\omega))_{t\in T}$ if for all $t \in T$ and $B \in \tilde{B} : \Sigma_{\omega\in B}P(\omega|B)[u_t(y_t(\omega)) - u_t(x_t(\omega))] \geqq 0$ with at least one strict inequality.

An allocation $(x_t(\omega))_{t\in T}$ is said to be $\tilde{B}$-Pareto efficient if it is not $\tilde{B}$-dominated by any other allocation. This definition asks only that an allocation $(x_t(\omega))_{t\in T}$ be efficient within the events of the partition without requiring efficiency across events of the partition. Clearly, if an allocation is efficient with respect to a partition, it is efficient with respect to any finer partition, but not conversely. Efficiency with respect to the coarsest partition, $\tilde{B} = (\Omega)$, is our stage-one efficiency. Generally, agents should desire efficiency with respect to the coarsest partition that is sensible. As mentioned earlier, comparisons across some events may be tantamount to interpersonal comparison of utility, hence undesirable. Second, choices may be made after agents know certain events are impossible, hence gains in efficiency across such events may not be applicable.

In the definition of efficiency with respect to a partition, every agent's utility is calculated with respect to the same partition. This avoids the problem in the above example in which evaluations were made over different partitions. If an allocation is efficient with respect to a partition, the agents will not wish to recontract if they are informed of an event contained in some element of the partition.

If the only uncertainty in the economy involves the agents' neoclassical preferences, and if our assumption that an agent's private information contains at least his own preferences, the coarsest partition with respect to which efficiency can be evaluated is the joint information partition. This follows from the fact that for any two elements in this partition at least one agent has different neoclassical preferences in the two events by assumption. Hence, any coarser partition involves interpersonal comparison of utility. If there is uncertainty involving some aspect other than neoclassical preferences, there may be coarser partitions that avoid interpersonal comparisons of utility. The efficiency concept for a differential information environment should thus be with respect to some partition at least as coarse as the joint information partition.

This entire discussion is motivated by a desire to define efficiency properties of performance correspondences (social choice correspondences) independently of the possibility of implementing them. We wish to define our criteria independently of the possibility of satisfying them.

References

Akerlof, G. 1970. The market for lemons. *Quarterly Journal of Economics* 84:488–500.
Aumann, R., Y. Katznelson, R. Radner, R. Rosenthal, and B. Weiss. 1983. Approximate purification of mixed strategies. *Mathematics of Operations Research* 8:327–341, and errata 9:317.
Dubey, P., A. Mas-Colell, and M. Shubik. 1980. Efficiency properties of strategic market games: An axiomatic approach. *Journal of Economic Theory* 22:339–362.
Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41:587–602.
Harris, M., and R. Townsend. 1981. Resource allocation under asymmetric information. *Econometrica* 49:33–64.
Harsanyi, J. 1967–68. Games with incomplete information played by Bayesian players. *Management Science* 14:159–182, 320–334, 486–502.
Holmstrom, B., and R. Myerson. 1983. Efficient and durable decision rules with incomplete information. *Econometrica* 51:1799–1819.
Hurwicz, L. 1972. On informationally decentralized systems. In *Decision and organization*, ed. C. B. McGuire and R. Radner, 297–336. Amsterdam: North-Holland Publishing Co.
Hurwicz, L., E. Maskin, and A. Postlewaite. 1982. Feasible implementation of social choice correspondences by Nash equilibria. Mimeo.
Kreps, D. 1977. A note on "fulfilled expectations" equilibria. *Journal of Economic Theory* 14:32–43.
Maskin, E. 1977. Nash equilibria and welfare optimality. Massachusetts Institute of Technology. Mimeo.
Myerson, R. 1979. Incentive compatibility and the bargaining problem. *Econometrica* 47:61–74.
Nash, J. 1950. Equilibrium points of *n*-person games. *Proceedings of the National Academy of Science*, USA 36:48–49.
———. 1951. Non-cooperative games. *Annals of Mathematics* 54:286–295.
Postlewaite, A., and D. Schmeidler. 1978. Approximate efficiency of non-Walrasian Nash equilibria. *Econometrica* 46:127–135.
Rosenthal, R. 1978. Arbitration of two-party disputes under uncertainty. *Review of Economic Studies* 45:595–604.
Rothschild, M., and J. Stiglitz. 1976. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics* 90:629–650.
Savage, J. 1954. *Foundations of statistics*. 2d ed. 1967. New York: John Wiley and Sons.
Schmeidler, D. 1980. Walrasian analysis via strategic outcome functions. *Econometrica* 48:1585–1593.

13

Incentives, Information, and Iterative Planning

John Roberts

1. Introduction

In his 1972 Ely Lecture, Leonid Hurwicz identified the "proper integration of the information and incentive aspects of resource allocation models [as] perhaps the major unsolved problem in the theory of mechanism design" (1972a, 27). In the intervening years, a rich literature has developed on the incentive problem, and the information issue of the size of the message needed to achieve satisfactory performance has also received significant attention. Still, the desired integration hardly can be considered to have been achieved. Indeed, with the exception of some recent work (Green 1982; Green and Laffont 1987; Reichelstein 1984a, 1984b), there has been little formal analysis of the interplay between the demands of incentive compatibility and the desirability of economizing on the number and size of the messages sent. Moreover, as Green (1982) has argued, recent tendencies to view the revelation principle as a basis for mechanism design are a retreat from the realization of an objective of finding mechanisms that not only obtain efficient outcomes under self-interested behavior, but do so in an efficient manner.

One important line of work where the issue of economizing on the costs of information transfer has been fundamental is the study of iterative planning procedures (see Heal 1973; Tulkens 1978; Green and Laffont 1979; and, especially, Rochet for broad-ranging discussions or explicit surveys of this field). The basic idea underlying these procedures is that instead of transmitting a single, huge message (such as a utility function), the participants send many small messages (such as marginal rates of substitution).

I would like to thank Paul Champsaur, Joel Demski, Peter Hammond, participants in seminars at UC Berkeley, USC, and Stanford's economics department and business school, and especially the participants in the Conference on Informational Decentralization and Demand Theory for their comments and suggestions. This work was supported by funds from CORE, the Stanford GSB, and the National Science Foundation (grants SOC-7907542, SES-8108226, and SES-8308723). I am grateful to each of these institutions for their support and to CORE for its hospitality during my visit in May 1982 when much of the research reported here was done.