

DISTRIBUTIONAL STRATEGIES FOR GAMES WITH INCOMPLETE INFORMATION*†

PAUL R. MILGROM‡ AND ROBERT J. WEBER§

We study games with incomplete information from a point of view which emphasizes the empirical predictions arising from game-theoretic models. Using the notion of "distributional" strategies, we prove four main theorems: (i) a mixed-strategy Nash equilibrium existence theorem, (ii) a pure-strategy equilibrium existence theorem, (iii) a pure-strategy ϵ -equilibrium existence theorem, and (iv) a theorem describing how the set of equilibria of a game varies with the parameters of the game.

1. Introduction. In 1961, William Vickrey introduced games with incomplete information into the mainstream of economic theory in a study of competitive bidding [30]. The variety of applications of these games has widened considerably over the years (as evidenced by [6], [13], [18], [21]–[23], [25], [26], [31]), and many contributions have been made to the underlying theory of information [15]–[17], [20].

Despite these advances, the most fundamental questions which arise in applications remain unanswered. Do Nash equilibria exist for general games with incomplete information? When will such games have equilibria in pure strategies? How sensitive are equilibrium outcomes to modeling assumptions? For example, can a small variation in the assumed information structure lead to a large change in the equilibrium strategies?

Our aim is to provide partial answers to all of these questions, for one-stage simultaneous-move games. We prove an equilibrium existence theorem for a broad class of these games. We also prove that, with the appropriate concepts of closeness for information structures, payoff functions, and strategies, the correspondence that maps the specifications of a game into its set of Nash equilibria is upper-hemicontinuous. The ideas underlying this continuity theorem have been used elsewhere [13], [21] to simplify and unify the solutions of certain bidding games and to gain insights into the nature of the equilibria of these games; here we use them to study a game of timing that can serve as a model of competition between two animals, or of strike behavior. For games in which the players' informational variables have atomless distributions, we show that each player's set of pure strategies is dense in his complete set of strategies. For such games, mixed strategies are empirically indistinguishable from pure strategies and so the common objection that "one never observes people adopting mixed strategies" has no force. Finally, we identify a large class of games in which every mixed strategy equilibrium has a "purification", i.e., a pure-strategy equilibrium at which each player has the same expected payoff and the same distribution of

* Received October 13, 1980; revised August 3, 1984.

AMS 1980 subject classification. Primary: 90D10.

IAOR 1973 subject classification. Main: Games. Cross references: Bidding.

OR/MS Index 1978 subject classification. Primary: 238 Games/group decisions/noncooperative.

Key words. Incomplete information, Nash equilibrium, distributional strategies.

† This research was supported in part by U.S. Office of Naval Research contracts ONR-N000-14-77-C-0518 and ONR-N000-14-79-C-0685, National Science Foundation grants SOC-77-06000-A1 and SES-80-01932, and the Center for Advanced Studies in Managerial Economics at Northwestern University.

‡ Yale University.

§ Northwestern University.

observable behavior as at the mixed-strategy equilibrium in each of his informational states.

The games which we study can be described as follows. Each player i observes an informational variable (or *type*) t_i whose values lie in some complete, separable metric space T_i . After observing his type, player i selects an *action* a_i from some compact metric space A_i of feasible actions.

To accommodate a large variety of applications, we allow each player's payoff to depend on the actions chosen by all the players, on all the players' types, and also on some environmental variable¹ t_0 chosen by Nature from the set T_0 . We designate player i 's payoff by $U_i = U_i(t_0, \dots, t_n, a_1, \dots, a_n)$, where n is the number of players. Within this formulation, we define an *information structure* η for the game to be a joint probability distribution on $T_0 \times \dots \times T_n$ (where the measurable structure on each T_i is its collection of Borel subsets).

The conventional analysis of games involves three types of strategies: pure, mixed, and behavioral. A *pure strategy* is a measurable function $p_i: T_i \rightarrow A_i$. This has the interpretation that when player i learns that his type is t_i , he selects the action $p_i(t_i)$. Aumann [2] has observed that to define a mixed strategy properly (when T_i is "large") a randomizing device must be introduced for each player. Thus, let \bar{s}_i be uniformly distributed on $[0, 1]$. A *mixed strategy* for player i is a measurable function $\sigma_i: [0, 1] \times T_i \rightarrow A_i$. The interpretation is that when player i observes his type t_i and his randomizing variable s_i , he selects the action $\sigma_i(s_i, t_i)$. Let A_i be the collection of Borel subsets of A_i . A *behavioral strategy* is a function $\beta_i: A_i \times T_i \rightarrow [0, 1]$ with these two properties: (i) For every $B \in A_i$, the function $\beta_i(B, \cdot): T_i \rightarrow [0, 1]$ is measurable. (ii) For every $t_i \in T_i$, the function $\beta_i(\cdot, t_i): A_i \rightarrow [0, 1]$ is a probability measure. The interpretation of a behavioral strategy is that when player i observes t_i , he selects an action in A_i according to the measure $\beta_i(\cdot, t_i)$.

These conventional characterizations of strategies are not well suited to our purposes. Instead, we define a *distributional strategy* for player i to be a joint probability distribution on $T_i \times A_i$ for which the marginal distribution on T_i is the one specified by the information structure. We shall later show that distributional strategies are simply another way of representing mixed and/or behavioral strategies. While no meaningfully distinct strategies are added or deleted through this representation, the representation is more convenient for studying the relationship between the data of a game and the game's equilibria.

The remainder of this paper is organized as follows. In §2, we present an example to illustrate the ideas of convergence and purification. §3 contains a formal description of our model and presents the Existence Theorem, which provides sufficient conditions for the existence of an equilibrium in distributional strategies. The Convergence Theorem, which is developed in §4, asserts that when the data that specify a game are varied "continuously", the set of equilibria varies upper-hemicontinuously.

¹It is always possible to reduce a general game in which the payoffs may depend on t_0 (as well as on all the players' types and all the players' actions) to a canonical game where the payoffs do not depend on t_0 . Given a game with payoff function U_i for player i , the payoff function V_i in the canonical game is obtained by integrating out t_0 , as follows:

$$V_i(t_1, \dots, t_n, a_1, \dots, a_n) = E[U_i(t_0, \dots, t_n, a_1, \dots, a_n) | t_1, \dots, t_n].$$

This canonical form is the one studied by Harsanyi [15]. It is usually the appropriate form for the comparison of theoretical results, since it eliminates from consideration such spurious generality as the inclusion of chance events about which no information is available.

For applications, however, it is sometimes convenient to include t_0 in the explicit formulation of the game. For example, the assumptions in [25] are most easily stated and the economic insights are most easily explained when t_0 is included in the formulation. In [26], a major portion of the analysis is devoted to studying the effects of variations in the players' information about the state variable t_0 .

We take up the matter of pure-strategy equilibria in §5. Two theorems and a corollary summarize our results. Theorem 3 (the Denseness Theorem) states that a player's set of pure strategies is dense in his set of distributional strategies whenever his type has an atomless distribution. A corollary asserts the existence of approximate equilibria in pure strategies. To study exact equilibria, we need the following definition. A mixed strategy δ is said to have a *purification* if there is a pure strategy σ with these three properties: (i) σ is a best reply to the opposing strategies whenever δ is, (ii) for whatever strategies the other players may adopt, it is always true that the expected payoff of each type of each opposing player is the same against σ as against δ , and (iii) for whatever strategies the other players may adopt, the joint distribution of the state t_0 and all the players' actions is the same whether σ or δ is played. Theorem 4 (the Purification Theorem) gives sufficient conditions to ensure that every mixed strategy has a purification. When the conditions are satisfied, every mixed-strategy equilibrium corresponds to some pure-strategy equilibrium in which each player faces the same decision problem (at equilibrium) and earns the same expected payoff as in the mixed-strategy equilibrium. §6 indicates how the assumption that the type spaces are metric can be relaxed and how the case of "inconsistent beliefs" can be treated.

2. **An example: The war of attrition.** In this section, we present an example which illustrates the notions of "closeness" of information and payoff structures, and of strategies in (possibly different) games with incomplete information, notions which are central to the results of this paper. The game we analyze is known as the "War of Attrition",² and has been used to study conflict among animals [6]. It can also serve to model other conflicts, such as labor-management disputes involving strikes.

In the animal-conflict interpretation, two animals face one another in competition for a valuable prize, such as food or the opportunity to mate. The competition is nonviolent (for example, the animals engage in a display ritual); the cost of competition is proportional to the time spent in conflict (time which alternatively could be employed in the search for a comparable prize elsewhere, time during which both competitors are vulnerable to attack from predators). Let s_i ($i = 1, 2$) be the greatest length of time animal i would be willing to compete for the prize, if he were certain to receive the prize at the end of that time. We take the payoff functions of the game to be of the form

$$U_i(s_1, s_2, a_1, a_2) = \begin{cases} s_i - a_j & \text{if } a_i > a_j, \\ -a_i & \text{otherwise,} \end{cases}$$

where $(i, j) = (1, 2)$, and a_i and a_j are the lengths of time for which the two animals choose to compete. Suppose that \bar{s}_1 and \bar{s}_2 are independent, identically distributed random variables with common distribution F . It has been observed by several authors that, when F is concentrated at a point \bar{v} , so that the prize has commonly-known value \bar{v} to both competitors, the game has a unique symmetric Nash equilibrium point; the strategy of each at equilibrium is to choose the time at which he will concede according to the exponential distribution with mean \bar{v} . On the other hand, if F is a continuous distribution with density $f \equiv F'$, it is known that there is a unique symmetric Nash equilibrium point, in which each uses a pure strategy which calls for competition until time $b(s_i)$, where

$$b(s_i) = \int_0^{s_i} \frac{tf(t)}{1-F(t)} dt.$$

²This game is also known by various other names, including "Both-Pay Auction" and "Dollar Auction."

How do the equilibria in these two cases compare? Let G be the distribution of quitting times in the second case; then

$$G(a) = \Pr(\sigma(\bar{s}) < a) = \Pr\left(\int_0^{\bar{s}} \frac{f(s)}{1-F(s)} ds < a\right).$$

Assume that F concentrates its mass on a neighborhood $(\bar{v} - \epsilon, \bar{v} + \epsilon)$. Then the integrand in the expression above is nonzero only if $s < \bar{v} + \epsilon$, and so

$$\begin{aligned} G(a) &> \Pr\left((\bar{v} + \epsilon) \int_0^{\bar{s}} \frac{f(s)}{1-F(s)} ds < a\right) \\ &= \Pr(F(\bar{s}) < 1 - \exp[-a/(\bar{v} + \epsilon)]) \\ &= 1 - \exp[-a/(\bar{v} + \epsilon)], \end{aligned}$$

since the random variable $F(\bar{s})$ has the uniform distribution on $(0, 1)$. Similarly, since the integrand is only nonzero if $s > \bar{v} - \epsilon$, $G(a) < 1 - \exp[-a/(\bar{v} - \epsilon)]$. From these two bounds on $G(a)$, it is clear that, as $\epsilon \rightarrow 0$, G converges weakly to the exponential distribution with mean \bar{v} . That convergence is comforting to the behavioral theorist, since it would be quite unsettling if a slight change in the specifications of the model led to a large change in predicted behavior.

We shall next present a general analysis of the War of Attrition, taking a point of view which is focused on the equilibrium distribution of observed behavior. Let t_1 and t_2 be independent and uniformly distributed on $(0, 1)$, and let $\sigma \equiv F^{-1}$. (Here we denote by F^{-1} the generalized inverse of F , defined by $F^{-1}(t) = \sup\{s | F(s) < t\}$.) Then we can equivalently view the game as one in which the value of the prize to animal i is the random variable $v(t_i)$ (i.e., $v(t_i)$ has the same distribution as \bar{s}_i), and in which the quitting time chosen by animal i depends directly on his "type" t_i . Let $\sigma: (0, 1) \rightarrow [0, \infty]$ be a symmetric equilibrium strategy of the game. Without loss of generality, we may take σ to be nondecreasing;³ note that σ^{-1} is the distribution of quitting times for each animal. If animal i , with type t_i , competes until time a , his expected payoff is

$$\int_0^a (v(t_i) - s) d\sigma^{-1}(s) - a(1 - \sigma^{-1}(a)).$$

The first-order optimality condition is then

$$\frac{1}{v(t_i)} = \frac{d\sigma^{-1}/da}{1 - \sigma^{-1}(a)},$$

and equilibrium requires that $a = \sigma(t_i)$, or equivalently, that $t_i = \sigma^{-1}(a)$. Substituting the latter equality into the first-order condition yields:

$$\frac{1}{v(\sigma^{-1}(a))} = \frac{d\sigma^{-1}/da}{1 - \sigma^{-1}(a)}.$$

The left-hand side of this equation is a nonincreasing function of a , and the right-hand side is the hazard rate of the distribution of individual quitting times at equilibrium. From this we obtain a new result: The hazard rate of the duration of conflict is

³Indeed, for every strategy of one competitor, the other has a nondecreasing best response. In [6], the assumption that σ is increasing is motivated by that statement that "an animal is sometimes hungry and sometimes less so. It is common sense that it should be willing to compete more strongly for food when hungry."

nonincreasing at equilibrium. This is an empirical prediction of the model which is independent of the specification of v , and hence of F .

Substituting $a = \sigma(t_i)$ directly into the first-order condition yields a differential equation which, combined with the necessary boundary condition that $\sigma(0) = 0$, in turn yields the equilibrium strategy:

$$\sigma(t_i) = \int_0^{t_i} \frac{v(s)}{1-s} ds.$$

Note that this analysis, which is focused on the distribution of equilibrium behavior, offers a unified treatment of the game, independent of the character of F . Furthermore, because σ^{-1} is precisely the distribution of choices predicted at equilibrium, the analysis makes it relatively easy to deduce the empirical implications of the model. (This fact has already been illustrated by our discovery of the declining-hazard-rate property.) Finally, in this "distributional" form the dependence of equilibrium behavior on the specifications of the model can be clearly seen. For example, we can observe that σ increases monotonically with the function v , i.e., the distribution σ^{-1} of quitting times increases stochastically with the distribution $v^{-1} = F$ of values. Also, σ varies continuously with v (in the sense of almost-everywhere convergence), so σ^{-1} varies continuously with F (in the sense of weak convergence of probability distributions). These relations are less transparent in the traditional nondistributional analysis.

This model possesses a special monotonic structure, which we exploited in concluding that animals with higher types would choose later quitting times. In more general models, a description of the equilibrium distributions of actions is not sufficient to specify the strategies of the players: It is necessary to indicate which types will take the various actions. In order to preserve the analytical advantages noted in the preceding paragraph, in the following sections we will represent the strategies of players by the joint distributions of their types and actions, and will endow each space of "distributional" strategies with the topology of weak convergence of probability measures. In this setting, the convergence of some sequences of pure strategies to mixed strategies emerges naturally, and implies convergence of the associated distributions of actions.⁴

From the characterization of equilibrium strategies given above, it can be shown that the game has a symmetric equilibrium point in pure strategies if and only if the distribution F is atomless. We shall further explore the relationship between atomless information structures and the existence of pure-strategy equilibria in §5.

3. The formal model and the existence theorem. There are six formal elements in our model. The first four are:

- (i) the set of *players*: $N = \{1, 2, \dots, n\}$.
- (ii) the set of *types* for each player: $\{T_i\}_{i \in N}$. Each T_i is a complete, separable metric space.
- (iii) the set of *actions* available to each player: $\{A_i\}_{i \in N}$. Each A_i is a compact metric space.
- (iv) the set of possible *states*: T_0 , a complete, separable metric space.
Let $T \equiv T_0 \times \dots \times T_n$ and let $A \equiv A_1 \times \dots \times A_n$. Then the last two elements are:
 - (v) the *payoff functions*: $\{U_i\}_{i \in N}$. Each U_i is a bounded, measurable function from $T \times A$ into \mathbb{R} .
 - (vi) the *information structure*: η , a probability measure on the Borel subsets of T . Associated with the information structure η is a marginal distribution on each T_i .

⁴Other instances of the "point-opening" approach which unifies pure and mixed strategies are reported in [13] and in [21, footnote 8].

How do the equilibria in these two cases compare? Let G be the distribution of quitting times in the second case; then

$$G(a) = \Pr(\sigma(\bar{s}_i) < a) = \Pr\left(\int_0^{\bar{s}_i} \frac{sf(s)}{1-F(s)} ds < a\right).$$

Assume that F concentrates its mass on a neighborhood $(\bar{v} - \epsilon, \bar{v} + \epsilon)$. Then the integrand in the expression above is nonzero only if $s < \bar{v} + \epsilon$, and so

$$\begin{aligned} G(a) &> \Pr\left((\bar{v} + \epsilon) \int_0^{\bar{s}_i} \frac{f(s)}{1-F(s)} ds < a\right) \\ &= \Pr(F(\bar{s}_i) < 1 - \exp[-a/(\bar{v} + \epsilon)]) \\ &= 1 - \exp[-a/(\bar{v} + \epsilon)], \end{aligned}$$

since the random variable $F(\bar{s}_i)$ has the uniform distribution on $(0, 1)$. Similarly, since the integrand is only nonzero if $s > \bar{v} - \epsilon$, $G(a) < 1 - \exp[-a/(\bar{v} - \epsilon)]$. From these two bounds on $G(a)$, it is clear that, as $\epsilon \rightarrow 0$, G converges weakly to the exponential distribution with mean \bar{v} . That convergence is comforting to the behavioral theorist, since it would be quite unsettling if a slight change in the specifications of the model led to a large change in predicted behavior.

We shall next present a general analysis of the War of Attrition, taking a point of view which is focused on the equilibrium distribution of observed behavior. Let t_1 and t_2 be independent and uniformly distributed on $(0, 1)$, and let $v \equiv F^{-1}$. (Here we denote by F^{-1} the generalized inverse of F , defined by $F^{-1}(t) = \sup\{s \mid F(s) < t\}$.) Then we can equivalently view the game as one in which the value of the prize to animal i is the random variable $v(t_i)$ (i.e., $v(t_i)$ has the same distribution as \bar{s}_i), and in which the quitting time chosen by animal i depends directly on his "type" t_i . Let $\sigma: (0, 1) \rightarrow [0, \infty]$ be a symmetric equilibrium strategy of the game. Without loss of generality, we may take σ to be nondecreasing;³ note that σ^{-1} is the distribution of quitting times for each animal. If animal i , with type t_i , competes until time a , his expected payoff is

$$\int_0^a (v(t_i) - s) d\sigma^{-1}(s) - a(1 - \sigma^{-1}(a)).$$

The first-order optimality condition is then

$$\frac{1}{v(t_i)} = \frac{d\sigma^{-1}/da}{1 - \sigma^{-1}(a)},$$

and equilibrium requires that $a = \sigma(t_i)$, or equivalently, that $t_i = \sigma^{-1}(a)$. Substituting the latter equality into the first-order condition yields:

$$\frac{1}{\sigma(\sigma^{-1}(a))} = \frac{d\sigma^{-1}/da}{1 - \sigma^{-1}(a)}.$$

The left-hand side of this equation is a nonincreasing function of a , and the right-hand side is the hazard rate of the distribution of individual quitting times at equilibrium. From this we obtain a new result: The hazard rate of the duration of conflict is

³Indeed, for every strategy of one competitor, the other has a nondecreasing best response. In [6], the assumption that σ is increasing is motivated by that statement that "an animal is sometimes hungry and sometimes less so. It is common sense that it should be willing to compete more strongly for food when hungry."

nonincreasing at equilibrium. This is an empirical prediction of the model which is independent of the specification of v , and hence of F .

Substituting $a = \sigma(t_i)$ directly into the first-order condition yields a differential equation which, combined with the necessary boundary condition that $\sigma(0) = 0$, in turn yields the equilibrium strategy:

$$\sigma(t_i) = \int_0^{t_i} \frac{v(s)}{1-s} ds.$$

Note that this analysis, which is focused on the distribution of equilibrium behavior, offers a unified treatment of the game, independent of the character of F . Furthermore, because σ^{-1} is precisely the distribution of choices predicted at equilibrium, the analysis makes it relatively easy to deduce the empirical implications of the model. (This fact has already been illustrated by our discovery of the declining-hazard-rate property.) Finally, in this "distributional" form the dependence of equilibrium behavior on the specifications of the model can be clearly seen. For example, we can observe that σ increases monotonically with the function v , i.e., the distribution σ^{-1} of quitting times increases stochastically with the distribution $v^{-1} = F$ of values. Also, σ varies continuously with v (in the sense of almost-everywhere convergence), so σ^{-1} varies continuously with F (in the sense of weak convergence of probability distributions). These relations are less transparent in the traditional nondistributional analysis.

This model possesses a special monotonic structure, which we exploited in concluding that animals with higher types would choose later quitting times. In more general models, a description of the equilibrium distributions of actions is not sufficient to specify the strategies of the players: It is necessary to indicate which types will take the various actions. In order to preserve the analytical advantages noted in the preceding paragraph, in the following sections we will represent the strategies of players by the joint distributions of their types and actions, and will endow each space of "distributional" strategies with the topology of weak convergence of probability measures. In this setting, the convergence of some sequences of pure strategies to mixed strategies emerges naturally, and implies convergence of the associated distributions of actions.⁴

From the characterization of equilibrium strategies given above, it can be shown that the game has a symmetric equilibrium point in pure strategies if and only if the distribution F is atomless. We shall further explore the relationship between atomless information structures and the existence of pure-strategy equilibria in §5.

3. The formal model and the existence theorem. There are six formal elements in our model. The first four are:

- (i) the set of players: $N = \{1, 2, \dots, n\}$.
- (ii) the set of types for each player: $\{T_i\}_{i \in N}$. Each T_i is a complete, separable metric space.
- (iii) the set of actions available to each player: $\{A_i\}_{i \in N}$. Each A_i is a compact metric space.
- (iv) the set of possible states: T_0 , a complete, separable metric space.

Let $T \equiv T_0 \times \dots \times T_n$ and let $A \equiv A_1 \times \dots \times A_n$. Then the last two elements are:

- (v) the payoff functions: $\{U_i\}_{i \in N}$. Each U_i is a bounded, measurable function from $T \times A$ into \mathbb{R} .
- (vi) the information structure: η , a probability measure on the Borel subsets of T . Associated with the information structure η is a marginal distribution on each T_i .

⁴Other instances of the "point-opening" approach which unifies pure and mixed strategies are reported in [13] and in [21, footnote 8].

which we denote by η_i . Thus, if S is a Borel subset of T_1 , then $\eta_i(S) = \eta(T_0 \times S \times T_2 \times \dots \times T_n)$.

A pure strategy is a measurable function $p_i: T_i \rightarrow A_i$. In the case where T_i is uncountable, it was observed by Aumann [2] that a mixed strategy cannot be acceptably defined as a measure on the set of pure strategies. To find a more appropriate definition, he reasoned as follows.

Let us recall the intuitive meaning of a mixed strategy: It is a method for choosing a pure strategy by the use of a random device. Physically, one tosses a coin, and according to which side comes up chooses a corresponding pure strategy; or, if one wants to randomize over a continuum of pure strategies, one uses a continuous roulette wheel. Mathematically, the random device—the sides of the coin, or the set of points on the edge of the roulette wheel—constitutes a probability measure space, sometimes called a sample space; a mixed strategy is a function from this sample space to the set of all pure strategies. In other words, what we have here is precisely a random variable whose values are pure strategies. We previously attempted to work with something corresponding to the distribution of this random variable; now we propose to use the random variable itself.³

It is this idea which underlies the definition of a mixed strategy for player i as a measurable function $\sigma_i: [0, 1] \times T_i \rightarrow A_i$. Our approach of defining a (distributional) strategy as a measure on $T_i \times A_i$ provides another way of avoiding measurability problems.

DEFINITION. A *distributional strategy* for player i is a probability measure μ_i on the subsets⁶ of $T_i \times A_i$, for which the marginal distribution of T_i is η_i . Formally, this restriction on the marginal distribution is that for all $S \subset T_i$, $\mu_i(S \times A_i) = \mu_i(S)$. When the players adopt distributional strategies (μ_1, \dots, μ_n) , the expected payoff π_i to player i is defined to be:

$$\pi_i(\mu_1, \dots, \mu_n) = \int U_i(t, a) \mu_i(da_i | t_i) \dots \mu_n(da_n | t_n) \eta_i(dt_i).$$

There is a simple correspondence between a player's behavioral strategies and his distributional strategies. Given a behavioral strategy β_i , the corresponding distributional strategy μ_i is defined for each $S \times B \subset T_i \times A_i$ (and, hence, for all Borel subsets of $T_i \times A_i$) by

$$\mu_i(S \times B) = \int_S \beta_i(B, t_i) \eta_i(dt_i).$$

In the reverse direction, for any given distributional strategy μ_i the corresponding behavioral strategies are the regular conditional distributions (see [7] for definitions): $\beta_i(B, t_i) = \mu_i(B | t_i)$.

Aumann [2] has shown that there is many-to-one mapping from mixed to behavioral strategies that preserves the players' expected payoffs. We have just seen that there is another many-to-one payoff-preserving mapping from behavioral strategies to distributional strategies, i.e., each distributional strategy corresponds to an equivalence class of behavioral strategies. Since any pair of distinct distributional strategies will generally lead to distinct payoffs and since distinct distributional strategies represent different predictions about a player's behavior in the game (when his type and selected action are observable), distributional strategies give the most parsimonious representation possible of a player's meaningful strategic options.

Consider the following regularity conditions for the games we are studying.

R1: Equicontinuous Payoffs. For every player i and every $\epsilon > 0$, there is a subset

E of T such that $\eta(E) > 1 - \epsilon$ and such that the family of functions $\{U_i(t, \cdot) | t \in E\}$ is equicontinuous.⁷

R2: Absolutely Continuous Information. The measure η is absolutely continuous⁸ with respect to the measure $\hat{\eta} \equiv \eta_0 \times \dots \times \eta_n$. We denote the density of η with respect to $\hat{\eta}$ by f .

A principal requirement imposed by R1 is that for any t the players' payoffs must be continuous functions of their actions. This aspect of the condition is genuinely restrictive: it rules out many bidding games and games of timing. One cannot, however, prove an existence theorem without some such restriction. For example, there are bidding games for which no equilibrium exists.⁹

The following proposition indicates that a large number of models do meet the requirements of R1.

PROPOSITION 1. Each of the following three conditions is sufficient to imply R1.

- (a) For each i , A_i is finite.
- (b) For each i , $U_i: T \times A \rightarrow \mathbb{R}$ is a uniformly continuous function.
- (c) For each i , and for each t in T , $U_i(t, \cdot)$ is (uniformly) continuous with modulus of continuity $\delta(t, \cdot)$, and for every $\epsilon > 0$, $\delta(\cdot, \epsilon)$ is measurable.

The following important consequence of the equicontinuous payoffs condition can be proved using Lusin's Theorem.

PROPOSITION 2. In a game with equicontinuous payoffs, the following condition is satisfied:

R1*: For every player i and every $\epsilon > 0$, there are a continuous function $V_i: T \times A \rightarrow \mathbb{R}$ and a subset K of T such that (i) $\eta(K) > 1 - \epsilon$, (ii) V_i has the same bound as U_i , and (iii) V_i and U_i agree on $K \times A$.

Condition R2 is a fairly weak requirement on the joint information of the players. It is always satisfied when the variables t_0, \dots, t_n are independent, as well as when T is finite. It is also satisfied in many applied models.¹⁰ Nevertheless, R2 is a potent assumption. It allows us to express the players' expected payoffs in a convenient manner:

$$\pi_i(\mu_1, \dots, \mu_n) = \int_{T \times A} U_i(t, a) f(t) d\eta_0 d\mu_1 \dots d\mu_n. \quad (3.1)$$

The frequent applicability of R2 is emphasized by the following proposition.

PROPOSITION 3. Each of the following three conditions is sufficient to imply R2.

- (a) For each i , T_i is finite or countable.
- (b) The variables t_0, \dots, t_n are independent.

⁷A family of functions $\{f_n\}$ is equicontinuous if for every x and every $\epsilon > 0$, there is a $\delta > 0$ such that $|x - y| < \delta$ implies $|f_n(x) - f_n(y)| < \epsilon$ for every n .

⁸A measure P is absolutely continuous with respect to another measure Q if for every S , $Q(S) = 0$ implies $P(S) = 0$. The Radon-Nikodym Theorem then asserts that there is a density f of P with respect to Q , such that for every S , $P(S) = \int_S f dQ$.

⁹Consider the two-person game in which $T_1 = \{10\}$, $T_2 = \{10, 20\}$, $A_1 = A_2 = \{0, 30\}$, and

$$U_i(t_i, a_i, a_j) = \begin{cases} t_i - a_i & \text{if } a_i > a_j, \\ \frac{1}{2}(t_i - a_i) & \text{if } a_i = a_j, \\ 0 & \text{otherwise.} \end{cases}$$

Suppose $\Pr\{t_2 = 10\} = \Pr\{t_2 = 20\} = \frac{1}{2}$. This simple bidding game has no Nash equilibrium.

¹⁰For an example in which R2 does not hold, let P be the uniform distribution on the unit square and let Q be the uniform distribution on the diagonal of the square. Then $\eta = (P + Q)/2$ is not an absolutely continuous information structure.

³Quoted from [2, p. 633].

⁶Throughout this paper, we deal only with Borel sets and Borel-measurable functions. The adjectives "Borel" and "measurable" are suppressed hereafter.

(c) There exists some product measure $\hat{\eta} = \lambda_0 \times \dots \times \lambda_n$ on $T_0 \times \dots \times T_n$ such that η is absolutely continuous with respect to $\hat{\eta}$.

THEOREM 1 (Existence Theorem). *If a game has equicontinuous payoffs and absolutely continuous information (i.e., if it satisfies R1 and R2), then there exists an equilibrium point in distributional strategies.*

PROOF. We verify that conditions hold that are sufficient for the application of Glicksberg's existence theorem.¹¹

In view of the tightness¹² of η (and hence of each η_i) and the compactness of the action spaces, each player's set of distributional strategies is a tight set of probability measures; also, it is easy to check that the set is closed in the weak topology. By Prohorov's Theorem,¹³ it follows that the strategy sets are compact metric spaces in the weak topology. Convexity of these sets is also easy to check.

Since the density f of η with respect to $\hat{\eta}$ is $\hat{\eta}$ -integrable, there exists a sequence $\{f_b\}$ of bounded continuous functions such that

$$\int_T |f(t) - f_b(t)| \hat{\eta}(dt) \rightarrow 0.$$

Also, using R1*, we can approximate any U_i by a continuous function V_i . Let B be a bound on U_i and let $\{(\mu_1^k, \dots, \mu_n^k)\}$ be a sequence¹⁴ of strategy n -tuples converging to (μ_1, \dots, μ_n) . Then using (3.1),

$$\pi_i(\mu_1^k, \dots, \mu_n^k) = \int_{T \times A} V_i(t, a) f_b(t) \eta_0(dt_0) d\mu_1^k \dots d\mu_n^k + R^k(b, \epsilon), \quad \text{where (3.2)}$$

$$|R^k(b, \epsilon)| < B \int_T |f(t) - f_b(t)| \hat{\eta}(dt) + 2\epsilon B.$$

An expression similar to (3.2) can be written for $\pi_i(\mu_1, \dots, \mu_n)$. Since the integrand in (3.2) is bounded and continuous, it follows for all pairs (b, ϵ) that

$$\limsup_k |\pi_i(\mu_1^k, \dots, \mu_n^k) - \pi_i(\mu_1, \dots, \mu_n)| < 2B \int_T |f(t) - f_b(t)| \hat{\eta}(dt) + 4\epsilon B.$$

For large b and small ϵ , this bound approaches zero. Hence, π_i is continuous. From (3.1), π_i is linear.

In summary, when distributional strategies are topologized by weak convergence, the players' strategy sets are compact, convex metric spaces and the payoff functions are continuous and linear. By Glicksberg's theorem, an equilibrium exists. ■

4. The convergence theorem. Having proved the existence of a Nash equilibrium we turn our attention to sequences of games to study how variations in the specifications of a game affect the game's equilibria. Throughout the analysis, we hold the type space T fixed and we assume that R1 and R2 hold. We index games in the sequence by k . In the k th game, η^k is the distribution of types, and we define $\hat{\eta}^k = \eta_0^k \times \dots \times \eta_n^k$ and $f^k = d\eta^k / d\hat{\eta}^k$. The set of actions available to player i is a compact set A_i^k and

his payoff function is U_i^k . The corresponding items in the $*$ -game are η^* , $\hat{\eta}^*$, f^* , A_i^* and U_i^* . Let $(\mu_1^k, \dots, \mu_n^k)$ be an equilibrium point of the k th game.

THEOREM 2 (Convergence Theorem). *Suppose that each game has equicontinuous payoffs (R1) and absolutely continuous information (R2). If for all $i \in N$,*

- (i) $\{\mu_i^k\}$ converges weakly to μ_i^* , and hence $\{\hat{\eta}^k\}$ converges weakly to $\hat{\eta}^*$,
- (ii) $\{U_i^k\}$ converges uniformly to U_i^* ,
- (iii) $\{f^k\}$ converges uniformly to f^* on every compact subset of T ,
- (iv) U_i^* is continuous on $T \times A$ and f^* is continuous almost everywhere [$\hat{\eta}^*$], and
- (v) $\{A_i^k\}$ converges in the Hausdorff metric to A_i^* , then $(\mu_1^*, \dots, \mu_n^*)$ is an equilibrium of the $*$ -game.

PROOF. Suppose, contrary to the theorem, that player 1 has a pure strategy σ^* in the $*$ -game which raises his expected payoff by some positive amount α over his payoff from playing μ_1^* . Notice that a pure strategy in distributional form is simply a probability measure concentrated on the graph of a classical pure strategy. Then, by condition (v) of the theorem, there exists a sequence $\{\sigma^k\}$ of pure strategies, viewed as functions, that converges uniformly to σ^* , where σ^k is a feasible strategy in the k th game.

Arguing as in the proof of the Existence Theorem, one can show that:

- (a) $\lim_{k \rightarrow \infty} \pi_1^k(\mu_1^k, \dots, \mu_n^k) = \pi_1^*(\mu_1^*, \dots, \mu_n^*)$, and
- (b) $\lim_{k \rightarrow \infty} \pi_1^k(\sigma^k, \mu_2^k, \dots, \mu_n^k) = \pi_1^*(\sigma^*, \mu_2^*, \dots, \mu_n^*)$.

Also, by assumption,

$$(c) \pi_1^*(\sigma^*, \mu_2^*, \dots, \mu_n^*) > \pi_1^*(\mu_1^*, \dots, \mu_n^*) + \alpha.$$

From (a), (b), and (c) it follows that for all sufficiently large k , the strategy σ^k is better than μ_1^k in the k th game, contradicting our hypothesis that each $(\mu_1^k, \dots, \mu_n^k)$ is an equilibrium point. ■

Condition (iii) of the theorem is noteworthy: it is not sufficient that the η^k 's converge weakly to η^* , as the following example shows.

EXAMPLE 1. A Bayesian statistical decision problem is a game pitting one strategic player (the *statistician*) against Nature. We pose the standard estimation problem in which the statistician must estimate an unknown parameter t_0 . Let $T_0 = (0, 1)$, $T_1 = A_1 = [0, 1]$. In this problem, one often supposes that there is a quadratic loss function: $U(t_0, a) = -(t_0 - a)^2$. We define a sequence of games, in which the information structure for the k th game is concentrated on $2k$ points:

$$\Pr\{\tilde{t}_0 = 0, \tilde{t}_1 = j/k\} = 1/(2k) \quad \text{for } j = 1, \dots, k,$$

$$\Pr\{\tilde{t}_0 = 1, \tilde{t}_1 = (2j-1)/(2k)\} = 1/(2k) \quad \text{for } j = 1, \dots, k.$$

The information structure for each game conveys perfect information about \tilde{t}_0 . If $2k\tilde{t}_1$ is even, then $\tilde{t}_0 = 0$; if it is odd, then $\tilde{t}_0 = 1$. Obviously, the optimal strategy in the k th game is:

$$\sigma^k(t_1) = \begin{cases} 0 & \text{if } 2kt_1 \text{ is even,} \\ 1 & \text{if } 2kt_1 \text{ is odd.} \end{cases}$$

Passing to the weak limit, the information structure becomes:

$$\Pr\{\tilde{t}_0 = 0, \tilde{t}_1 < \alpha\} = \Pr\{\tilde{t}_0 = 1, \tilde{t}_1 < \alpha\} = \alpha/2.$$

For this information structure, \tilde{t}_0 and \tilde{t}_1 are independent! Thus, \tilde{t}_1 conveys no information about \tilde{t}_0 . The optimal strategy under this null information structure is¹⁵

¹¹We refer to the following result, which can be extracted from [14]; related results appear in [9] and [19]. Let the players' strategy spaces be nonempty compact, convex subsets of convex Hausdorff linear topological spaces. Let the payoff functions be continuous on the product of the strategy spaces, and let each player's payoff function be quasiconcave in his strategy. Then an equilibrium point exists.

¹²A set of probability measures on a metric space is called *tight* if for every $\epsilon > 0$ there is a compact set K such that for every P in the set of measures, $P(K) > 1 - \epsilon$. Any single probability measure on a complete separable metric space is tight. See [5, Theorem 1.4].

¹³See [5, Theorem 6, p. 240].

¹⁴It suffices to consider sequences (rather than nets) because the domain of π_j is a finite product of metric spaces.

¹⁵It is well known that in an estimation problem with a quadratic loss function, the optimal estimate is the posterior expectation of the unknown parameter (cf. [10, p. 228]).

$\sigma(t_i) \equiv \frac{1}{2}$. The weak limit σ^* of the sequence $\{\sigma^k\}$ is quite different (and nonoptimal), calling for the player to choose his estimate to be either 0 or 1, each with probability $\frac{1}{2}$.

The following example highlights the role of assumption R2.

EXAMPLE 2. Consider the following variant of the "Battle of the Sexes" game. Let $T_1 = T_2 = [0, 1]$ and let $A_1 = A_2 = \{1, 2\}$. Assume that the payoffs are independent of the types, and are given by the following table:

	1	2
1	2, 1	0, 0
2	0, 0	1, 2

Suppose that the information structure is given by $\Pr(\tilde{t}_1 < u, \tilde{t}_2 < v) = \min(u, v)$, where u and v are numbers in $[0, 1]$. Thus, $t_1 \equiv t_2$, and these variables are uniformly distributed. Now consider the pure strategies

$$\sigma^k(t) = \begin{cases} 1 & \text{if the integer part of } kt \text{ is odd,} \\ 2 & \text{otherwise.} \end{cases}$$

If both players adopt the strategy σ^k , perfect coordination is achieved and the strategy pair is an equilibrium point. The limit of this sequence of pure strategies is the following distributional strategy for player i :

$$\Pr(\tilde{t}_i < u, a_i = 1) = \Pr(\tilde{t}_i < u, a_i = 2) = u/2. \tag{4.1}$$

Equation (4.1) asserts that the player ignores his information and randomizes his choice of action, choosing each action with probability $\frac{1}{2}$. This "limit" is not an equilibrium: a better response for player i would be to choose action i with certainty. Thus, the set of equilibria of this game is *not* closed in the weak topology, and hence the Convergence Theorem cannot apply to this game.

5. **Pure strategies.** Game-theoretic models are often criticized for their reliance on mixed-strategy equilibrium points. Critics argue that mixed strategies have no role in a behavioral theory: people do not base their decisions on the roll of a die or the toss of a coin.

There are several kinds of responses one might make to such criticisms. First, one can challenge the premise that mixed strategies are not actually observed. Close decisions are often made on the basis of minor distinctions or simple whimsy, factors which are hardly less random than roulette wheels. Second, one can claim that the critics have failed to show that there is any observable difference between mixed and pure strategic behavior. Third, models without pure-strategy equilibria may nevertheless have pure-strategy ϵ -equilibria¹⁶ for every positive ϵ . If these are "close" to the mixed-strategy equilibria in some appropriate sense, and if the ϵ -equilibrium concept seems empirically justifiable, then mixed strategies can be viewed as a convenient technical device for behavioral modeling. Finally, one can accede to the critics and try to identify classes of games for which pure-strategy equilibria exist.

The two theorems that we offer in this section address this whole range of possible responses. The Denseness Theorem asserts that if a player's type has an atomless distribution,¹⁷ then his set of pure strategies is dense in his entire strategy set. It then follows that if one can only observe points $T_i \times A_i$ subject to some continuous measurement error, pure strategies and mixed strategies are empirically indistinguish-

able. Moreover, starting at any mixed-strategy equilibrium point, the Denseness Theorem implies that one can locate a nearby pure-strategy ϵ -equilibrium point.

The Purification Theorem identifies a class of games with the property that every mixed strategy has a corresponding pure strategy such that (i) the pure strategy is optimal whenever the mixed strategy is, (ii) substituting the pure strategy for the mixed strategy leaves the other players' decision problems unchanged, and (iii) an observer seeing only t_0 and a would be unable to distinguish the pure strategy from the mixed strategy.

THEOREM 3 (Denseness Theorem). *Suppose that η_i is atomless. The player i 's set of pure strategies is dense in his set of distributional strategies.*

PROOF. Fix a distributional strategy μ for player i , and fix $\epsilon > 0$. Since A_i is compact, there exists a finite ϵ -partition B_1, \dots, B_k of A_i (i.e., a partition such that each B_l has radius less than ϵ). Since T_i is complete and separable, η_i is tight [5, Theorem 4.1]. Therefore T_i can be partitioned into (K, S_0) , where K is compact and $\eta_i(S_0) < \epsilon$. Also, K has a finite ϵ -partition (S_1, \dots, S_m) . Since η_i is atomless, each S_j can in turn be partitioned into sets S_{j1}, \dots, S_{jk} , such that $\eta_i(S_{jl})/\eta_i(S_j) = \mu(B_l|S_j)$ for $l = 1, \dots, k$ (cf. [12], or [8, §2.2, problem 23]).

Fix any points b_1, \dots, b_k in B_1, \dots, B_k , and define a pure strategy $\sigma_i: T_i \rightarrow A_i$ by $\sigma_i(t) = b_l$ for all t in S_{jl} . It is routine to verify that as $\epsilon \rightarrow 0$, σ_i converges weakly to μ (cf. [4, p. 603]). ■

In the statement and proof of the Denseness Theorem, we have assumed neither equicontinuous payoffs nor absolutely continuous information. In the course of proving the Existence Theorem, these two conditions were shown to imply that each player's expected payoff is a continuous function of the n -tuple of strategies. Thus, these continuity conditions, together with the Denseness Theorem, ensure that there are pure-strategy ϵ -equilibrium points arbitrarily near any mixed-strategy equilibrium point. In view of the Existence Theorem, we have the following result.

COROLLARY. *If a game satisfies the equicontinuous payoffs and absolutely continuous information conditions (R1 and R2), if each η_i is atomless, and if the action spaces are compact, then for every $\epsilon > 0$ there exists a pure-strategy ϵ -equilibrium point.*

Adapting terminology introduced by Radner and Rosenthal [27] to our model, we say that a pure strategy σ_1 is a *purification* of the strategy μ_1 if two conditions are met:

(5.1) For almost every t_1 , $\sigma_1(t_1)$ lies in the support of $\mu_1(\cdot|t_1)$. (Consequently, if μ_1 is a best response to some $(n-1)$ -tuple of strategies (μ_2, \dots, μ_n) and if R1 holds, then σ_1 is also a best response.)

(5.2) For every player $i \neq 1$ and every $(n-1)$ -tuple (μ_2, \dots, μ_n) of strategies for players $2, \dots, n$, substituting σ_1 for μ_1 preserves i 's expected payoff: $\pi_i(\mu_1, \dots, \mu_n) = \pi_i(\sigma_1, \mu_2, \dots, \mu_n)$.

It is clear from the definition that if (μ_1, \dots, μ_n) is an equilibrium point and σ_1 is a purification of μ_1 , then $(\sigma_1, \mu_2, \dots, \mu_n)$ is also an equilibrium point. Radner and Rosenthal have shown that if (i) the players' types are mutually independent, (ii) each η_i is atomless, (iii) each player's payoff depends only on his own type t_i and the list of actions a (that is, $U_i = U_i(t_i, a)$), and (iv) the action spaces are all finite, then each strategy of each player has a purification.

In a paper studying statistical decision problems, Dvoretzky, Wald and Wolfowitz [11] proved that if T_0 is a finite set and η_1 is atomless, then for every strategy μ_1 there is a pure strategy σ_1 satisfying condition (5.1) and the following condition:

(5.3) Conditional on any t_0 , the distributions induced on A_1 by μ_1 and σ_1 are identical, i.e., for any subset B of A_1 ,

$$\eta_1(\sigma_1^{-1}(B)|t_0) = \int \mu_1(B|t_1)\eta_1(dt_1|t_0).$$

¹⁶An ϵ -equilibrium point of a game is an n -tuple (μ_1, \dots, μ_n) of players' strategies, such that for every player i and every alternative strategy μ'_i , $\pi_i(\mu_1, \dots, \mu_n) + \epsilon > \pi_i(\mu_1, \dots, \mu'_i, \dots, \mu_n)$.

¹⁷A probability measure η is *atomless* if for every B with $\eta(B) > 0$, there is a $C \subset B$ for which

As a corollary to the Dvoretzky–Wald–Wolfowitz result and to the Existence Theorem, we obtain the Purification Theorem.

THEOREM 4 (Purification Theorem). *If (i) conditional on t_0 , the players' types are independent, (ii) each η_i is atomless, (iii) each player's payoff depends only on the state variable t_0 , his own type t_i , and the list of actions a (that is, $U_i = U_i(t_0, t_i, a)$), (iv) each player's action set A_i is finite,¹⁸ (v) payoffs are equicontinuous (R1 holds), and (vi) T_0 is a finite set, then each strategy of each player has a purification satisfying conditions (5.1), (5.2), and (5.3). Furthermore, the game has an equilibrium point, and hence has an equilibrium point in pure strategies.*

PROOF. It is direct to verify that conditions (i), (iii), (iv), and (5.3) imply (5.2), so the existence of purifications follows from the Dvoretzky–Wald–Wolfowitz theorem. Also, it is direct to show that conditions (i) and (vi) imply R2, so existence follows from the Existence Theorem. ■

Theorem 4 extends the Radner–Rosenthal purification result [27] to allow some players to have information about variables that appear in other players' payoff functions. Models with this latter feature are known as “adverse selection” models, and play an important role in information economics.

6. Complements and comments. Our formulation of games with incomplete information contains the assumption that an exogenously-specified metric on the type space T is available. It might appear preferable to simply treat the players' types as points in a general measurable space, without assuming any topological structure. Indeed, the critical conditions of equicontinuous payoffs (R1) and absolutely continuous information (R2) depend only on measure-theoretic properties of T . Yet the topology on types was necessary in order to define the weak topology on distributional strategies, and this topology played a crucial role in the Existence, Convergence, and Denseness Theorems.

How might we have proceeded, if only a measurable structure on T had been given? A natural approach would have been to define endogenously a metric on T which reflects the nature of the game. In general, a player's type has two aspects. First, it influences his payoffs, as well as the payoffs of others. Additionally, it affects his beliefs about the types of his competitors, and hence about their behavior. As noted, for example, in [20] and [24], both of these effects are metrizable. We here define two metrics (actually, pseudometrics) on T which correspond to the two effects. For the sake of expositional simplicity, our analysis will be in terms of the canonical form of the game (cf. footnote 1; we assume that the state has been integrated out of the payoff functions, and that $T = T_1 \times \dots \times T_n$).

Assume that the players' payoff functions are bounded, and are continuous on A for each t in T . For any player i and types t'_i and t''_i in T_i , define

$$d_i^1(t'_i, t''_i) = \sum_{j=1}^n \sup_{a \in A} \sup_{t_{-i} \in T_{-i}} |U_j(t'_i, t_{-i}, a) - U_j(t''_i, t_{-i}, a)|,$$

and for t' and t'' in T , define $d^1(t', t'') = \sum_{i=1}^n d_i^1(t'_i, t''_i)$. With respect to the product topology on $T \times A$ induced by this metric on T and the originally-given topology on A , all of the players' payoff functions are continuous. (Of course, this statement is trivial if, for example, the metric d^1 induces the discrete topology on T .)

For any player i and type t_i in T_i , let $\eta_{-i}(\cdot | t_i)$ denote the conditional distribution¹⁹

¹⁸Yakar Kannai has provided an example which, when suitably extended, shows that this condition cannot be weakened to allow arbitrary compact action spaces.

¹⁹We assume that a regular conditional distribution exists—as it does, for example, if T is a Borel space (see [7, Chapter 4]).

on T_{-i} induced by η . For any t'_i and t''_i in T_i , define

$$d_i^2(t'_i, t''_i) = \sup_{B \subset T_{-i}} |\eta_{-i}(B | t'_i) - \eta_{-i}(B | t''_i)|,$$

and for t' and t'' in T , define

$$d^2(t', t'') = \sum_{i=1}^n d_i^2(t'_i, t''_i).$$

The metric d_i^2 might be termed the “continuous beliefs” metric for player i .

Given a metric d on T , we say that the d -topology is measurable if all d -open sets are measurable. A probability measure η on T is d -tight if the d -topology on T is measurable, and if for every $\epsilon > 0$ there is a d -compact set of measure at least $1 - \epsilon$.

PROPOSITION 4. (a) *If η is d^1 -tight, then the game has equicontinuous payoffs.*
(b) *If η is d^2 -tight, then the game has absolutely continuous information.*

A natural endogenously-determined topology on any T_i is that induced by the metric $d_i^1 + d_i^2$. All of the results of the paper could have been derived in terms of these metrics. Since essentially all proofs of the existence of equilibria for classes of games depend on the compactness of the players' strategy spaces, and since the tightness of η is required to ensure the compactness of the sets of distributional strategies in a game with incomplete information, it seems unlikely that our Existence Theorem can be substantially generalized.

A proof of Proposition 4 is given in [24]. That paper also discusses the connections between this paper, [1], and [27]. Those latter two papers approach games with incomplete information with objectives different from ours, and contain, among other results, special cases of our Existence Theorem.

Harsanyi [16] studied the perturbation of a complete-information game by the introduction of payoff uncertainty, and showed that almost any mixed-strategy equilibrium in the original game can be distributionally approximated by pure-strategy equilibria in the perturbed games. This may be viewed as a lower-hemicontinuity result which complements our Convergence Theorem.

Throughout this paper, we have assumed that all the players agree on the information structure η . Let us now suppose instead that, according to player i , the joint distribution on $T_0 \times \dots \times T_n$ is η^i . Let $\eta = (\eta^1 + \dots + \eta^n)/n$. It is straightforward to check that if each η^i is an absolutely continuous information structure, then each η^i is absolutely continuous with respect to η . Let $f^i = d\eta^i/d\eta$ be the density of η^i with respect to η . If we replace f everywhere it appears by f^i or f^j , as appropriate, then all of our arguments retain their validity. Thus, nothing essential is affected by the consistency assumption used in §§1–5.

Acknowledgement. We are grateful to R. J. Aumann, R. Radner, R. W. Rosenthal, and a referee for their helpful comments.

References

- [1] Armbruster, W. and Böge, W. (1979). Bayesian Game Theory. *Game Theory and Related Topics*. O. Moeschlin and D. Pallaschke (eds.). North-Holland Publishing Co., Amsterdam, 17–28.
- [2] Aumann, R. J. (1964). Mixed and Behavior Strategies in Infinite Extensive Games. *Advances in Game Theory*, Ann. of Math. Studies No. 52, Princeton University Press, Princeton, N.J., 627–650.
- [3] ——— (1974). Subjectivity and Correlation in Randomized Strategies. *J. Math. Economics* 1 67–96.
- [4] Bellman, R. and Blackwell, D. (1949). Some Two-Person Games Involving Bluffing. *Proc. Nat. Acad. Sci. U.S.A.* 35 600–605.
- [5] Billingsley, P. (1968). *Convergence of Probability Measures*. John Wiley and Sons, New York.
- [6] Bishop, D. T., Cannings, C. and Smith, J. Maynard. (1978). The War of Attrition with Random Rewards. *J. Theoret. Biol.* 74 377–388.

- [7] Breiman, L. (1968). *Probability*. Addison-Wesley, Menlo Park.
- [8] Chung, K. (1974). *A Course in Probability Theory*. Second ed., Academic Press, New York.
- [9] Debreu, G. (1952). A Social Equilibrium Existence Theorem. *Proc. Nat. Acad. Sci. U.S.A.* 38 886-893.
- [10] DeGroot, M. (1970). *Optimal Statistical Decisions*. McGraw-Hill, New York.
- [11] Dvoretzky, A., Wald, A. and Wolfowitz, J. (1951). Elimination of Randomization in Certain Statistical Decision Procedures and Zero-Sum Two-Person Games. *Ann. Math. Statist.* 22 1-21.
- [12] ———, ——— and ———. (1951). Relations among Certain Ranges of Vector Measures. *Pacific J. Math.* 1 59-74.
- [13] Engelbrecht-Wiggans, R., Milgrom, P. and Weber, R. J. (1983). Competitive Bidding with Propriety Information. *J. Math. Economics* 11 161-169.
- [14] Glicksberg, I. (1952). A Further Generalization of Kakutani's Fixed Point Theorem with Application to Nash Equilibrium Points. *Proc. Nat. Acad. Sci. U.S.A.* 38 170-172.
- [15] Harsanyi, J. C. (1967-68). Game with Incomplete Information Played by Bayesian Players. *Management Sci.* 14 159-182, 320-334, 486-502.
- [16] ———. (1973). Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points. *Internat. J. Game Theory* 2 1-23.
- [17] Kreps, D. M. and Wilson, R. (1982). Sequential Equilibria. *Econometrica* 50 863-894.
- [18] ——— and ———. (1982). Reputation and Imperfect Information. *J. Economic Theory* 27 253-279.
- [19] Ky Fan. (1952). Fixed-Point and Minimax Theorems in Locally Convex Topological Linear Spaces. *Proc. Nat. Acad. Sci. U.S.A.* 38 121-126.
- [20] Milgrom, P. (1979). The Structure of Information in Competitive Bidding. Ph.D. dissertation, Stanford University (1978), Garland Publ. Co., New York.
- [21] ———. (1981). Rational Expectations, Information Acquisition, and Competitive Bidding. *Econometrica* 49 921-943.
- [22] ——— and Roberts, D. J. (1982). Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis. *Econometrica* 50 443-459.
- [23] ——— and ———. (1982). Predation, Reputation, and Entry Deterrence. *J. Economic Theory* 27 280-312.
- [24] ——— and Weber, R. J. (1981). Topologies on Information and Strategies in Games with Incomplete Information. *Game Theory and Mathematical Economics*. O. Moeschlin and D. Pallaschke (eds.), North-Holland Publishing Company, Amsterdam.
- [25] ——— and ———. (1982). A Theory of Auctions and Competitive Bidding. *Econometrica* 50 1089-1122.
- [26] ——— and ———. (1982). The Value of Information in a Sealed-Bid Auction. *J. Math. Economics* 10 105-114.
- [27] Radner, R. and Rosenthal, R. W. Private Information and Pure-Strategy Equilibria. *Math. Oper. Res.* (to appear).
- [28] Selten, R. (1975). Re-Examination of the Perfectness Concept for Equilibrium Points in Extensive Games. *Internat. J. Game Theory* 4 22-55.
- [29] ———. (1978). The Chain-Store Paradox. *Theory and Decision* 9 127-159.
- [30] Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed-Tenders. *J. Finance* 16 8-37.
- [31] Wilson, R. (1971). A Bidding Model of Perfect Competition. *Rev. Economic Studies* 44 511-518.

MILGROM: SCHOOL OF ORGANIZATION AND MANAGEMENT, YALE UNIVERSITY, BOX 1A, NEW HAVEN, CONNECTICUT 06520

WEBER: J. L. KELLOGG GRADUATE SCHOOL OF MANAGEMENT, NORTHWESTERN UNIVERSITY, EVANSTON, ILLINOIS 60201