

Milgrom, Paul; Weber, Robert

Working Paper

Distributional Strategies for Games with Incomplete Information

Discussion Paper, No. 428R

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Milgrom, Paul; Weber, Robert (1981) : Distributional Strategies for Games with Incomplete Information, Discussion Paper, No. 428R, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/220788>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Discussion Paper #428R

DISTRIBUTIONAL STRATEGIES FOR GAMES
WITH INCOMPLETE INFORMATION

by

Paul R. Milgrom and Robert J. Weber

August, 1981

J.L. Kellogg Graduate School of Management

Northwestern University

Evanston, Illinois 60201

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.

ABSTRACT

We study games with incomplete information from a point of view which emphasizes the empirical predictions arising from a model. 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 varies with the parameters of a game. We illustrate the application of the distributional point of view to the computation of equilibria and the determination of their properties.

DISTRIBUTIONAL STRATEGIES FOR GAMES WITH INCOMPLETE INFORMATION

by

Paul R. Milgrom and Robert J. Weber

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 strike behavior, or competition between animals. 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 which always have pure-strategy equilibrium points.

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 \tilde{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 \mathfrak{A}_i be the collection of Borel subsets of A_i . A behavioral strategy is a function $\beta_i: \mathfrak{A}_i \times T_i \rightarrow [0,1]$ with these two properties: (i) For every $B \in \mathfrak{A}_i$, 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 Section 2, we present an example to illustrate the desirability of a distributional view of strategic behavior. Section 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 Section 4, asserts that when the data that specify a game are varied "continuously," the set of equilibria varies upper-hemicontinuously. Two examples serve to clarify the conditions of the theorem.

We take up the matter of pure-strategy equilibria in Section 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 the set of all his distributional strategies when his type has an atomless distribution. Its corollary asserts the existence of approximate equilibria in pure strategies. To study exact equilibria, we need the following definition. A mixed strategy $\hat{\sigma}$ is said to have a purification if there is a pure strategy σ with these two properties: (i) σ is a best reply to the opposing strategies whenever $\hat{\sigma}$ is, and (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 $\hat{\sigma}$. 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. Section 6 uses the distributional strategy approach to solve completely the "War of Attrition" game introduced in Section 2, and Section 7 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 how the standard approach to games with incomplete information obscures the important relationship between the data specifying a game and the game's equilibria. The game we analyze is known as the War of Attrition.² It has been used to study conflict among animals [6], although it applies equally well to other conflicts, such as labor-management disputes involving strikes.

In the animal conflict interpretation, two animals battle for a valuable prize, such as food or the opportunity to mate. Animal i ($i = 1, 2$) would be willing to fight for up to t_i minutes to acquire the prize if he could be certain the battle would end with victory. If he actually fights m minutes and drives off his competitor, his payoff is $t_i - m$. If he fights m minutes and then retreats, his payoff is $-m$. (That is, there is an opportunity cost associated with the time spent in conflict.) Each animal, knowing only his own type t_i , must decide how long he is willing to persevere in the battle.

Formally, we have $T_1 = T_2 = A_1 = A_2 = R_+$. Also,

$$U_i(t_1, t_2, a_1, a_2) = \begin{cases} t_i - a_j & \text{if } a_i > a_j \\ -a_i & \text{otherwise} \end{cases},$$

where $\{i, j\} = \{1, 2\}$. Suppose \tilde{t}_1 and \tilde{t}_2 are independent and identically distributed with common distribution F . Further suppose that F is absolutely continuous, with continuous density $F' = f$. The standard treatment of this game proceeds in the following manner.

Since the game is symmetric, it is natural to look for a symmetric Nash equilibrium point³ in which both competitors use some pure strategy σ . Assume animal 1 "believes" that an increasing, differentiable⁴ strategy σ will be used by animal 2. Then if 1's type is t and if he plans to persevere until a , his expected payoff is

$$\int_0^{\sigma^{-1}(a)} [t - \sigma(s)] f(s) ds - a[1 - F(\sigma^{-1}(a))] .$$

If σ is to be a symmetric equilibrium strategy, this expected payoff must be maximized when $a = \sigma(t)$. The resulting first-order condition, together with the boundary condition $\sigma(0) = 0$, leads to

$$(2.1) \quad \sigma(t) = \int_0^t \frac{s f(s)}{1 - F(s)} ds ;$$

the strategy pair (σ, σ) is known to be the unique symmetric equilibrium point of the game.

If, instead of assuming that F is continuous, we assume that F concentrates all its mass at some value v , then no symmetric pure-strategy equilibrium point will exist. Consequently, we search for a symmetric mixed-strategy equilibrium. Let $G^*(a)$ be the probability at equilibrium that an animal will not perservere beyond time a . It is straightforward to show that the distribution G^* must be atomless and that its support must be convex (cf. [6]).

Let π be the equilibrium expected payoff. Then, for every action in the support of G^* ,

$$\begin{aligned} \pi &= \int_0^a (v - x) dG^*(x) - a(1 - G^*(a)) \\ (2.2) \quad &= vG^*(a) - \int_0^a (1 - G^*(x))dx . \end{aligned}$$

(The latter equality is obtained using integration by parts.) Since $G^*(0) = 0$ (because G^* is atomless), setting $a = 0$ in (2.2) leads to the conclusion $\pi = 0$. With $\pi = 0$, the unique solution of the integral equation (2.2) is

$$(2.3) \quad G^*(a) = 1 - \exp(-a/v) .$$

Thus the equilibrium strategy calls for an animal to randomize his efforts using an exponential distribution whose mean v is the value of the object.

These two variants of the War of Attrition appear rather different on the surface. In the first game, there is a pure-strategy equilibrium

which specifies how long an animal should persevere as a function of how hungry it is for the prize. In the second game, the animals randomize over actions.

In order to compare the predicted behavior in these two variants, we determine the distribution G of the first animal's actions which is induced by the equilibrium strategy in the first game. To compute G from the data of the game we begin with the observation that

$$(2.4) \quad G(a) = \Pr(\sigma(\tilde{t}_1) < a) = \Pr\left(\int_0^{\tilde{t}_1} \frac{s f(s)}{1-F(s)} ds < a\right) .$$

Assume that F concentrates its mass on a neighborhood $(v-\epsilon, v+\epsilon)$. Then the integrand in (2.4) is nonzero only if $s < v+\epsilon$, and so

$$\begin{aligned} G(a) &> \Pr\left((v+\epsilon) \int_0^{\tilde{t}_1} \frac{f(s)}{1-F(s)} ds < a\right) \\ &= \Pr(F(\tilde{t}_1) < 1-\exp[-a/(v+\epsilon)]) \\ &= 1-\exp[-a/(v+\epsilon)] , \end{aligned}$$

since the random variable $F(\tilde{t}_1)$ has the uniform distribution on $(0,1)$ (cf. [8], page 38). Similarly, since the integrand is only nonzero if $s > v-\epsilon$, $G(a) < 1-\exp[-a/(v-\epsilon)]$. From these two bounds on $G(a)$, it is clear that G approaches G^* pointwise as $\epsilon \rightarrow 0$. 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 return to this example in Section 6. There, we illustrate how a direct distributional analysis of the game unifies the above results, yields new insights into the nature of equilibrium behavior, and clarifies the dependence of the equilibrium strategies on the underlying parameters of the game.

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 which we denote by η_i . Thus, if S is a Borel subset of T_1 , then

$$\eta_1(S) = \eta(T_0 \times S \times T_2 \times \dots \times T_n) .$$

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.⁵

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 on T_i is η_i . Formally, this restriction on the marginal distribution is that for all $S \subset T_i$, $\mu_i(S \times A_i) = \eta_i(S)$. When the players adopt distributional strategies (μ_1, \dots, μ_n) , the expected payoff π_i to player i is defined to be:

$$\pi_i(\mu_1, \dots, \mu_n) = \int U_i(t, a) \mu_1(da_1 | t_1) \dots \mu_n(da_n | t_n) \eta(dt) .$$

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 $S \times B \subset T_i \times A_i$ by⁷

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

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 a 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 mapping from behavioral strategies to distributional strategies; this, too, preserves

payoffs. 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.

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 continuous payoffs condition can be proved using Lusin's Theorem.

Proposition 2: In a game with continuous payoffs, the following condition is satisfied:

R1*: For every player i and every $\epsilon > 0$, there is a continuous function $V_\epsilon: T \times A \rightarrow \mathbb{R}$ and a subset K of T such that (i) $\eta(K) > 1 - \epsilon$, (ii) V_ϵ has the same bound as U_i , and (iii) V_ϵ 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 $\tilde{t}_0, \dots, \tilde{t}_n$ 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:

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

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.
- (b) The variables $\tilde{t}_0, \dots, \tilde{t}_n$ are independent.
- (c) There exists some product measure $\hat{\lambda} = \lambda_0 \times \dots \times \lambda_n$ on $T_0 \times \dots \times T_n$ such that η is absolutely continuous with respect to $\hat{\lambda}$.

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_ε . 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),

$$\begin{aligned} & \pi_i(\mu_1^k, \dots, \mu_n^k) \\ (3.2) \quad & = \int_{T \times A} V_\varepsilon(t, a) f_b(t) \eta_0(dt_0) d\mu_1^k \dots d\mu_n^k + R^k(b, \varepsilon) , \end{aligned}$$

where

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

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

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

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. Q.E.D.

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_1^* is continuous on $T \times A$ and f^* is continuous almost everywhere $[\hat{n}^*]$, 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:

$$(a) \quad \lim_{k \rightarrow \infty} \pi_1^k(\mu_1^k, \dots, \mu_n^k) = \pi_1^*(\mu_1^*, \dots, \mu_n^*) \quad , \text{ and}$$

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

Also, by assumption,

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

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. Q.E.D.

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

$$T_0 = \{0, 1\}, \quad T_1 = A_1 = [0, 1] \quad .$$

In this problem, one often supposes that there is a quadratic loss function:

$$U(t_0, a) = -(t_0 - a)^2 \quad .$$

We define a sequence of games, in which the information structure for the k -th game is concentrated on $2k$ points:

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

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

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 $\sigma(t_1) \equiv 1/2$. 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 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

$$\Pr\{\tilde{t}_1 \leq u, \tilde{t}_2 \leq v\} = \min(u, v) \quad ,$$

where u and v are numbers in $[0,1]$. Thus, $\tilde{t}_1 \equiv \tilde{t}_2$, and these variables are uniformly distributed. Now consider the pure strategies

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

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 :

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

Equation (4.1) asserts that the player ignores his information and randomizes his choice of action, choosing each action with probability $1/2$. This "limit" is not an equilibrium: a better response for player i would be to choose action 1 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 in $T_1 \times A_1$ subject to some continuous measurement error, pure strategies and mixed strategies are empirically indistinguishable. 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 4: (Denseness Theorem) Suppose that η_i is atomless. Then 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 1.4). 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 $\eta_i(S_{j\ell})/\eta_i(S_j) = \mu(B_\ell | S_j)$ for $\ell = 1, \dots, k$ (cf. [12], or [8], Section 2.2, problem 23).

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

Q.E.D.

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,¹⁷ (iii) player i '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 ,

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

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 mutually independent, (ii) each η_i is atomless, (iii) player i 's payoffs depend 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) payoffs are equicontinuous (R1 holds), and (v) 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), 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 (v) imply R2, so existence follows from the Existence Theorem. Q.E.D.

Theorem 4 extends the Radner-Rosenthal result to include the case of compact actions spaces and, more importantly, 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. Applying the Distributional Point of View

In order to illustrate the insights into a problem which can be gained through the distributional approach, and also to demonstrate the computational techniques associated with the approach, we return our attention to the War of Attrition. It warrants noting that neither the Existence Theorem nor the Convergence Theorem applies to this game, since the payoff functions are not continuous in the players' acts. Hence, our analysis will demonstrate the applicability of our ideas in an even broader setting than that in which those theorems were developed.

To facilitate the distributional analysis, we shall reformulate the War of Attrition game. In our new formulation, each animal i has a type \tilde{t}_i which is uniformly distributed on $(0,1)$. If animals 1 and 2

resolve to perservere until times m_1 and m_2 , respectively, then their payoffs are given in the following table.

<u>Condition</u>	<u>Animal 1</u>	<u>Animal 2</u>
$m_1 > m_2$	$\alpha(t_1) - m_2$	$-m_2$
$m_2 > m_1$	$-m_1$	$\alpha(t_2) - m_1$
$m_1 = m_2$	$1/2 \alpha(t_1) - m_1$	$1/2 \alpha(t_2) - m_2$

Here, $\alpha(t_i)$ represents the value of the prize to animal i . To model the case where the value (in the original formulation of the game) has some strictly increasing continuous distribution F , we may take $\alpha = F^{-1}$. If F concentrates all its mass at a point v , we may take $\alpha \equiv v$. In general, for any desired value distribution F , there is some nondecreasing function α such that $\alpha(\tilde{t}_i)$ has that distribution.¹⁸ Thus, we may assume without loss of generality that α is nondecreasing.

Let $\sigma : [0,1] \rightarrow \mathbb{R}_+$ denote a symmetric equilibrium strategy for the game. As in Section 2, we may take σ to be nondecreasing. Notice that σ^{-1} is the distribution of quitting times for each animal. Thus, if animal 1 with type t_1 fights until time a , his expected payoff is

$$(6.1) \quad \int_0^a (\alpha(t_1) - s) d\sigma^{-1}(s) - a(1 - \sigma^{-1}(a)) .$$

The first-order optimality condition is:

$$(6.2) \quad \frac{1}{\alpha(t_1)} = \frac{d\sigma^{-1}/da}{1 - \sigma^{-1}(a)} ,$$

and equilibrium requires that $a = \sigma(\tau_1)$, or equivalently, that $\tau_1 = \sigma^{-1}(a)$. Substituting the latter equality into (6.2) yields:

$$(6.3) \quad \frac{1}{\alpha(\sigma^{-1}(a))} = \frac{d\sigma^{-1}/da}{1 - \sigma^{-1}(a)} .$$

The left-hand side of (6.3) 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 nonincreasing. This is an empirical prediction of the model which is independent of the specification of α .

Substituting $a = \sigma(\tau_1)$ into (6.2) yields the differential equation:

$$(6.4) \quad \sigma'(\tau_1) = \frac{\alpha(\tau_1)}{1 - \tau_1} .$$

Since an animal of type $\tau_1 = 0$ has an expected payoff at equilibrium of $-\sigma(0)$, and since he can obtain at least zero by never fighting, we must have $\sigma(0) = 0$. Together with (6.4), that yields the equilibrium strategy:

$$(6.5) \quad \sigma(\tau_1) = \int_0^{\tau_1} \frac{\alpha(s)}{1 - s} ds .$$

The analysis presented here improves upon the traditional analysis given in Section 2 in at least three respects. First, it offers a

unified treatment of the game, independent of the character of F . Second, because σ^{-1} is precisely the distribution of choices predicted at equilibrium, formulas such as (6.5) make 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. Third, in distributional form the dependence of equilibrium behavior on the specifications of the model is clearly seen. For example, from (6.5) we see that σ increases monotonically with the function α , i.e., the distribution σ^{-1} of quitting times increases stochastically with the distribution α^{-1} of values. Also, σ varies continuously with α (in the sense of almost-everywhere convergence), so σ^{-1} varies continuously with α^{-1} (in the sense of weak convergence). These inferences are more difficult to draw from the traditional analysis.

The kind of analysis employed in this section is especially useful for studying incomplete information models in which the types and actions are real numbers and the equilibrium strategies are monotone functions. Models with these properties are common in economic applications [6, 13, 21, 30, 31], and we have used the techniques presented here to study several such models [13, 22, 23, 25].

7. 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 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 $\hat{\eta} = (\eta^1 + \dots + \eta^n)/n$. 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 $\hat{\eta}$. Let $f^i = d\eta^i/d\hat{\eta}$ be the density of η^i with respect to $\hat{\eta}$. 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 Sections 1-5.

FOOTNOTES

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:

$$V_i(t_1, \dots, t_n, a_1, \dots, a_n) = E[U_i(\tilde{t}_0, \dots, \tilde{t}_n, a_1, \dots, a_n) | t_1, \dots, t_n] .$$

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 .

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

³The theoretical biologists call these "evolutionarily stable strategies" (ESS's).

⁴It can be proved that these properties must hold at a symmetric equilibrium point of the game. In [6], the assumption that σ is increasing is motivated by the 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."

⁵Quoted from [2], page 633.

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

⁷It is well known that a measure on a product space is uniquely determined by the measure it assigns to "rectangles." Thus, the definition completely determines μ_i .

⁸A measure P is absolutely continuous with respect to another measure Q if for every set 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

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

Suppose $\Pr\{\tilde{t}_2 = 10\} = \Pr\{\tilde{t}_2 = 20\} = 1/2$. 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.

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, page 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], page 228).

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 $\eta(B) > \eta(C) > 0$.

¹⁸See [8], page 38, problem 4.

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

REFERENCES

- [1] W. Armbruster and W. Boge, "Bayesian game theory," Game Theory and Related Topics, O. Moeschlin and D. Pallaschke (eds.), North-Holland Publishing Co. (1979), 17-28.
- [2] R.J. Aumann, "Mixed and behavior strategies in infinite extensive games," Advances in Game Theory, Annals of Mathematics Studies No. 52, Princeton University Press (1964), 627-650.
- [3] R.J. Aumann, "Subjectivity and correlation in randomized strategies," Journal of Mathematical Economics 1 (1974), 67-96.
- [4] R. Bellman and D. Blackwell, "Some two-person games involving bluffing," Proc. Nat. Acad. Sci. (U.S.A.) 35 (1949), 600-605.
- [5] P. Billingsley, Convergence of Probability Measures, John Wiley and Sons, New York (1968).
- [6] D.T. Bishop, C. Cannings, and J. Maynard Smith, "The War of Attrition with random rewards," J. Theor. Biol. 74 (1978), 377-388.
- [7] L. Breiman, Probability, Addison Wesley, Menlo Park (1968).
- [8] K. Chung, A Course in Probability Theory, second edition, Academic Press, New York (1974).
- [9] G. Debreu, "A social equilibrium existence theorem," Proc. N.A.S. 38 (1952), 886-893.
- [10] M. DeGroot, Optimal Statistical Decisions, McGraw-Hill, New York (1970).
- [11] A. Dvoretzky, A. Wald, and J. Wolfowitz, "Elimination of randomization in certain statistical decision procedures and zero-sum two-person games," Ann. Math. Stat. 22 (1951), 1-21.
- [12] A. Dvoretzky, A. Wald, and J. Wolfowitz, "Relations among certain ranges of vector measures," Pac. J. Math. 1 (1951), 59-74.
- [13] R. Engelbrecht-Wiggans, P. Milgrom, and R.J. Weber, "Competitive bidding with proprietary information," CMSEMS Discussion Paper No. 465 (1981), Northwestern University.
- [14] I. Glicksberg, "A further generalization of Kakutani's fixed point theorem with application to Nash equilibrium points," Proc. Nat. Acad. Sci. (U.S.A.) 38 (1952), 170-172.
- [15] J.C. Harsanyi, "Game with incomplete information played by Bayesian players," Management Science 14 (1967-8), 159-182, 320-334, 486-502.

- [16] J.C. Harsanyi, "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points," International Journal of Game Theory 2 (1973), 1-23.
- [17] D.M. Kreps and R. Wilson, "Sequential equilibria," mimeo (1980), Stanford University.
- [18] D.M. Kreps and R. Wilson, "On the Chain-Store Paradox and predation: Reputation for 'toughness'," mimeo (1980), Stanford University.
- [19] Ky Fan, "Fixed-point and minimax theorems in locally convex topological linear spaces," Proc. Nat. Acad. Sci. (U.S.A.) 38 (1952), 121-126.
- [20] P. Milgrom, The Structure of Information in Competitive Bidding (Ph.D. dissertation, Stanford University, 1978), Garland Publ. Co., New York (1979).
- [21] P. Milgrom, "Rational expectations, information acquisition, and competitive bidding," Econometrica (to appear).
- [22] P. Milgrom and D.J. Roberts, "Equilibrium limit pricing doesn't limit entry," CMSEMS Discussion Paper No. 399 (1980), Northwestern University.
- [23] P. Milgrom and D.J. Roberts, "Predation, reputation, and entry deterrence," CMSEMS Discussion Paper No. 427 (1980), Northwestern University.
- [24] P. Milgrom and R.J. Weber, "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 (to appear).
- [25] P. Milgrom and R.J. Weber, "A theory of auctions and competitive bidding," CMSEMS Discussion Paper No. 447R (1981), Northwestern University.
- [26] P. Milgrom and R.J. Weber, "The value of information in a sealed-bid auction," CMSEMS Discussion Paper No. 463 (1981).
- [27] R. Radner and R.W. Rosenthal, "Private information and pure-strategy equilibria," Math. Oper. Res. (to appear).
- [28] R. Selten, "Re-examination of the perfectness concept for equilibrium points in extensive games," Int. J. Game Theory 4 (1975), 22-55.

- [29] R. Selten, "The Chain-Store Paradox," Theory and Decision 9 (1978), 127-159.
- [30] W. Vickrey, "Counterspeculation, auctions, and competitive sealed-tenders," Journal of Finance 16 (1961), 3-37.
- [31] R. Wilson, "A bidding model of perfect competition," Review of Economic Studies 44 (1977), 511-518.

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