GAME THEORY AND MATHEMATICAL ECONOMICS
O. Moeschlin, D. Pallaschke (editors)

TOPOLOGIES ON INFORMATION AND STRATEGIES IN GAMES WITH INCOMPLETE INFORMATION

Paul Milgrom and Robert J. Weber
J.L. Kellogg Graduate School of Management
Northwestern University
Evanston, Illinois U.S.A.

Introduction. Recently, several authors have dealt with the issue of the existence of equilibrium points in normal-form games with incomplete information. Their papers have had different focuses. Armbruster and Böge [AB] investigated a variety of non-standard equilibrium notions for games in which the players have inconsistent beliefs, and treated the classical Nash equilibrium point primarily for reasons of contrast. Rosenthal and Radner [RR] studied situations in which the mixed-strategy equilibria of a game correspond to "equivalent" pure-strategy equilibria (the occurrence of these situations is related to the nonatomicity of the game's information structure). Milgrom and Weber [MW] were concerned with the manner in which the equilibria of a game vary as the information structure is changed.

The most general of the game models considered was that developed in [MW], wherein a general existence theorem was also obtained. Our purpose in this paper is to discuss the general model, and to explore the relationships among the three existence theorems mentioned above. A central feature of this paper is the treatment of topological issues concerning the information in the game and the players' strategies.

Definitions and examples. In an n-person game with incomplete information, an initial chance move determines an n-tuple of player "types." Each player is told his own type, and must then select from a set of permissible "actions." Finally, each receives a payoff depending on the n-tuples of types and selected actions. The formal model consists of six elements, all of which are common knowledge among the players:

(i) a set $N = \{1,2,\dots,n\}$ of players;
(ii) for each player i, a compact metric space A_i of actions;
(iii) for each player i, a measurable space T_i of types;
(iv) a measurable space T_0 of states.

Let $A \equiv A_1 \times\dots\times A_n$ and $T \equiv T_0 \times\dots\times T_n$. The next element of the model is:

(v) the information structure of the game, a probability measure η on T.

The marginal distribution of η on player i's type space is denoted by η_i ; that is, for every measurable set S in T_i, $\eta_i(S) \equiv \eta(T_0 \times \ldots \times T_{i-1} \times S \times T_{i+1} \times \ldots \times T_n)$. The final element of the model is:

(vi) for each player i, a payoff function u_i: $T \times A \to R$, which is bounded, measurable on T, and continuous on A.

Example 1 (preference uncertainty). Assume that each player in a game knows his own preferences, but is uncertain about the preferences (and hence, about the strategic motivations) of his competitors. Harsanyi proposed that this situation be modeled as a game of incomplete information, in which each payoff function can be written in the form $u_i(t_i,a)$. The realization of the random variable $\tilde{t}_i$ is player i's "type," known to him but unknown to the other players. (We are dealing here with the case in which the beliefs of the players are consistent.) Frequently, it is further assumed that the types of the players are distributed independently, i.e., $\eta = \eta_1 \times \ldots \times \eta_n$. It is this independent-types, preference-uncertainty model which is the principal object of study in both [AB] and [RR].

When i's type is t_i, his payoff function $u_i(t_i,\cdot)$ is a von Neumann-Morgenstern utility function representing his preferences. Such utility functions are determined only up to positive affine transformations. Therefore, it is desirable that the analysis of the game should be independent of the particular choice of utility representations. That is, the game with payoff functions $v_i(t_i,a) = \alpha(t_i)u_i(t_i,a) + \beta(t_i)$, where α and β are measurable functions and $\alpha > 0$, should be treated as equivalent to the original game.

Example 2 (quality uncertainty). Assume that the preferences of the players are known to all, but that the payoffs are affected by some chance event represented by the random variable $\tilde{t}_0$; that is, each payoff function can be written in the form $u_i(t_0,a)$. The variable $\tilde{t}_i$ represents a private signal received by player i prior to his choice of an action. Note that a player's signal may be informative about the signals of others, as well as about the chance event, through the joint distribution η.

In this case, the expected payoff of a player, given that the vector of signals $t = (t_1,\ldots,t_n)$ has arisen and the players have selected the vector of actions a, is $v_i(t,a) = E[u_i(\tilde{t}_0,a)|t_1,\ldots,t_n]$.

Similarly, in any game in which the state variable appears in the payoff functions, that variable can be "integrated out" of the game, yielding an equivalent game from the standpoint of expected payoffs. (Actually, this corresponds to moving an unobservable chance move from the beginning of the extensive form of the game to the end, and then replacing terminal payoffs by expected payoffs.) Therefore, throughout the remainder of this paper we shall assume without loss of generality that the state variable does not appear in the payoff functions, and that $T = T_1 \times \ldots \times T_n$.

Strategies. A pure strategy for player i is a measurable function p_i: $T_i \to A_i$. It is well-known that even in the case of complete information (i.e., $|T| = 1$), many games fail to have equilibria if the players are restricted to using pure strategies. Therefore, we wish to extend the players' options to include randomized strategies.

In games of complete information, a randomized strategy is defined as a probability measure on a player's pure-strategy space, or equivalently, on his action space. This approach leads to technical difficulties when we try to extend it to general games with incomplete information (for a discussion of these difficulties, see [Au]). One alternative is to define a mixed strategy for player i to be a measurable function σ_i: $[0,1] \times T_i \to A_i$. The first argument of σ_i is interpreted as a uniformly distributed random variable $\tilde{s}$; hence, $\sigma_i(\tilde{s},t_i)$ is a random variable on A_i. The expected payoff to player i, when the players use mixed strategies $\sigma_1,\ldots,\sigma_n$, is

$$\bar{\pi}_i(\sigma_1,\ldots,\sigma_n) = \int u_i(t,[\sigma_1(s_1,t_1),\ldots,\sigma_n(s_n,t_n)])\, ds_1\ldots ds_n\, \eta(dt).$$

Another approach, developed in [MW], is to define a distributional strategy for player i to be a probability measure μ_i on $T_i \times A_i$, for which the marginal distribution on T_i is η_i (that is, for all measurable sets S in T_i, $\mu_i(S \times A_i) = \eta_i(S)$). We denote player i's set of distributional strategies by D_i. When the players adopt distributional strategies $\mu_1,\ldots,\mu_n$, the expected payoff to player i is

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

It is shown in [MW] that mixed and distributional strategies provide the players with identical strategic opportunities. However, each distributional strategy corresponds to many mixed strategies. For example, in the case of complete information, many mixed strategies induce the same distribution on actions; only a single distributional strategy is this distribution.

Equilibrium points. An n-tuple $(\mu_1,\ldots,\mu_n)$ of distributional strategies is an equilibrium point of a game if each player's strategy is a best response to the strategies of the others, i.e., if for every player i and every μ_i' in D_i, $\pi_i(\mu_1,\ldots,\mu_i,\ldots,\mu_n) \geq \pi_i(\mu_1,\ldots,\mu_i',\ldots,\mu_n)$.

A general theorem asserting the existence of equilibria for classes of games is due to Glicksberg:

Theorem. Let the players' strategy spaces be nonempty compact, convex linear Hausdorff 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.

In order to apply this theorem to games with incomplete information, it is necessary to topologize each space of distributional strategies in such a way as to resolve two conflicting tensions: the topology must be weak enough to make the strategy spaces compact, yet strong enough to make the expected payoff functions continuous. How might we do this?

Assume that each T_i is endowed with a complete separable metric topology, and that the measurable structure of T_i contains the Borel sets. Then a natural topology on D_i is the topology of weak convergence of probability measures.

Recall that a family of probability measures on a topological space is tight if for every $\varepsilon > 0$ there is a compact set K_ε on which each of the measures places probability at least $1 - \varepsilon$. Every single probability measure on a complete separable metric space is tight; hence, η_i is tight on T_i. A_i is compact, and for any K in T_i and every μ_i in D_i, $\mu_i(K \times A_i) = \eta_i(K)$; hence, D_i is a tight family and, by Prohorov's Theorem, is compact in the weak topology.

The existence theorem. In order to assure that the expected payoff functions are continuous in the players' distributional strategies when each D_i is given the weak topology, we make two further assumptions.

Continuous payoffs: For every player i and every $\varepsilon > 0$, there is a subset E of T such that $\eta(E) > 1 - \varepsilon$ and such that the family of functions $\{u_i(t,\cdot): t \in E\}$ is equicontinuous.

Continuous information: The measure η is absolutely continuous with respect to the product measure $\hat{\eta} \equiv \eta_1 \times\ldots\times \eta_n$.

The continuous payoffs condition holds, for example, when each u_i is uniformly continuous on $T \times A$, or when each A_i is finite. More generally, this condition permits us (via Lusin's Theorem) to find, for every player i and every $\varepsilon > 0$, a set K in T with $\eta(K) > 1 - \varepsilon$ and a bounded, continuous function v_i on $T \times A$ such that v_i coincides with u_i on $K \times A$.

The continuous information condition holds, for example, when the type spaces are finite, or when the types are independently distributed. Assuming the condition to hold, let f be the Radon-Nikodym derivative of η with respect to $\hat{\eta}$. Then the expected payoff functions can be rewritten in the following form:

$$\pi_i(\mu_1,\ldots,\mu_n) = \int u_i(t,a)\, f(t)\, d\mu_1(t_1,a_1)\ldots d\mu_n(t_n,a_n).$$

(One can define new payoff functions $v_1,\ldots,v_n$, by $v_i(t,a) = u_i(t,a)f(t)$. The expected payoff functions in the game with these payoffs, and with information structure $\hat{\eta}$, are the same as the expected payoff functions in the original game; hence, the equilibria of the two games are identical. Several authors have made use of this identity in the finite-types, finite-actions case, to enable them to restrict their attention to games with independent types.)

With these assumptions, it is shown in [MW] that the expected payoff functions are continuous when each D_i is given the weak topology.

Theorem. If a game with incomplete information has continuous payoffs and continuous information, then an equilibrium point exists.

The existence theorem of [AB] treats the case in which each type space is compact and the payoff functions are continuous in both types and actions; in this case, the continuous payoffs assumption is satisfied. Furthermore, they deal only with the case of independent types, so the continuous information assumption is satisfied as well. Hence, the above theorem encompasses their result.

In the existence theorem of [RR], on the other hand, no topology on the type spaces is assumed to be given. We discuss their approach in the next section.

The finite-actions, independent-preferences case. Assume that the type spaces are general measurable spaces, that the types are independent, that each A_i is finite, and that each u_i depends only on t_i and a.

(Actually, [RR] also assume only that each $u_i(\cdot,a)$ is η_i-integrable. However, using the rescaling discussed in Example 1, there is no loss of generality in assuming that these functions are bounded.) They endow each D_i with the topology of pointwise convergence: a (generalized) sequence $\{\mu_i^\alpha\}$ converges to μ_i if for every measurable C in $T_i \times A_i$, $\{\mu_i^\alpha(C)\}$ converges to $\mu_i(C)$. Interpreting this as a weak*-topology, they obtain the compactness of D_i via Alaoglu's Theorem, and demonstrate the continuity of the expected payoff functions.

An alternative approach is available in their setting. Consider the mapping $t_i \mapsto u_i(t_i,\cdot)$, which takes T_i into $R^{|A|}$. Topologize T_i with the Euclidean topology on $R^{|A|}$, identifying types which are mapped into the same point. (With respect to this topology, the payoff function u_i is continuous on $T_i \times A$. Furthermore, the measurability of each $u_i(\cdot,a)$ implies that all Borel subsets of T_i are measurable.) Use this topology to define the weak topology on D_i, and observe that the continuous payoffs and continuous information conditions are satisfied by the games under consideration. Existence of equilibria then follows from the theorem given in the previous section.

How do the two topologies defined on D_i compare? Up to identifications, they are identical. For any a_i in A_i and μ_i in D_i, let the partial distribution of μ_i conditioned on a_i be the measure $\mu_i[a_i]$ on T_i defined for all measurable S in T_i by $\mu_i[a_i](S) = \mu_i(S \times \{a_i\})$. A sequence of strategies converges weakly (when there are only finitely many actions available) if and only if the partial distributions conditioned on acts all converge weakly. And since all of these partial distributions must sum to, and therefore are bounded by, the marginal distribution η_i, weak convergence is equivalent to the pointwise convergence defined above.

Endogenously-defined topologies on types. For games of the type treated by [RR], it was possible to construct a "suitable" topology on T_i when none was originally given. How might this be done for more general games with incomplete information? A player's type has two features. It plays a direct role in his payoff function, and it affects his beliefs about the types of his competitors. We shall define a metric (actually, a pseudometric) topology on types which captures these two aspects separately.

For any player i and types t_i' and t_i'' in T_i, define

$$d_i^1(t_i',t_i'') = \sum_{j=1}^{n} \sup_{a\in A} \sup_{t_{-i}\in T_{-i}} |u_j(t_i',t_{-i},a) - u_j(t_i'',t_{-i},a)|$$

(where $T_{-i} \equiv T_1\times\cdots\times T_{i-1}\times T_{i+1}\times\cdots\times T_n$), 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. (The topology defined on T_i in the preceding section is the d_i^1-topology .)

For any player i and type t_i in T_i, let $\eta_{-i}(\cdot|t_i)$ denote the conditional distribution on T_{-i} induced by η. (We assume here that regular versions of conditional probability exist for η.) For any t_i' and t_i'' in T_i, define

$$d_i^2(t_i',t_i'') = \sup_{S\subset T_{-i}} |\eta_{-i}(S|t_i') - \eta_{-i}(S|t_i'')|,$$

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

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

When the players' types are independent, this metric is identically zero.

Given a metric d on T, a probability measure η on T is d-tight if all of the Borel sets with respect to d are measurable, and if for every $\varepsilon > 0$ there is a d-compact set of measure at least $1 - \varepsilon$.

Theorem.

(a) If η is d^1-tight, then the game has continuous payoffs.

(b) If η is d^2-tight, then the game has continuous information.

Proof. Part (a) is an immediate consequence of the Arzela -Ascoli Theorem (see [Bi], page 55). The remainder of the proof deals with part (b).

Since η is d^2-tight, η_1 is d_1^2-tight. Fix any $\varepsilon > 0$. From tightness it follows that there is a countable collection $\{B^k\}$ of disjoint measurable subsets of T_1 such that each B^k has d_1^2-diameter less than ε and $\Sigma\, \eta_1(B^k) = 1$. Let $\{t^k\}$ be a collection of points such that each t^k is in B^k. For each k, define

$$\eta^{\varepsilon,k} \equiv \eta_{-1}(\cdot|t^k) \times \eta_1|_{B^k}, \text{ where } \eta_1|_{B^k}(C) \equiv \eta_1(C\cap B^k),$$

and define $\eta^\varepsilon \equiv \Sigma\, \eta^{\varepsilon,k}$. For every S in T, $|\eta^\varepsilon(S) - \eta(S)| < \varepsilon$. Also, $\eta^\varepsilon \ll \eta_{-1}^\varepsilon \times \eta_1$.

Define $\bar{\eta} \equiv \Sigma\, 2^{-m} \cdot \eta^{1/m}$ Then $\eta << \bar{\eta}$, and $\bar{\eta} << \bar{\eta}_{-1} \times \eta_1$. Hence, η is absolutely continuous with respect to a product measure, and therefore (taking the Lebesgue decomposition of $\bar{\eta}_{-1}$ with respect to η_{-1}) is absolutely continuous with respect to $\eta_{-1} \times \eta_1$. The theorem follows by induction.

In view of this theorem, a natural topology on each T_i is that induced by the metric $d_i^1 + d_i^2$. Tightness of η with respect to the product topology induced by these metrics is sufficient to guarantee the existence of equilibrium points.

An example which sheds light on part (b) of the theorem is the following: Consider a two-person game, in which both players jointly observe a single random variable which is uniformly distributed on [0,1]. In this case, the continuous-information condition fails to hold; the d_i^2-topology is discrete, and η_i is not d_i^2-tight. Still, the existence of an equilibrium point follows by treating the games associated with different values of the random variable separately. It is not known whether an existence result can be obtained, for example, for three-player games in which there are three random variables, $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$, and the players observe $\tilde{x}$ and $\tilde{y}$, $\tilde{x}$ and $\tilde{z}$, and $\tilde{y}$ and $\tilde{z}$, respectively.

Acknowledgement. The preparation of this paper was supported by the Center for Mathematical Studies in Managerial Economics at Northwestern University, and by grants from the U.S. Office of Naval Research and the National Science Foundation.

References.

[AB] W. Armbruster and W. Böge, "Bayesian game theory," Game Theory and Related Topics (ed. O. Moeschlin and D. Pallaschke), North-Holland Publ. Co., 1979, 17-28.

[Au] R.J. Aumann, "Mixed and behavior strategies in infinite extensive games," Annals of Mathematics Studies No. 52, Princeton University Press, 1964, 627-650.

[Bi] P. Billingsley, Convergence of Probability Measures, John Wiley and Sons, 1968.

[Gl] 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, 170-174.

[MW] P. Milgrom and R.J. Weber, "Distributional strategies for games with incomplete information," CMSEMS Discussion Paper No. 428, Northwestern University, 1980, 47 pp.

[RR] R. Radner and R.W. Rosenthal, "Private information and pure strategy equilibria," Bell Laboratories Technical Report, 1980, 28 pp.