Discussion Paper No. 840

Correlated equilibria with payoff uncertainty*

Kevin D. Cotter
Department of Economics
Northwestern University

July 1989

Corresponding address: Department of Economics
Wayne State University
Detroit, MI 48202

Aumann's notion of correlated equilibrium is extended to games with payoff uncertainty. In this paper, an action correlated equilibrium is defined to be a probability distribution over types and actions which is consistent with the prior distribution over types and is self-fulfilling when each player observes its type and action from the distribution. This definition is broader, and mathematically simpler, than the one previously studied by the author, which was a correlated equilibrium for the *ex ante* game in behavioral strategies. The action correlated equilibrium correspondence is shown to be continuous with respect to the prior distribution of types, proving existence.

KEYWORDS: Game theory, Bayesian games, correlated equilibrium

1. Introduction

Some authors have claimed that the Nash equilibrium concept is too restrictive to include all outcomes of a game which follow from the hypothesis that the Bayesian rationality of players is common knowledge. Aumann (1974, 1987) introduced the notion of correlated equilibrium as a remedy to this problem. A correlated equilibrium allows players to base their actions on a set of observations, known as a correlation device, which are not included in the original description of the game. Aumann (1987) showed that a correlated equilibrium is a probability distribution over the set of players' actions such that each player optimally chooses to follow its observed action from that distribution. Similar notions of equilibrium with correlated strategies have been based on the hypothesis of common knowledge of rational behavior [Pierce (1984), Bernheim (1984), Brandenburger and Dekkel (1987)], which was also used by Aumann (1987) to motivate correlated equilibrium.

One question left unanswered by the above literature is how to define correlated equilibrium for games with payoff uncertainty. Without practical loss of generality each player's information about that uncertainty can be taken to be a signal, with its realization known as the player's *type*. Though Aumann's definition of correlated equilibrium includes a state space and a signal for each player, these are part of the equilibrium rather than the original description of the game. If there are states of nature which affect payoffs, however, then any correlated equilibrium must include those states.

The purpose of this paper is to provide a definition of correlated equilibrium for games with payoff uncertainty which preserves the Bayesian rationality of players. Another such definition has been previously discussed by Cotter (1989b). In that paper a correlated equilibrium was defined to be a correlated equilibrium for the game in type-dependent (behavioral) strategies. This solution concept was shown to have two advantages over the standard Bayesian-Nash equilibrium. First, the set of correlated equilibria is robust with respect to misspecifications of the underlying state space. Bayesian-Nash equilibria, unlike strategy correlated equilibria, are effected by the degree of correlation of player types, which in turn is very sensitive to small perturbations in the game. Second, the existence of Bayesian-Nash equilibrium, even in otherwise well-behaved games, is an open question, while correlated equilibria were shown to exist for such games.

Two problems arise with this definition, known in this paper as strategy correlated equilibrium. First, the set of behavioral strategy profiles is not generally compact with respect to any topology for which expected payoffs are continuous. This complicates the use of probability distributions on the set of strategy profiles, which is overcome at the expense of some advanced mathematics. Second, there may be additional opportunities for correlation which are not included in the definition of strategy correlated equilibrium. Since correlation is defined over behavioral strategies, any correlation of actions which does not result from the correlation of types must be independent of the type space.

In this paper, a further extension of correlated equilibrium for games with payoff uncertainty is introduced. An action correlated equilibrium is a probability distribution over the set of types and actions such that each player, drawing its type and action from the distribution, optimally chooses to follow that action. Action correlated equilibria have the same advantages over Bayesian-Nash equilibrium that were previously described for strategy correlated equilibria, and are mathematically simpler.

Both correlated equilibrium concepts can be supported by the use of a correlation device in addition to the types of players. The difference between them is that in a strategy correlated equilibrium, the correlation device must be independent of the types of players, i.e., the prior uncertainty in the model. By contrast, the correlation device is permitted to depend on players' types in an action correlated equilibrium. One interpretation of such a correlation device is that it represents players' information which was not properly specified in the model. If a correlation device is interpreted as a mechanism to which players report their types, however, then the equilibrium would require that players truthfully reveal their types to the mechanism. This incentive compatibility issue will not be discussed in this paper.

The outline of this paper is as follows. Section 2 provides an example of a game with a discussion of its strategy and action correlated equilibria, and presents some basic concepts for later reference. The formal model, with a review of known results, is presented in Section 3. Action correlated equilibrium is defined in Section 4, while the relationship between the prior beliefs about the game and the resulting set of equilibria is studied in Section 5. All proofs are in Section 6.

2. An example

Consider a game with two players, 1 and 2, and two states of nature, x and y. Player 1 believes both states are equally likely, while player 2 observes the state before play begins. Let $A_1 = \{U,D\}$ and $A_2 = \{L,R\}$ be the sets of possible actions for the players. The payoff matricies are

x	L	R
U	(1,1)	(0,0)
D	(0,0)	(1,1)

y	L	R
U	(0,0)	(1,1)
D	(1,1)	(0,0)

Following standard terminology, each player's *type* is its private observation about the state of nature. In this game player 1's type space $T_1 = \{t_1\}$ (say), while player 2's type space $T_2 = \{x,y\}$.

Aumann (1975, 1987) defined a correlated equilibrium for a game without payoff uncertainty to consist of a correlation device $(Z_1 \times Z_2, \eta)$, with η a probability distribution on $Z_1 \times Z_2$, and for each player i, a function $g_i : Z_i \rightarrow A_i$ such that $g_i(z_i)$ maximizes player i's expected payoff given z_i and the function g_{-i} chosen by the other player. Aumann (1987) showed that without loss of generality, $Z_i = A_i$ and g_i is the identity function. Therefore a correlated equilibrium is a probability distribution υ over $A_1 \times A_2$ such that each player optimally follows its observation a_i.

In this paper, a correlated equilibrium for games with payoff uncertainty is defined in the same way, except that each z_i must include player i's type. Assume without loss of generality that $Z_i = T_i \times A_i$ and g_i is the projection onto A_i. An *action correlated equilibrium* (a.c.e.) is a probability distribution σ over the set of types and actions, $\{x,y\} \times \{U,D\} \times \{L.R\}$[1], which satisfies the following conditions:

(A1) the marginal distribution of σ on $\{x,y\}$ equals the prior distribution,

(A2) each player, drawing its own type and action from the distribution, optimally follows that action.

Consider an a.c.e. $\sigma = (\sigma_{xUL}, \sigma_{xUR}, \sigma_{xDL}, \sigma_{xDR}, \sigma_{yUL}, \sigma_{yUR}, \sigma_{yDL}, \sigma_{yDR})$, with

$$\sigma_{xUL} = \sigma(x,U,L) \quad \sigma_{xUR} = \sigma(x,U,R) \quad \sigma_{xDL} = \sigma(x,D,L) \quad \sigma_{xDR} = \sigma(x,D,R)$$

$$\sigma_{yUL} = \sigma(y,U,L) \quad \sigma_{yUR} = \sigma(y,U,R) \quad \sigma_{yDL} = \sigma(y,D,L) \quad \sigma_{yDR} = \sigma(y,D,R)$$

Player 1 draws its action from the distribution, since its type space is degenerate, while player 2 draws its action and the state of nature. Condition (A1) can be written

$$\sigma_{xUL}+\sigma_{xUR}+\sigma_{xDL}+\sigma_{xDR} = \sigma_{yUL}+\sigma_{yUR}+\sigma_{yDL}+\sigma_{yDR} = 1/2. \tag{1}$$

Condition (A2) requires that if player 1 observes U, then the conditional expected payoff from choosing U is at least as large as from choosing D:

$$P[(x,L) \mid U] + P[(y,R) \mid U] \geq P[(x,R) \mid U] + P[(y,L) \mid U] \text{ so } \sigma_{xUL}+\sigma_{yUR} \geq \sigma_{xUR}+\sigma_{yUL} \tag{2a}$$

and conversely if D is observed:

$$P[(x,R)\mid D] + P[(y,L)\mid D] \geq P[(x,L)\mid D] + P[(y,R)\mid D] \text{ so } \sigma_{xDR}+\sigma_{yDL} \geq \sigma_{xDL}+\sigma_{yDR}. \quad (2b)$$

If player 2 observes (x,L) or (y,L), then the conditional expected payoff from choosing L must be at least as large as from choosing R:

$$P[U\mid(x,L)] \geq P[D\mid(x,L)] \qquad \text{so } \sigma_{xUL} \geq \sigma_{xDL} \qquad (2c)$$

$$P[D\mid(x,R)] \geq P[U\mid(x,R)] \qquad \text{so } \sigma_{xDR} \geq \sigma_{xUR} \qquad (2d)$$

and conversely if (x,R) or (y,R) is observed:

$$P[D\mid(y,L)] \geq P[U\mid(y,L)] \qquad \text{so } \sigma_{yDL} \geq \sigma_{yUL} \qquad (2e)$$

$$P[U\mid(y,R)] \geq P[D\mid(y,R)] \qquad \text{so } \sigma_{yUR} \geq \sigma_{yDR}. \qquad (2f)$$

Any σ satisfying (1-2) is an a.c.e. for this game.

Two special cases of a.c.e., both of which restrict the correlation device, have been studied in the literature. A *Bayesian-Nash equilibrium* (b.n.e.), introduced by Harsanyi (1967-8), is an a.c.e. for which Z_i is independent of Z_{-i} and T_{-i}. A b.n.e. satisfies (A1-A2) and the following:

(A3) Conditional on any profile of types, the actions of players 1 and 2 are independent.

(A4) For each player i, the distribution of its action given the profile of types of all players depends only on player i's type.

Since this equilibrium concept is well known, the details of calculating the b.n.e. for this game are omitted. They are the two pure strategy equilibria

$$\sigma_{xUL} = \sigma_{yUR} = 1/2 \qquad (3a)$$

$$\sigma_{xDR} = \sigma_{yDL} = 1/2 \tag{3b}$$

and all distributions of the form

$$\sigma_{xUL} = \sigma_{xDL} = \sigma_{yUL} = \sigma_{yDL} \qquad \sigma_{xUR} = \sigma_{xDR} = \sigma_{yUR} = \sigma_{yDR}. \tag{3c}$$

The other special case of an a.c.e., introduced by Cotter (1989a), is a *strategy correlated equilibrium* (s.c.e.). An s.c.e. is an a.c.e. for which the correlation device consists of the type space and a second correlation device which is independent of the type space. This restriction cannot be stated in terms of the distribution σ over types and actions, but requires the use of state-contingent behavioral strategies[2]. A behavioral strategy for player i is a function from its set of types to its actions[3]. Player 1's set of behavioral strategies is its set of actions $\{U,D\}$, while player 2's set is $\{(L_x,L_y),(L_x,R_y),(R_x,L_y),(R_x,R_y)\}$. A new game can be defined for which each player chooses a behavioral strategy and its payoff function is the expected payoff from both players' strategies. Cotter (1989b) proved that an s.c.e. is a correlated equilibrium for this game in behavioral strategies. Each player observes its strategy from the distribution and its type, then optimally chooses the corresponding action.

Consider an s.c.e. $p = (p_{ULL},p_{ULR},p_{URL},p_{URR},p_{DLL},p_{DLR},p_{DRL},p_{DRR})$ for the above game, with

$$p_{ULL} = p(U,L_x,L_y) \qquad p_{ULR} = p(U,L_x,R_y) \qquad p_{URL} = p(U,R_x,L_y) \qquad p_{URR} = p(U,R_x,R_y)$$

$$p_{DLL} = p(D,L_x,L_y) \qquad p_{DLR} = p(D,L_x,R_y) \qquad p_{DRL} = p(D,R_x,L_y) \qquad p_{DRR} = p(D,R_x,R_y)$$

Player 1 always follows its observed strategy if

$$\begin{aligned} &P[(L_x,L_y)\mid U] + 2P[(L_x,R_y)\mid U] + P[(R_x,R_y)\mid U] \\ &\qquad \geq P[(R_x,R_y)\mid U] + 2P[(R_x,L_y)\mid U] + P[(L_x,L_y)\mid U] \quad \text{so } p_{ULR} \geq p_{URL} \end{aligned} \tag{4a}$$

$$P[(R_x,R_y)\,|\,D]+2P[(R_x,L_y)\,|\,D]+P[(L_x,L_y)\,|\,D]$$
$$\geq P[(L_x,L_y)\,|\,D]+2P[(L_x,R_y)\,|\,D]+P[(R_x,R_y)\,|\,D] \qquad \text{so } p_{DRL} \geq p_{DLR} \qquad (4b)$$

while player 2 always follows its observed strategy if

$$P[U\,|\,(L_x,L_y)] \geq P[D\,|\,(L_x,L_y)] \quad \text{and} \quad P[D\,|\,(L_x,L_y)] \geq P[U\,|\,(L_x,L_y)] \quad \text{so } p_{ULL} = p_{DLL} \qquad (4c)$$

$$P[D\,|\,(R_x,R_y)] \geq P[U\,|\,(R_x,R_y)] \quad \text{and} \quad P[U\,|\,(R_x,R_y)] \geq P[D\,|\,(R_x,R_y)] \quad \text{so } p_{URR} = p_{DRR} \qquad (4d)$$

$$P[U\,|\,(L_x,R_y)] \geq P[D\,|\,(L_x,R_y)] \qquad \text{so } p_{ULR} \geq p_{DLR}. \qquad (4e)$$

$$P[D\,|\,(R_x,L_y)] \geq P[U\,|\,(R_x,L_y)] \qquad \text{so } p_{DRL} \geq p_{URL}. \qquad (4f)$$

Any p satisfying (4) is an s.c.e.

The distributions p and σ are related in the following way:

$$\sigma_{xUL} = (p_{ULL} + p_{ULR})/2 \qquad \sigma_{xUR} = (p_{URL} + p_{URR})/2$$

$$\sigma_{xDL} = (p_{DLL} + p_{DLR})/2 \qquad \sigma_{xDR} = (p_{DRL} + p_{DRR})/2$$

$$\sigma_{yUL} = (p_{ULL} + p_{URL})/2 \qquad \sigma_{yUR} = (p_{ULR} + p_{URR})/2$$

$$\sigma_{yDL} = (p_{DLL} + p_{DRL})/2 \qquad \sigma_{yDR} = (p_{DLR} + p_{DRR})/2 \qquad (5)$$

The difference between an a.c.e. and an s.c.e. is that in an s.c.e., each player observes its own type and behavioral strategy, as opposed to its own type and action in an a.c.e. The latter conveys less information since it is a many-to-one function of the former. There are several consequences of this distinction.

First, if p is an s.c.e., then the distribution σ given by (5) is an a.c.e.. To prove this, equation (1) holds automatically, and (2a,b) are implied by (4a,b) respectively. Equation (2c) is implied by (4c,e), (2d) by (4d,f), (2e) by (4c,f), and (2f) by (4d,e).

Second, given a distribution σ over types and actions, there need not exist a distribution p over behavioral strategies which satisfies (5). Condition (A4) is required, which implies the following restrictions:

$$P[U \mid x] = P[U \mid y] \qquad \text{so } \sigma_{xUL} + \sigma_{xUR} = \sigma_{yUL} + \sigma_{yUR} \qquad (6a)$$

$$P[D \mid x] = P[D \mid y] \qquad \text{so } \sigma_{xDL} + \sigma_{xDR} = \sigma_{yDL} + \sigma_{yDR} \qquad (6b)$$

Third, if σ satisfies (A1, A4), there will generally exist more than one p which satisfies (5). The reason is that there are eight equations in eight unknowns, but two linear dependencies, so the system is undetermined. Since $p \geq 0$ is also required, the solution may be unique in special cases, as shown below.

Finally, even if σ is an a.c.e. which satisfies (A4), there need not exist an s.c.e. p which satisfies (5). Consider the following a.c.e.:

$$\sigma_{xUL} = \sigma_{xDR} = 1/4 \quad \sigma_{xUR} = \sigma_{xDL} = 0 \quad \sigma_{yUL} = \sigma_{yUR} = \sigma_{yDL} = \sigma_{yDR} = 1/8. \qquad (7)$$

Despite the fact that σ satisfies (A1,A2,A4), there is no s.c.e. which generates it. The unique solution to (5) and $p \geq 0$ is

$$p_{ULL} = p_{ULR} = p_{DRL} = p_{DRR} = 1/4 \qquad p_{URL} = p_{URR} = p_{DLL} = p_{DLR} = 0. \qquad (8)$$

This p is not an s.c.e. since it violates (5c,d).

One objection to the a.c.e. given by (7) is that it cannot be interpreted as a mechanism to which players truthfully report their types. Player 2 can manipulate its information about the observed action by player 1 through its report to the mechanism. If player 2 reports x, then the mechanism would send (U,L) with probability 1/2 and (D,R) with probability 1/2, so player 2 would know which action was received by player 1. If player 2 reports y, then the mechanism would send each of (U,L), (U,R), (D,L), (D,R) with probability 1/4, so player 2 would not know which message player 1 received. Therefore player 2 should always report x.

Despite this failure of incentive compatibility, there are two reasons for studying action correlated equilibria. First, they reflect outcomes in games for which the state space may have been misspecified in the original model. Assuming that each player's type space may actually be $T_i \times Z_i$ rather than T_i, there is no reason to believe that Z_i must be independent of the types of other players. In this interpretation of an a.c.e., the mechanism already "knows" players' types, so it does not depend on truthful reporting of those types. A second reason is that when each player has a finite number of types, an a.c.e. is also a correlated equilibrium for the nonstochastic game constructed by Harsanyi where each type of each player is a distinct player. This statement does not hold for s.c.e.. Finally, a.c.e. are mathematically simpler objects than s.c.e., or for that matter, b.n.e..

3. The model

Consider a game with uncertainty and a finite set of players $I \equiv \{1,\ldots,I\}$. To economize notation I denotes both the set and the number of players, and $i \in I$ is a generic player. Each i has an

action space A_i, a compact metric space.

privately observed type $t_i \in T_i$, a complete separable metric space.

payoff function $u_i: T \times A \to \Re$, where $A = \prod_{i \in I} A_i$ and $T = \prod_{i \in I} T_i$.

For any metric space X, let $\Delta(X)$ be the set of probability[4] measures on X with the usual topology of weak convergence. By Theorems II.6.2 and II.6.4 of Parthasarathy (1967), $\Delta(X)$ is a compact (resp. separable) metric space if and only if X is compact (resp. separable). The information of players about the types of others is given by an

information structure $\mu \in \Delta(T)$.

Let μ_i be the marginal of μ on T_i. The assumptions about the payoff function are straightforward.

Assumption 3.1: For each i,

(a) the mapping $t \to u_i(\cdot,a)$ is measurable for each $a \in A$,

(b) the mapping $a \to u_i(t,\cdot)$ is continuous for each $t \in T$,

(c) the mapping $t \to \sup_{a \in A} |u_i(\cdot,a)|$ is integrable.

The standard way of modelling equilibrium for this game has been to transform it into a nonstochastic game in type-dependent strategies, then apply standard equilibrium concepts to the transformed game. Player i's strategy as a function of t_i can be defined as either a

distributional strategy $\sigma_i \in \Delta(T_i \times A_i)$ such that the marginal of σ_i on T_i is μ_i, or a

behavioral strategy $s_i : T_i \to \Delta(A_i)$ measurable.

Distributional strategies have been used by Milgrom and Weber (1985), while behavioral strategies have been studied by Radner and Rosenthal (1982), Balder (1986), and Cotter (1989a).

Given a distributional strategy σ_i, there exists, by Theorem V.8.1 of Parthasarathy (1967) and the fact that T_i is a complete separable metric space, a regular conditional probability distribution $\sigma_i[\cdot \mid t_i]$ on A_i given t_i, unique up to a.e. equivalence. Therefore any distributional strategy is equivalent to a unique behavioral strategy. The converse result also holds. These statements are formalized below[5].

Theorem 3.2: (a) For any distributional strategy σ_i there exists a unique (up to a.e. equivalence) behavioral strategy s_i such that $s_i(t_i)(B_i) = \sigma_i[B_i \mid t_i]$ for measurable $B_i \subset A_i$ and a.e. t_i.

(*b*) *For any behavioral strategy* s_i *there exists a unique distributional strategy* σ_i *such that* $\sigma_i(W_i \times B_i) = \int_{W_i} s_i(t_i)(B_i)\mu_i(dt_i)$ *for measurable* $W_i \subset T_i$ *and* $B_i \subset A_i$.

Let S_i be the set of distributional (or behavioral) strategies for player *i*, and $S^p = \prod_{i \in I} S_i$. Define $U_i:S^p \to \Re$ to be the expected payoff to player *i*:

$$U_i(s_1,\ldots,s_I) = \int_T \int_{A_I} \ldots \int_{A_1} u_i(t,a) s_1(t_1)(da_1)\ldots s_I(t_I)(da_I)\mu(dt). \tag{9a}$$

$$U_i(\sigma_1,\ldots,\sigma_I) = \int_T \int_{A_I} \ldots \int_{A_1} u_i(t,a)\sigma_1[da_1 \mid t_1]\ldots\sigma_I[da_I \mid t_I]\mu(dt). \tag{9b}$$

In effect, U_i is defined in terms of behavioral strategies regardless of whether behavioral or distributional strategies are used.

The original game with payoff uncertainty can therefore be transformed into a nonstochastic game for which player *i*'s payoff function and strategy space are U_i and S_i respectively. Most definitions of equilibrium for Bayesian games have been constructed by applying standard equilibrium concepts to the transformed game. Two such definitions are given below.

Use the convention that for any $s \in S^p$ (resp. $a \in A$, $t \in T$), s_{-i} (resp. a_{-i}, t_{-i}) is the profile of strategies (resp. actions, types) of players other than *i*.

Definition 3.3: A *Bayesian-Nash equilibrium* [Milgrom and Weber (1985), Radner and Rosenthal (1982)] is a Nash equilibrium for the transformed game, i.e., $s^* \in S^p$ such that for each *i* and $s_i \in S_i$, $U_i(s^*) \geq U_i(s_i, s^*_{-i})$.

Definition 3.4: A *strategy correlated equilibrium* [Cotter (1989b)] is a correlated equilibrium for the transformed game, i.e., $\upsilon \in \Delta(S^p)$ such that for each *i* and measurable $\delta_i:S_i \to S_i$,

$$\int_{S^p} U_i(s)\upsilon(ds) \geq \int_{S^p} U_i(\delta_i(s_i), s_{-i})\upsilon(ds). \tag{10}$$

Equivalently [Cotter (1989b)], an s.c.e. is a product space of separable metric spaces $Z = \prod_{i \in I} Z_i$, a probability distribution $\eta \in \Delta(Z)$, and for each i, a measurable function $r_i^*: T_i \times Z_i \to \Delta(A_i)$ such that for all measurable $r_i: T_i \times Z_i \to \Delta(A_i)$,

$$\int_Z \int_T \int_A u_i(t,a)[r_i^*(t_i,z_i)(da_i) - r_i(t_i,z_i)(da_i)] r_{-i}^*(t_{-i},z_{-i})(da_{-i}) \mu(dt) \eta(dz) \geq 0. \tag{11}$$

Equation (11) states that an s.c.e. is a b.n.e. for the game where each player's type space is $T_i \times Z_i$ and the information structure on $T \times Z$ is $\mu \times \eta$. Therefore an s.c.e. allows players to base their actions on observations which are not part of the original model and independent of the types of other players.

4. Action correlated equilibria

A major problem with solution concepts that rely on the transformed game $\{(U_i, S_i) \mid i \in I\}$ is that the expected payoff function U_i is not generally continuous with respect to any metric on S^p for which S^p is compact. For example, give S_i the topology of weak convergence as a subset of $\Delta(T_i \times A_i)$, and S^p the product topology. Then S_i and S^p are compact metric spaces but U_i is not continuous with respect to S^p. An example was constructed by Milgrom and Weber (1985, Example 2). See Cotter (1989a) for a more detailed interpretation of this problem.

As a consequence, existence of Bayesian-Nash equilibria for this model remains an open question. Existence of strategy correlated equilibria was proven by Cotter (1989b), but at the expense of some advanced mathematics which makes the concept more difficult to apply.

To motivate an alternate definition of correlated equilibrium, suppose players use an arbitrary correlation device which includes their type spaces, as described in Section 2.

Definition 4.1: A *joint distributional strategy* is $\sigma \in \Delta(T \times A)$ such that the marginal distribution of σ on T is μ.

Definition 4.2: A *joint behavioral strategy* is a measurable $s{:}T \to \Delta(A)$.

By Theorem 3.2, there is a one-to-one equivalence between joint distributional and joint behavioral strategies.

Definition 4.3: An *action correlated equilibrium* (a.c.e.) is a joint distributional strategy σ, or equivalently, a joint behavioral strategy s, such that for each i and measurable function $\alpha_i{:}T_i \times A_i \to A_i$,

$$\int_{T\times A} u_i(t,a)\sigma(dt\times da) \geq \int_{T\times A} u_i(t,\alpha_i(t_i,a_i),a_{-i})\sigma(dt\times da), \text{ or} \tag{12a}$$

$$\int_{T\times A} u_i(t,a)s(t)(da)\mu(d\omega) \geq \int_{T\times A} u_i(t,\alpha_i(t_i,a_i),a_{-i})s(t)(da)\mu(d\omega). \tag{12b}$$

Let υ be an s.c.e.. Player i observes (t_i,s_i) from the distribution on $T\times S^p$, then optimally chooses $s_i(t_i)$. Consider the mapping $\Phi{:}T\times S^p \to T\times A$ defined by $\Phi(t_1,\ldots,t_I,s_1,\ldots,s_I) = (t_1,\ldots,t_I,s_1(t_1),\ldots,s_I(t_I))$, and the resulting distribution σ on $T\times A$. It is easy to verify that σ is defined by, for $W\subset T$ and $B_i \subset A_i$ for each i,

$$\sigma(W\times B_1\times\ldots\times B_I) = \int_{S^p} [\int_W s_1(t_1)(B_1)\ldots s_I(t_I)(B_I)]\upsilon(ds). \tag{13}$$

If each player i were to observe $(t_i,s_i(t_i))$ rather than (t_i,s_i), then its observed action would still be optimal since less information is conveyed. This proves the following result.

Theorem 4.4: If υ is an s.c.e. and σ is the joint distributional strategy defined by (13), *then σ is an a.c.e. In particular, every Bayesian-Nash equilibrium is an a.c.e.*

Since the mapping Φ is many-to-one, there are in general many distributions on S^p which generate a particular joint distributional strategy, and not every a.c.e. can be generated by an s.c.e.. This was demonstrated in Section 2.

Another way of explaining the difference between action and strategy correlated equilibria is that an a.c.e. satisfies equation (11) except that t and z need not be independent. This is seen by taking $Z_i = A_i$.

5. Similarity of game characteristics

It is useful to know the extent to which the equilibria of the game depends on its underlying characteristics. Let G be the set of parameters of the game, and $\Delta(T \times A)$ the set of possible outcomes. The dependence of the game on its parameters can be posed in terms of the equilibrium correspondence $\xi: G \to\to \Delta(T \times A)$. The most important such properties are upperhemicontinuity and nonempty-valuedness. Though many possible definitions of G exist, in this section attention will be restricted to $G = \Delta(T)$, the set of information structures μ on T. Consider a sequence $\{\mu^n\} \subset \Delta(T)$. Milgrom and Weber (1985) required some complicated requirements on $\{\mu^n\}$ to obtain convergence of a corresponding sequence of Bayesian-Nash equilibria. Therefore the Bayesian-Nash equilibrium correspondence has closed graph with respect to a complicated topology on $\Delta(T)$. In this section, the action equilibrium correspondence is shown to be *continuous*, as well as nonempty and compact valued.

The following lemma provides a test for a.c.e. which may be easier to verify than Definition 4.3.

Lemma 5.1: A joint distributional strategy $\sigma \in \Delta(T \times A)$ *is an a.c.e. if and only if for all continuous functions* $\delta_i^c: T_i \times A_i \to \Delta(A_i)$,

$$\int_T \int_A u_i(t,a)\sigma(dt \times da) \geq \int_T \int_A \int_{A_i} u_i(t,\alpha_i,a_{-i})\delta_i^c(t_i,a_i)(d\alpha_i)\sigma(dt \times da). \qquad (14)$$

In addition, if for each i, A_i is a convex subset of $\Re^{\ell}$ for some ℓ, then σ is an a.c.e. if and only if (12a), *or equivalently* (12b), *holds for all continuous α_i.*

The following assumption is required to ensure that expected payoffs are defined for all probability measures on T.

Assumption 5.2: In addition to Assumption 3.1, u_i is uniformly bounded.

The main results of this section follows.

Theorem 5.3: The action equilibrium correspondence $\xi:\Delta(T) \to\to \Delta(T\times A)$ is upperhemicontinuous. In addition, ξ has convex graph, i.e., if $\mu, \mu' \in \Delta(T)$ and $\sigma \in \xi(\mu)$, $\sigma' \in \xi(\mu')$, then $\lambda\sigma + (1-\lambda)\sigma' \in \xi(\lambda\mu + (1-\lambda)\mu')$ for all $0 \leq \lambda \leq 1$.

Using the fact that the set of probability measures with finite support is dense in $\Delta(T)$ and that any such information structure has a Bayesian-Nash equilibrium, the following result is immediate.

Corollary 5.4: An a.c.e. exists, so ξ is nonempty-valued.

A consequence of Theorem 5.3 and Corollary 5.4 is the following.

Theorem 5.5: The action equilibrium correspondence ξ is lowerhemicontinuous.

6. Proofs

A key technical result is the following.

Lemma 6.1: For each $\eta \in \Delta(A)$, the mapping $t \to \int_A u_i(t,a)\eta(da)$ is integrable.

Proof: By Theorem II.6.3 of Parthasarathy (1967), there exists a sequence $\{\eta^n\} \subset \Delta(A)$, each with finite support, that converges to η. By Assumption 2.1(a),

$\int_A u_i(t,a)\eta^n(da)$ is integrable. Since $\{\int_A u_i(t,a)\eta^n(da)\}$ converges to $\int_A u_i(t,a)\eta(da)$ for each t, the latter is integrable by Assumption 2.1(c) and the dominated convergence theorem. $\therefore$

Proof of Theorem 3.2: Given a distributional strategy σ_i there exists, by Theorem V.8.1 of Parthasarathy (1967) and the fact that T_i is a complete separable metric space, a regular conditional probability distribution $\sigma_i[\cdot \mid t_i]$ on A_i given t_i, unique up to null sets. Define $s_i: T_i \to \Delta(A_i)$ as in (a). Given a continuous $c_i: A_i \to \Re$ and $d \in \Re$, let $O_i = \{\eta_i \in \Delta(A_i) \mid \int_{A_i} c_i(a_i)\eta_i(da_i) < d\}$. Then $s_i^{-1}(O_i) = \{t_i \mid \int_{A_i} c_i(a_i)\sigma_i[da_i \mid t_i] < d\}$ which is a measurable subset of T_i since $\sigma_i[\cdot \mid t_i]$ is regular. Since all sets of the form O_i generate the weak topology on $\Delta(A_i)$, s_i is measurable. This proves (a).

Conversely, let s_i be a behavioral strategy. Given an closed set $K_i \subset A_i$, let $K_i^n = \{a_i \in A_i \mid \rho_i(a_i, K_i) < 1/n\}$, which is open [Parthasarathy (1967, Theorem I.1.1)]. By the Tietze extension theorem [Munkres (1975, Theorem 3.2, p. 212)] there exists $c_i^n \in C_i(A_i)$ with $0 \le c_i^n \le 1$, $c_i^n = 0$ outside K_i^n, and $c_i^n = 1$ on K_i. Then for every $t_i \in T_i$, $s_i(t_i)(K_i) \le \int_{A_i} c_i^n(a_i)s_i(t_i)(da_i) \le s_i(t_i)(K_i) + s_i(t_i)(K_i^n \backslash K_i)$. Since $\{K_i^n \backslash K_i\}$ is decreasing to the empty set, $\lim_{n\to\infty} s_i(t_i)(K_i^n \backslash K_i) = 0$ so $s_i(t_i)(K_i) = \lim_{n\to\infty} \int_{A_i} c_i^n(a_i)s_i(t_i)(da_i)$. Therefore the map $t_i \to s_i(t_i)(K_i)$ is measurable as the limit of measurable mappings. Now let $\mathcal{A}_i = \{B_i \subset A_i \mid B_i$ is measurable and the map $t_i \to s_i(t_i)(B_i)$ is measurable$\}$. As just shown, $\mathcal{A}_i$ contains the closed sets, and is obviously closed under complementation and finite unions, and contains the empty set. Now let $\{B_i^n\} \subset \mathcal{A}_i$ be increasing to B_i, a measurable subset of A_i. Then for every $t_i \in T_i$, $s_i(t_i)(B_i) = \lim_{n\to\infty} s_i(t_i)(B_i^n)$, so $B_i \in \mathcal{A}_i$.

Therefore $\mathcal{A}_i$ is a monotone class containing the closed sets, which is the Borel sets of A_i, so $t_i \to s_i(t_i)(B_i)$ is measurable for every measurable $B_i \subset A_i$. Defining σ_i as in (b) completes the proof. $\therefore$

Proof of Theorem 4.4: Let $\alpha_i{:}T_i{\times}A_i \to A_i$ be measurable and define $\delta_i{:}T_i{\times}A_i \to \Delta(A_i)$ such that for all (t_i,a_i), $\delta_i(t_i,a_i)$ places probability one on $\alpha_i(t_i,a_i)$. Then

$$\int_T\int_A u_i(t,\alpha_i(t_i,a_i),a_{-i})\sigma(dt\times da) = \int_T\int_A\int_{A_i} u_i(t,\alpha_i,a_{-i})\delta_i(t_i,a_i)(d\alpha_i)\sigma(dt\times da)$$

$$= \int_{S^p}[\int_T\int_{A_{-i}}\int_{A_i}\int_{A_i} u_i(t,\alpha_i,a_{-i})\delta_i(t_i,a_i)(d\alpha_i)s_i(t_i)(da_i)s_{-i}(t_{-i})(da_{-i})\mu(dt)]\upsilon(ds)$$

$$\leq \int_{S^p}[\int_T\int_{A_{-i}}\int_{A_i} u_i(t,a)s_i(t_i)(da_i)s_{-i}(t_{-i})(da_{-i})\mu(dt)]\upsilon(ds)$$

$$= \int_T\int_A u_i(t,a)\sigma(dt\times da) \tag{15}$$

which completes the proof. $\therefore$

The following intermediate results will be needed to prove Lemma 5.1:

Lemma 6.2: A joint distributional strategy $\sigma \in \Delta(T{\times}A)$ is an a.c.e. if and only if for each i and measurable $\delta_i{:}T_i{\times}A_i \to \Delta(A_i)$,

$$\int_T\int_A u_i(t,a)\sigma(dt\times da) \geq \int_T\int_A\int_{A_i} u_i(t,\alpha_i,a_{-i})\delta_i(t_i,a_i)(d\alpha_i)\sigma(dt\times da). \tag{16}$$

Proof: Sufficiency is obvious. To prove necessity, let σ be an a.c.e. and $\delta_i{:}T_i{\times}A_i \to \Delta(A_i)$ be measurable. Let $\phi_i{:}T_i{\times}A_i{\times}A_i \to \Re$ be defined by

$$\phi_i(t_i,a_i,\alpha_i) = \int_{T_{-i}}\int_{A_{-i}} u_i(t,\alpha_i,a_{-i})\sigma_{-i}[dt_{-i}\times da_{-i} \mid t_i,a_i] \tag{17}$$

where $\sigma_{-i}[dt_{-i}\times da_{-i} \mid t_i,a_i]$ is the conditional probability distribution on $T_{-i}\times A_{-i}$. By hypothesis, $\phi_i(t_i,a_i,\alpha_i) \le \phi_i(t_i,a_i,a_i)$ for a.e. (t_i,a_i). Then

$$\int_T\int_A\int_{A_i} u_i(t,\alpha_i,a_{-i})\delta_i(t_i,a_i)(d\alpha_i)\sigma(dt\times da) = \int_{T_i}\int_{A_i}\int_{A_i}\phi_i(t_i,a_i,\alpha_i)\delta_i(t_i,a_i)(d\alpha_i)\sigma_i(dt_i\times da_i)$$

$$\le \int_{T_i}\int_{A_i}\phi_i(t_i,a_i,a_i)\sigma_i(dt_i\times da_i) = \int_T\int_A u_i(t,a)\sigma(dt\times da) \tag{18}$$

completing the proof. ∴

Lemma 6.3: Let X be a compact metric space and Z be a metric space. Let $Y \subset Z$ be closed, and let $f:Y \to \Delta(X)$ be continuous. Then there exists a continuous $f^*:Z \to \Delta(X)$ such that $f^*(y) = f(y)$ for every $y \in Y$.

Proof: Let $\{c_1,c_2,\ldots\}$ be a countable dense subset of $C(X)$, the set of real-valued continuous functions on X with the uniform convergence metric. Then the mapping $\Phi:\Delta(X) \to [0,1]^\infty$ defined by $\Phi(\eta) = (\int_X c_1(x)\eta(dx),\int_X c_2(x)\eta(dx),\ldots)$ is a homeomorphism of $\Delta(X)$ onto a compact convex subset of $[0,1]^\infty$ [Parthasarathy (1967, proof of Theorem II.6.4)]. Identify $\Delta(X)$ with its image under Φ. Let $g:[-1,1] \to [0,1]$ be continuous, strictly convex, and satisfy $g(b) = 0$ if and only if $b = 0$, and let $G:[0,1]^\infty \to [0,1]$ be $G(w) = \sum_{n=1}^{\infty} 2^{-n}g(w^n)$ where w^n is the n^{th} component of w. Note that G is continuous. Let $\eta_1, \eta_2 \in \Delta(X)$ satisfy $\eta_1 \neq \eta_2$ and $G(w-\eta_1) = G(w-\eta_2)$. Then for $\lambda \in (0,1)$, $G(w-\lambda\eta_1-(1-\lambda)\eta_2) < \lambda G(w-\eta_1) + (1-\lambda)G(w-\eta_2) = G(w-\eta_1)$. Therefore for fixed w, the minimum of $\{G(w-\eta) \mid \eta\in \Delta(X)\}$ exists and is unique. Define $r:[0,1]^\infty \to \Delta(X)$ to be $r(w) = argmin\{G(w-\eta) \mid \eta\in \Delta(X)\}$. By the above arguments, r is defined everywhere, continuous, and maps $\Delta(X)$ into itself. The result follows from Munkres (1975, Exercise 11, p. 221). ∴

Proof of Lemma 5.1: Necessity follows from Lemma 6.2. To prove sufficiency, suppose (14) holds for continuous δ_i^c, and let $\delta_i{:}T_i \times A_i \to \Delta(A_i)$ be measurable. Choose $\varepsilon > 0$, and let γ be such that for all $M_i \subset T_i \times A_i$ with $\sigma(M_i) < \gamma$,

$$\left| \int_{M_i} \int_{A_i} \phi_i(t_i,a_i,\alpha_i)\delta_i(t_i,a_i)(d\alpha_i)\sigma_i(dt_i \times da_i) \right| < \varepsilon. \tag{19}$$

By Lusin's theorem [Royden (1968, Exercise 3.31)], there exists an open subset $M_i \subset T_i \times A_i$ with $\sigma_i(M_i) < \gamma$ such that δ_i restricted to the complement of M_i is continuous. By Lemma 6.2 there exists a continuous $\delta_i^c{:}T_i \times A_i \to \Delta(A_i)$ which equals δ_i outside of M_i. Then

$$\begin{aligned}\int_T \int_A u_i(t,a)\sigma(dt \times da) &\geq \int_{T_i} \int_{A_i} \phi_i(t_i,a_i,a_i)\sigma_i(dt_i \times da_i) \\ &\geq \int_{T_i} \int_{A_i} \int_{A_i} \phi_i(t_i,a_i,a_i)\delta_i^c(t_i,a_i)(d\alpha_i)\sigma_i(dt_i \times da_i) \\ &\geq \int_{T_i} \int_{A_i} \int_{A_i} \phi_i(t_i,a_i,a_i)\delta_i(t_i,a_i)(d\alpha_i)\sigma_i(dt_i \times da_i) - \varepsilon \\ &= \int_T \int_A \int_A u_i(t,\alpha_i,a_{-i})\delta_i(t_i,a_i)(d\alpha_i)\sigma(dt \times da) - \varepsilon.\end{aligned} \tag{20}$$

Taking $\varepsilon \to 0$ completes the proof of the first statement. If A_i is a convex subset of Euclidean space for each i, then the same proof can be applied to any function $\alpha_i{:}T_i \times A_i \to A_i$. ∴

Proof of Theorem 5.3: That ξ has convex graph is obvious. Let $\{\mu^n\} \subset \Delta(T)$ be a sequence converging to $\mu \in \Delta(T)$, and for each n, let σ^n be an a.c.e. for μ^n. Then given $\varepsilon > 0$ there exists a compact $K_\varepsilon \subset T$ such that $\mu^n(K_\varepsilon) < \varepsilon$ for every n

[Parthasarathy (1967, Theorem II.6.7)]. Therefore $\sigma^n(K_\varepsilon \times A) = \mu^n(K_\varepsilon) < \varepsilon$ for each n, so by the same theorem, $\{\sigma^n\}$ is relatively compact in $\Delta(T\times A)$. Therefore a subsequence of $\{\sigma^n\}$ converges to, say, $\sigma \in \Delta(T\times A)$. Renumbering if necessary, denote the subsequence $\{\sigma^n\}$.

The equilibrium correspondence ξ is obviously convex-valued, so it remains only to be shown that σ is an a.c.e. for μ. It is easily verified that σ is a joint distributional strategy for μ. Finally, to verify (14), choose i, and let $\delta_i^c: T_i \times A_i \to \Delta(A_i)$ be continuous. Given ε, for sufficiently large n,

$$\begin{aligned}\int_T\int_A u_i(t,a)\sigma(dt\times da) &\geq \int_T\int_A u_i(t,a)\sigma^n(dt\times da) - \varepsilon \\ &\geq \int_T\int_A\int_{A_i} u_i(t,a_i,a_{-i})\delta_i^c(t_i,a_i)(d\alpha_i)\sigma^n(dt\times da) - \varepsilon \\ &\geq \int_T\int_A\int_{A_i} u_i(t,a_i,a_{-i})\delta_i^c(t_i,a_i)(d\alpha_i)\sigma(dt\times da) - 2\varepsilon. \qquad (21)\end{aligned}$$

Equation (14) follows by letting $\varepsilon \to 0$, proving that ξ is upperhemicontinuous. ∴

Proof of Theorem 5.5 First the domain of ξ must be extended to include signed measures. For any metric space X let $\tilde{\Delta}(X)$ be the set of signed measures $\tilde{\mu}$ on X that are countably additive, regular, and are of bounded variation, such that $\tilde{\mu}(X) = 1$. Let $\hat{\Delta}(T\times A) = \{\hat{\sigma} \in \tilde{\Delta}(X) \mid$ for every measurable $W \subset T$ and $B \subset A$, and continuous bounded $f: T\times A \to \mathfrak{R}$, $[\int_W\int_B f(t,a)\hat{\sigma}(dt\times da)][\int_W\int_A f(t,a)\hat{\sigma}(dt\times da)] \geq 0\}$. Both are complete metric spaces. Define the action correlated equilibrium correspondence $\xi: \tilde{\Delta}(X) \to\to \hat{\Delta}(T\times A)$ where $\xi(\tilde{\mu}) = \{\hat{\sigma} \in \hat{\Delta}(T\times A) \mid \hat{\sigma}$ satisfies (12a) for every $\alpha_i\}$. The proofs of Lemmas 5.1, 6.2, and 6.3, Theorem 5.3, and Corollary 5.4 are unchanged. Note also that if $\mu \in \Delta(T)$, then $\xi(\mu) \subset \Delta(T\times A)$, so the definition of ξ is unaffected on $\Delta(T)$.

To prove that ξ is lowerhemicontinuous let $\mu \in \Delta(T)$, $\sigma \in \xi(\mu)$, and $\{\mu^n\} \subset \Delta(T)$ converge to μ. For each n let $\tilde{\mu}^n = n(\mu^n - \mu) + \mu$. Then $\tilde{\mu}^n \in \tilde{\Delta}(T)$. Let $\hat{\sigma}^n \in \xi(\tilde{\mu}^n)$, and $\sigma^n = (1/n)\hat{\sigma}^n + (1 - (1/n))\sigma$. By Theorem 5.3, $\sigma^n \in \xi(\mu^n)$, and it is easy to show that $\{\sigma^n\}$ converges to σ. This completes the proof. $\therefore$

Footnotes

[*]Support from NSF grant IRI-8609208 is gratefully acknowledged.

[1]Player 1's type is suppressed since it is degenerate.

[2]Distributional strategies can be used in the same way. The relationship between distributional and behavioral strategies will be discussed in Section 3.

[3]Strictly speaking, a behavioral strategy is a function from types to probability distributions over actions. Any randomization over actions can be handled by the correlation device, so the restriction to pure strategies entails no loss of generality.

[4]When no other qualification is stated, all measures are Borel, and measurability of sets and functions refers to Borel measurability.

[5]Milgrom and Weber (1985) claimed incorrectly that there is a many-to-one correspondence between behavioral and distributional strategies.

References

AUMANN, R. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67-96.

AUMANN, R. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1-18.

BALDER, E. (1986): "Generalized Equilibrium Results for Games with Incomplete Information," *Mathematics of Operations Research*, forthcoming.

BERNHEIM, D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007-1028.

BRANDENBURGER, A., AND E. DEKKEL (1987): "Rationalizability and Correlated Equilibria," *Econometrica*, 55, 1391-1402.

COTTER, K. (1989a): "Similarity of Games with Incomplete Information," unpublished manuscript, Northwestern University.

COTTER, K. (1989b): "Correlated Equilibria with Exogenous Uncertainty," unpublished manuscript, Northwestern University.

MILGROM, P., AND R. WEBER (1985): "Distributional strategies for games with incomplete information," *Mathematics of Operations Research*, 10, 619-632.

MUNKRES, J. (1975): *Topology*. Englewood Cliffs, NJ: Prentice-Hall.

PARTHASARATHY, K. (1967): *Probability Measures on Metric Spaces*. New York: Academic Press.

PEARCE, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029-1050.

RADNER, R., AND R. ROSENTHAL (1982): "Private Information and Pure Strategy Equilibria," *Mathematics of Operations Research*, 7, 401-409.

ROYDEN, H. (1968): *Real Analysis, 2nd. ed.* New York: MacMillan.