PRELIMINARY AND INCOMPLETE. Please do not cite or circulate.

Strategic Complementarities and Mixed Equilibria

Federico Echenique *

April 28, 2000

Abstract

I address two questions: Do the results on pure-strategy equilibria in games of strategic complementarities (GSC) hold for the full set of Nash equilibria? Are there reasons to believe that properly mixed equilibria are worse predictions in GSC than in other classes of games? The answer to both questions is, with some qualifications, yes. When strategy spaces are one-dimensional and mixed strategies are ordered by first-order stochastic dominance, the mixed extension of a GSC is a GSC. In particular, the full set of equilibria is a complete lattice and the extremal equilibria (smallest and largest) are in pure strategies. Second, in GSC, properly mixed equilibria are not likely to be good predictions of play because they are unstable under a broad class of fictitious-play-type learning processes. I also present results on global convergence of learning.

*Department of Economics, University of California at Berkeley. This work has been strongly influenced by conversations with Aaron Edlin. Future versions of this paper might be coauthored with him. I am very grateful to Robert Anderson, David Blackwell, Peter De Marzo, Ted O'Donoghue, Matthew Rabin, Ilya Segal, Chris Shannon and Clara Wang for comments and advise. All the errors are my responsibility.

	O	B
O	1,2	0,0
B	0,0	2,1

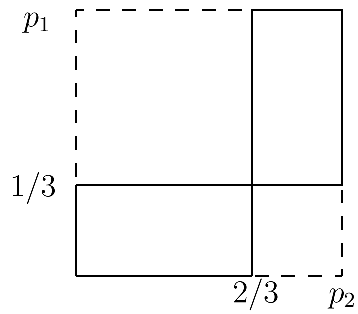

Figure 1: Battle of the Sexes

1 Introduction

Despite some controversy, many game theorists believe mixed-strategy equilibria to be good predictions of play in certain circumstances. However, analysis of games of strategic complementarities has until now focused on pure-strategy equilibria (see e.g. Topkis (1979), Vives (1990), Milgrom and Roberts (1990) Milgrom and Shannon (1994)). [1] Two questions seem natural: Do the conclusions in the literature on games of strategic complementarities hold when also mixed-strategy equilibria are considered? Are there reasons to believe that properly mixed equilibria are worse predictions here than in other classes of games? The answer to both questions is, with some qualifications, yes.

Consider the game "Battle of the Sexes" in Figure 1. Players 1 and 2 choose each simultaneously an element from $\{O, B\}$, the payoffs are specified in the bimatrix to the left. Player 1's best response to 2 playing B is to play B and her best response to 2 playing O is to play O. So, a change by 2 from B to O makes 1 change in the same direction. This is also true for player 2: a change by 1 from B to O makes 2 change in the same direction. Imposing an order on the players' strategies, we can say that O is "larger" than B. Then the best response of each player is increasing in the other player's choice of strategy, this means that Battle of the Sexes is a game of

[1] This is probably because strategic complementarities guarantee that equilibria in pure strategies exist.

strategic complementarities. There are two pure-strategy Nash equilibria of Battle of the Sexes, (O,O) and (B,B); and (O,O) is larger than (B,B).

Now, consider the mixed extension of Battle of the Sexes. That is the game obtained by allowing 1 and 2 to choose probability distributions over $\{O, B\}$ and obtaining payoffs as the expected value of the corresponding pure-strategy payoffs. Let p_i be the probability with which player i selects O. The best responses are now shown in Figure 1 on the right. When 2 plays O with probability smaller than 2/3, 1 sets $p_1 = 0$; when p_2 equals 2/3 player 1 is indifferent between O and B, so any choice of p_1 is a best response; when 2 sets p_2 larger than 2/3, 1 will optimally respond by choosing $p_1 = 1$. Hence, player 1's best response is increasing in 2's choice of p_2. The same is true for 2's optimal choice of p_2. This implies that also the mixed extension of Battle of the Sexes is a game of strategic complementarities.

There are three Nash equilibria of this game, indicated by the three points where the best-response functions intersect, they are $(0, 0), (1/3, 2/3)$ and $(1, 1)$. Note that $(0, 0)$ is smaller than $(1/3, 2/3)$, which is smaller than $(1, 1)$, and that the two extremal equilibria (the smallest and the largest) are in pure strategies.

In this paper I show that when strategy spaces are one-dimensional, and mixed strategies are ordered by first-order stochastic dominance, the mixed extension of a game of strategic complementarities is a game of strategic complementarities. If strategies are multi-dimensional this is not true. It follows that the existing characterization of the set of pure-strategy Nash equilibria of a game of strategic complementarities also applies to the full set of Nash equilibria of games of strategic complementarities that have one-dimensional choice sets. [2] The extremal equilibria (the largest and smallest in the sense of first order stochastic dominance) are in pure strategies. Since comparative statics results have been obtained for the extremal equilibria (see Sobel (1988), Milgrom and Roberts (1990), Milgrom and Shannon (1994)),

[2]The result for games of strategic complementarities is that the set of pure-strategy Nash Equilibria is a complete lattice, see Topkis (1998).

my results extend the comparative statics results to bind all Nash equilibria of a game of strategic complementarities.

I argue in this paper that, in games of strategic complementarities, properly mixed equilibria are not likely to be good predictions of play. Concretely, properly mixed equilibria are unstable under a broad class of fictitious-play-type learning processes. This result applies to arbitrary strategy spaces (and is extended to "smooth fictitious play", see below).

Suppose that our prediction of play for Battle of the Sexes is the properly mixed equilibrium $(p_1, p_2) = (1/3, 2/3)$. Suppose that the players' beliefs about their opponent's play are slightly wrong, suppose 1 believes 2 will select the larger action (O) with probability $2/3 + \epsilon$ and that 2 believes 1 will select the larger action with probability $1/3 + \epsilon$. By choosing $\epsilon > 0$ small enough, these beliefs are arbitrarily close to the equilibrium. Now, as can be seen from the picture, given these beliefs both players will select O with probability 1.

Given that both players observed their opponent choosing O they might infer that they were right in giving O larger weight than what $(1/3, 2/3)$ does. I will say that beliefs are monotone if they behave in this way. Suppose that the game is repeated. Given these new beliefs, with O receiving yet higher weight, play will still be (O,O). It is easy to see that repeated play of Battle of the Sexes will then always reinforce the initial deviation by ϵ of the Nash equilibrium beliefs $(1/3, 2/3)$–so $(1/3, 2/3)$ is unstable.

The question "is mixed strategy equilibria a good prediction of play?" has been addressed in the literature on learning mixed strategy equilibria. Most of the answers have been obtained for restrictive subclasses of games (2X2 games in Fudenberg and Kreps (1993), Benaim and Hirsch (1999) and Kaniovski and Young (1995). 3X3 games in Ellison and Fudenberg (1999)).

Third, I consider the global stability of learning in games of strategic complementarities. Milgrom and Roberts (1990) and Milgrom and Shannon (1994) show that the empirical distribution of adaptive play is in the limit bounded by the two extremal Nash equilibria of the game. Here I extend their result to convergence of "intended play". For learning mixed-strategy equilibria Fudenberg and Kreps (1993) propose the requirement that intended play converge. Consider best-response dynamics in "Matching pennies" it is easy to see that play will cycle, with each player choosing heads half the time and tails half the time. The criterion of convergence of the empirical distribution of play says that best-response dynamics converges to the mixed strategy equilibrium of matching pennies. The cycle is very simple, though, and it is likely that a real player would discover it and use it to improve her payoff. A player that recognizes the cycle could always extract her maximal payoff in every round of play. Note here that intended play does not converge, precisely because of the cycle in choices of heads/tails. [3]

I follow the literature on learning mixed-strategy equilibria in focusing on convergence of intended play. The result is that, for a class of fictitious-play-like learning processes, any subsequential limit (in the sense of intended play) is bounded by the extremal equilibria of the game.

2 Basic Results

2.1 Definitions

A set X with a transitive, reflexive, antisymmetric binary relation $\preceq$ is a **lattice** if whenever $x, y \in X$, both $x \wedge y = \inf\{x, y\}$ and $x \vee y = \sup\{x, y\}$ exist in X. It is **complete** if for every nonempty subset A of X, $\inf A, \sup A$ exist in X. A nonempty subset A of X is a **sublattice** if for all $x, y \in A$, $x \wedge_X y, x \vee_X y \in A$, where $x \wedge_X y$ and

[3] Fudenberg and Kreps (1993) present an example of fictitious play in the "Battle of the Sexes" game where beliefs converge to the mixed-strategy equilibrium of the game but where players' choices follow an exact cycle. In this example players are getting their worst possible payoffs in each stage of the game (i.e. they never "coordinate") and still fail to recognize the simple cycle in play.

$x \vee_X y$ are obtained taking the infimum and supremum as elements of X (as opposed to using the relative order on A). A nonempty subset $A \subset X$ is **subcomplete** if $B \subset A$, $B \neq \emptyset$ implies $\inf_X B, \sup_X B \in A$, again taking inf and sup of B as a subset of X.

Let X be a lattice. $A \subset X$ is **increasing** if, for all $x \in A$, $y \in X$ and $x \preceq y$ imply $y \in A$. $\mathcal{P}(X)$ denotes the set of (Borel) probability measures over X. For $\mu, \nu \in \mathcal{P}(X)$, μ is smaller than ν in the **first order stochastic dominance order** (denoted $\mu \leq_{st} \nu$) if for any increasing set $A \subset X$, $\mu(A) \leq \nu(A)$.

If X is a lattice and T a partially ordered set. $f : X \to \mathbf{R}$ is **supermodular** if $\forall x, y \in X$ $f(x) + f(y) \leq f(x \wedge y) + f(x \vee y)$; $f : X \times T \to \mathbf{R}$ has **increasing differences** in (x, t) if $x \prec x', t \prec t'$, then $f(x', t) - f(x, t) \leq f(x', t') - f(x, t')$; $f : X \times T \to \mathbf{R}$ has **strictly increasing differences** in (x, t) if $x \prec x', t \prec t'$, then $f(x', t) - f(x, t) < f(x', t') - f(x, t')$.

2.2 Basic Results

This section presents some simple results that are needed in the rest of the paper. Lemmas 3 and 4 may be known, but I am not aware of a reference for them, so I include complete proofs here.

Lemma 1 *If $X \subset \mathbf{R}$ is compact, then $\mathcal{P}(X)$ ordered by first order stochastic dominance is a complete lattice.*

Proof: I will use the fact that $X \subset \mathbf{R}$ to identify probability measures with their distribution functions. Let $F, G : X \to [0, 1]$ be two distribution functions on X. We know that F is smaller than G in first order stochastic dominance if and only if $G(x) \leq F(x)$ for all $x \in X$. It is easy to verify that $\leq_{st}$ is a partial order on $\mathcal{P}(X)$.

Define $H : X \to [0, 1]$ by $H(x) = F(x) \wedge G(x)$, H satisfies all requirements for being a distribution function. I will show that $H = F \vee G$ in the first order stochastic dominance order. First, H is larger than both F and G. Second, if H' is larger than

F and G, then for all x, $H'(x) \leq F(x) \wedge G(x) = H(x)$. Thus H' is also larger than H. These two claims imply that $H = F \vee G$. The argument that $F \wedge G$ exists is similar. This proves that the probability distributions are a lattice under $\leq_{st}$.

To prove that the lattice is complete, first I show that the weak topology on $\mathcal{P}(X)$ is finer than the order interval topology. For any $x \in X$, let $U_x = [x, \sup X]$. An order interval $[\mu, \nu]$ in $\mathcal{P}(X)$ is then

$$[\mu, \nu] = \bigcap_{\{U_x : x \in X\}} \left(\{p \in \mathcal{P}(X) : \mu(U_x) \leq p(U_x)\} \cap \{p \in \mathcal{P}(X) : p(U_x) \leq \nu(U_x)\} \right).$$

But for all x, $\{p \in \mathcal{P}(X) : \mu(U_x) \leq p(U_x)\}$ and $\{p \in \mathcal{P}(X) : p(U_x) \leq \nu(U_x)\}$ are weakly closed sets (Aliprantis and Border (1999) Theorem 14.6, note that μ and ν are fixed). Then, order intervals are weakly closed and since the order interval topology is the finest topology for which order intervals are closed, the weak is coarser than the order interval topology.

Now, since X is compact, $\mathcal{P}(X)$ is weakly compact. Then, $\mathcal{P}(X)$ is also compact in the order interval topology because it is finer than the weak topology. By the Birkhoff-Frink characterization of completeness, then, $\mathcal{P}(X)$ is a complete lattice. ■

Lemma 1 does not generalize to arbitrary sublattices $X \subset \mathbf{R}^n$. The following counterexample is taken from Kamae, Krengel, and O'Brien (1977), let $X = \{(0,0), (0,1), (1,0), (1,1)\}$ ordered as a subset of $\mathbf{R}^2$. Then $1/2(\delta_{(0,0)} + \delta_{(1,0)})$ and $1/2(\delta_{(0,0)} + \delta_{(0,1)})$ are maximal elements of the set of lower bounds on

$$\left\{1/2(\delta_{(0,0)} + \delta_{(1,1)}), 1/2(\delta_{(0,1)} + \delta_{(1,0)})\right\}.$$

This shows that if X is any complete lattice that contains two unordered elements, then $\mathcal{P}(X)$ is not a lattice when ordered by first order stochastic dominance.

This would create a problem when we study convergence of learning processes because we cannot argue that a monotone sequence converges to the supremum of its range (suprema are not everywhere well defined on partially ordered spaces that are not complete lattices). Fortunately, it is not hard to prove that monotone sequences in $\mathcal{P}(X)$ are convergent, as long as $X \subset \mathbf{R}^n$ is compact.

Lemma 2 *Let $X \subset \mathbf{R}^n$ be compact. If $\{\mu_n\}$ is a monotone increasing sequence in $\mathcal{P}(X)$, then it converges weakly to a probability measure μ and $\mu_n \leq_{st} \mu$ for all n.*

Proof: By Kamae, Krengel, and O'Brien's (1977) Proposition 4 there is a probability space $(\Omega, \mathcal{F}, p)$ and a sequence of random variables $\{Z_n\}$, taking values in X that is increasing a.s. and such that, for all n, Z_n is distributed as μ_n. By compactness of X, there is a random variable Z with $Z_n \uparrow Z$ a.s. Then, $\{\mu_n\}$ is weakly convergent to the distribution μ induced by Z on X and for any increasing integrable function $f : X \to \mathbf{R}$, $f(Z_n(\omega)) \leq f(Z(\omega))$ a.s for all n, so $\int f d\mu_n = \int f(Z_n)dp \leq \int f(Z)dp = \int f d\mu$. Hence, $\mu_n \leq_{st} \mu$ for all n. ■

It is easy to see that the integral of any (integrable) function is supermodular, when viewed as a function on the space of probability distributions over a subset of real numbers. Basically, if Y is a lattice that is also a vector space and where $y + z = y \vee z + y \wedge z$ for all $y, z \in Y$, then any linear function on Y is supermodular. Since integrals are linear in probability distributions, the result would follow. The set of probability distributions is not a vector space, though, so the proof is slightly more involved.

Lemma 3 *Let $f : X \subset \mathbf{R} \to \mathbf{R}$ be integrable. Then, $p \mapsto \int_X f(x)dp(x)$ is supermodular on $\mathcal{P}(X)$.*

Proof: Let $p_1, p_2 \in \mathcal{P}(X)$ define $E_1 = \{x \in X : p_2([x, \sup X]) \leq p_1([x, \sup X])\}$ and $E_2 = \{x \in X : p_1([x, \sup X]) < p_2([x, \sup X])\}$. Note that (E_1, E_2) is a measurable partition of X. On subsets of E_1, $p_1 \vee p_2$ coincides with p_1 (see the proof of Lemma 1) while $p_1 \wedge p_2$ coincides with p_2. On subsets of E_2, $p_1 \vee p_2$ coincides with p_2 while $p_1 \wedge p_2$ coincides with p_1. Then, $\int_X f(x)dp_1(x) + \int_X f(x)dp_2(x) =$

$$\int_{E_1} f(x)dp_1(x) + \int_{E_2} f(x)dp_1(x) + \int_{E_1} f(x)dp_2(x) + \int_{E_2} f(x)dp_2(x) =$$

$$\int_{E_1} f(x)dp_1 \vee p_2(x) + \int_{E_2} f(x)dp_1 \wedge p_2(x) + \int_{E_1} f(x)dp_1 \wedge p_2(x) + \int_{E_2} f(x)dp_1 \vee p_2(x)$$
$= \int_X f(x)dp_1 \vee p_2(x) + \int_X f(x)dp_1 \wedge p_2(x)$. ■

Lemma 4 *Let $X_1, X_2, \ldots X_n$ be a collection of subsets of* $\mathbf{R}$ *and $f : X = \times_{i=1}^n X_i \to \mathbf{R}$ be integrable. If f has increasing differences in (x_i, x_{-i}) for all i, then*

$$(p_1, \ldots p_n) \mapsto \int_X f(x)d(p_1 \times \ldots p_n)(x) \ : \times_{i=1}^n \mathcal{P}(X_i) \to \mathbf{R}$$

has increasing differences in (p_i, p_{-i}) for all i.

Proof: First note that repeated application of Fubini's theorem implies that $(p_1, \ldots p_n)$ is smaller (componentwise) than $(p_1', \ldots p_n')$ if and only if $\int_X g(x)d(p_1 \times \ldots p_n)(x) \leq \int_X g(x)d(p_1' \times \ldots p_n')(x)$ for every increasing integrable $g : X \to \mathbf{R}$. In other words that the first order stochastic dominance order on the set of product measures coincides with the product order.

Fix i and $p_i, p_i' \in \mathcal{P}(X_i)$ where $p_i \leq_{st} p_i'$. For any (componentwise) $x_{-i} \leq x_{-i}'$, $\hat{x}_i \mapsto f(\hat{x}_i, x_{-i}') - f(\hat{x}_i, x_{-i})$ is increasing. Then,

$$\int_{X_i} f(\hat{x}_i, x_{-i}')dp_i(\hat{x}_i) - \int_{X_i} f(\hat{x}_i, x_{-i})dp_i(\hat{x}_i) \leq \int_{X_i} f(\hat{x}_i, x_{-i}')dp_i'(\hat{x}_i) - \int_{X_i} f(\hat{x}_i, x_{-i})dp_i'(\hat{x}_i),$$

so that $x_{-i} \mapsto \int_{X_i} f(\hat{x}_i, x_{-i})dp_i'(\hat{x}_i) - \int_{X_i} f(\hat{x}_i, x_{-i})dp_i(\hat{x}_i)$ is increasing. Now, if $p_{-i} \leq_{st} p_{-i}'$, then,

$$\int_{X_{-i}} \left\{ \int_{X_i} f(\hat{x}_i, \hat{x}_{-i})dp_i'(\hat{x}_i) \right\} dp_{-i}(\hat{x}_{-i}) - \int_{X_{-i}} \left\{ \int_{X_i} f(\hat{x}_i, \hat{x}_{-i})dp_i(\hat{x}_i) \right\} dp_{-i}(\hat{x}_{-i})$$
$$\leq$$
$$\int_{X_{-i}} \left\{ \int_{X_i} f(\hat{x}_i, \hat{x}_{-i})dp_i'(\hat{x}_i) \right\} dp_{-i}'(\hat{x}_{-i}) - \int_{X_{-i}} \left\{ \int_{X_i} f(\hat{x}_i, \hat{x}_{-i})dp_i(\hat{x}_i) \right\} dp_{-i}'(\hat{x}_{-i}).$$

Using Fubini, $\int_X fd(p_i' \times p_{-i}) - \int_X fd(p_i \times p_{-i}) \leq \int_X fd(p_i' \times p_{-i}') - \int_X fd(p_i \times p_{-i}')$. This is just saying that $p_{-i} \mapsto \int_X fd(p_i' \times p_{-i}) - \int_X fd(p_i \times p_{-i})$ is increasing. ■

3 Mixed Strategies in Supermodular Games

A normal-form game is described by a set of players N together with strategy spaces S_i and payoffs $u_i : S = \times_{i \in N} S_i \to \mathbf{R}$ for all $i \in N$. Let $\Gamma = (N, \{(S_i, u_i) : i \in N\})$ be a normal-form game. Here I will assume that N is finite. The **mixed extension** of Γ is the normal-form game obtained when players $i \in N$ are allowed to choose randomizations $\sigma_i \in \Sigma_i = \mathcal{P}(S_i)$ over the strategies in S_i. The randomizations are assumed to be independent, so a strategy profile σ is a collection $(\sigma_i)_{i \in N}$ that induces a distribution over S by independently mixing the marginals σ_i. Payoffs are $U(\sigma) = \int_S u(s) d\sigma(s)$. Thus the mixed extension of Γ is the normal-form game $(N, \{(\Sigma_i, U_i) : i \in N\})$.

Definition 1 *A normal-form game* $\Gamma = (N, \{(S_i, u_i) : i \in N\})$ *is a* **supermodular game** *if, for all* $i \in N$*,*

1. S_i *is a complete lattice;*

2. u_i *is bounded,* $s_i \mapsto u_i(s_i, s_{-i})$ *is supermodular and* $(s_i, s_{-i}) \mapsto u_i(s_i, s_{-i})$ *has increasing differences;*

3. $s_i \mapsto u_i(s_i, s_{-i})$ *is upper semicontinuous and* $s_{-i} \mapsto u_i(s_i, s_{-i})$ *is continuous.*

If, in addition, for all $i \in N$ $S_i \subset \mathbf{R}$ *then it is called a* **simple supermodular game**. *If* $(s_i, s_{-i}) \mapsto u_i(s_i, s_{-i})$ *has strictly increasing differences then it is called a* **strict supermodular game**.

Theorem 1 *If* $(N, \{(S_i, u_i) : i \in N\})$ *is a simple supermodular game then its mixed extension* $(N, \{(\Sigma_i, U_i) : i \in N\})$ *is a supermodular game.*

Proof: Lemmas 1, 3 and 4 almost complete the proof. We just need to note that for all i, (weak) upper semicontinuity and continuity of $U_i(\sigma_i, \sigma_{-i})$ in σ_i and σ_{-i}, respectively, follow from standard results (see Aliprantis and Border (1999) Theorem 14.5). ■

Corollary 1 *If* $\Gamma = (N, \{(S_i, u_i) : i \in N\})$ *is a simple supermodular game then the set of mixed equilibria of* Γ *is a complete lattice and the extremal equilibria (largest and smallest) are in pure strategies.*

Proof: That the set of equilibria is a complete lattice follows from Zhou's (1994) fixed point theorem. Let $\sigma = (\sigma_1, \ldots \sigma_n)$ be the largest equilibrium and assume that it is not in pure strategies. Fix $i \in N$. Let $\beta_i(\sigma_{-i}) = \text{argmax}_{\sigma \in \Sigma_i} U(\sigma, \sigma_{-i})$ be i's best response correspondence.

The support of σ_i is $L_i = \text{argmax}_{s \in S_i} \int_{S_{-i}} u_i(s, s_{-i}) d\sigma_{-i}(s_{-i})$, the set of i's pure best responses to σ_{-i}. By Topkis's (1978) results, L_i is a subcomplete sublattice. Then, letting $\delta_{\sup L_i}$ be a point mass on $\sup L_i$, $\delta_{\sup L_i} = \sup \beta_i(\sigma_{-i})$ since by subcompleteness of L_i it is a best response to σ_{-i} and it is an upper bound on all mixed-strategy best responses.

Now, construct a sequence $\{x_k\}$ in $\Sigma = \times_{i \in N} \Sigma_i$ by $x_0 = (\sigma_1, \ldots \sigma_n)$, $x_1 = \times_{i \in N} \delta_{\sup L_i}$ and $x_k = \times_{i \in N} \sup \beta_i(x_{k-1})$ for all $k \geq 2$. Since x_0 is smaller than x_1 and each β_i is increasing (by Topkis's (1978) Theorem), $\{x_k\}$ is a componentwise monotone increasing sequence. By Lemma 2, for all i, $\{x_{ki}\}$ converges weakly to, say, x_i^*. Note that x_1 is smaller than $x^* = (x_1^*, \ldots x_n^*)$. But at least one player was assumed to have a properly mixed strategy under σ so $\sigma <_{st} x_1$. Then $\sigma <_{st} x^*$.

Finally, the claim is that x^* is an equilibrium. To see this, note that the aggregate best response correspondence $\beta(x) = \times_{i \in N} \beta_i(x_{-i})$ is upper hemicontinuous and therefore has a closed graph. The sequence $\{(x_k, x_{k-1})\}$ is in the graph of β and $x_k \to x^*$. Then, $x^* \in \beta(x^*)$ so x^* is an equilibrium. But, x^* is larger than σ, contradicting the assumption that σ is the supremum of the set of equilibria. The proof that the smallest equilibrium is pure follows analogously. ■

The counterexample in Section 2 shows that, when strategy spaces are utidimensional, the set of mixed strategies is not a lattice. This implies that we lack the mathematical structure needed for the current theory of complementarities. We need the lattice property of strategies to make sense of increasing best responses when they are not singleton valued. Mutiple best responses are always present when dealing with mixed equilibria and there does not seem to be a simple solution to the requirement that strategy spaces be lattices.

4 Learning Mixed Strategies

The model of learning presented here is similar to the one in Fudenberg and Kreps (1993). Learning takes place through repeated play of a stage game; the stage game is one of incomplete information. Player i's type is $\omega_i \in \Omega_i$, the type space Ω_i is assumed to be a compact topological space. The set of all type profiles is $\Omega = \times_{i \in N} \Omega_i$. In each stage of the repeated game, a type profile $\omega \in \Omega$ is drawn at random, player i is informed of ω_i and chooses a stage game strategy $s_i \in S_i$. When a strategy profile $s \in S$ is chosen, the payoff to i in the stage game is $u_i(\omega_i)(s)$. I will assume that the game that results from fixing ω is a strict supermodular game, i.e. for each $\omega \in \Omega$, $(N. \{(S_i, u_i(\omega_i)) : i \in N\})$ is a supermodular game. Also, let $\omega_i \mapsto u_i(\omega_i)(s)$ be a continuous function for all $i \in N$ and $s \in S$.

Let p be a probability measure over $\Omega^\infty = \Omega^{\mathbf{N}}$, the space of all sequences of draws (Ω^∞ is endowed with the canonical σ-algebra obtained from the Borel subsets of Ω). I will assume that sequences of type profiles $\omega^\infty \in \Omega^\infty$ are drawn according to p.

The present setup embeds two important special cases: literally mixed strategies and "purified" mixed strategies.

1. (Mixed Strategies) Let $u_i(\omega_i)$ be independent of ω_i. The type spaces represent only randomization devices. In this case, strategies are simply the mixed or correlated strategies that extend the stage game.

2. (Purification) Let $\Omega_i \subset \mathbf{R}^{S_i}$ and $u_i^\delta(\omega_i) = g_i + \delta\omega_i$ for $\delta > 0$ and an integrable $g_i \in \mathbf{R}^S$. This is the setup of Harsanyi's Purification Theorem. Let Γ^δ be

the resulting game of incomplete information. Harsanyi's Theorem says that–generically in finite games–for any mixed equilibrium σ of Γ^0 there is a collection (σ_δ) in Σ, where σ_δ is a (pure) equilibrium of Γ^δ, such that $\sigma = \lim_{\delta\to 0}\sigma_\delta$.

Fudenberg and Kreps (1993) argue that learning mixed strategies should be studied in the purification setup. Fudenberg and Kreps (1993) and Ellison and Fudenberg (1999) study learning of mixed equilibria in the purification setup.

At each stage a pure strategy profile $s \in S$ results from the players' choices. Histories of play $(s_1, \ldots s_t)$ are denoted h_t. The set of all histories of length t is $H_t = S^t$ and $H = \cup_{t=0}^{\infty} H_t$ is the set of all histories of finite length, including $H_0 = \emptyset$, the "null history".

Each player i chooses a strategy $\xi_i : \Omega_i^\infty \times H \to S_i$ and is endowed with beliefs $\mu_i : H \to \mathcal{P}(S_{-i})$. The interpretation is that, at each time t and history h_t, $\mu_i(h_t) \in \mathcal{P}(S_{-i})$ represents i's assessment of her opponents' play in the $t+1$ stage of the game. The assumptions below will guarantee that $\xi_i(\omega^\infty, h_t)$ depends only on ω_{it}, the period-t realization of the payoff shock to player i.

If $\xi = (\xi_i)_{i\in N}$ is a collection of strategies for all players and $\mu = (\mu_i)_{i\in N}$ is a collection of beliefs, then the pair (ξ, μ) is a **system of behavior and beliefs**. Note that I allow player i to believe that her opponents' play is correlated (for a discussion of the importance of this, see Fudenberg and Kreps (1993)).

Definition 2 *A system of behavior and beliefs* (ξ, μ) *is* **myopic** *if for all* $i \in N$, $h_t \in H$ *and* $\omega_i^\infty \in \Omega_i^\infty$, $\xi_i(\omega_i^\infty, h_t) \in argmax_{s_i\in S_i} \int u_i(\omega_{i,t})(s_i, \tilde{s}_{-i})\mu_i(h_t)(d\tilde{s}_{-i})$

The assumption of myopic behavior is very common in the literature on learning. It is restrictive because it implies that players do not attempt to manipulate the future behavior of their opponents by current actions. It is usually justified by assuming that, in each period of time, players are selected at random from a large population to play the stage game, so the likelihood that two particular players will meet more than once to play the stage game is negligible (see chapter 1 of Fudenberg and Levine (1998) for a discussion).

Definition 3 *Beliefs* μ *are* **weakly monotone** *if, for all* $i \in N$*, and* $h_t \in H$*,* $\sup$ *supp* $\mu_i(h_t) \leq s_{-i\tau}$ *for* $\tau = t+1, \ldots T$ *implies that* $\mu_i(h_t) \leq_{st} \mu_i(h_T)$*.*

Theorem 2 proves that, at any properly mixed Nash equilibrium σ, there are arbitrarily small perturbations which set off learning dynamics that are always outside of a neighborhood of σ. This is an instability result: small perturbations from σ are never "corrected" by subsequent dynamics. The perturbation takes the form of slightly wrong beliefs. It can also be done by perturbing behavior. That behavior ξ is always outside of a neighborhood W of σ means that the distribution of the random variable $\omega \mapsto \xi(\omega, \mu_t)$ will always lie in W^c.

Theorem 2 *(Mixed Setup) Let* Γ *be a strict supermodular game. Let* σ *be a properly mixed Nash Equilibrium and* (μ, ξ) *be a myopic system of behavior and beliefs, with weakly monotone beliefs. For any (weak) neighborhood* V *of* σ*, there is* $\mu' \in V$ *and a neighborhood* W *of* σ *such that if* $\mu_0 = \mu'$ *then* (μ, ξ) *is a.s. a sequence in* W^c*.*

Proof: Suppose, to simplify the notation, that all players play a nondegenerate mixed strategy under σ. Note that $(1-\epsilon)\sigma + \epsilon\delta_{\sup\beta(\sigma)} \to \sigma$ weakly as $\epsilon \to 0$ (since for any bounded, continuous, real-valued g, $(1-\epsilon)\int g d\sigma + \epsilon \int g d\delta_{\sup\beta(\sigma)} \to \int g d\sigma$). Then, given a neighborhood V of σ, there is $\epsilon \in (0,1)$ such that $\mu' = (1-\epsilon)\sigma + \epsilon\delta_{\sup\beta(\sigma)} \in V$.

I claim that $\sigma <_{st} \mu' <_{st} \delta_{\sup\beta(\sigma)}$. Let E be an increasing set and let $\beta(p) = \times_{i\in N}\text{argmax}_{s_i \in S_i} \int u_i(s_i, \tilde{s}_{-i}) dp_i(\tilde{s}_{-i})$ be the aggregate, pure-strategy, best-response correspondence. Note that β is non-empty-, subcomplete-, and sublattice-valued, so $\sup\beta(\sigma) \in \beta(\sigma)$. Also note that the support of σ is in $\beta(\sigma)$. If $E \cap \beta(\sigma) = \emptyset$, then $\sigma(E) = \delta_{\sup\beta(\sigma)}(E) = 0$, so $\mu'(E) = 0$. If $x \in E \cap \beta(\sigma)$ then $x \leq \sup\beta(\sigma)$ so $\sup\beta(\sigma) \in E$, because E is increasing. Then $\delta_{\sup\beta(\sigma)}(E) = 1$, which implies that

$$\begin{aligned} \sigma(E) &= (1-\epsilon)\sigma(E) + \epsilon\sigma(E) \leq (1-\epsilon)\sigma(E) + \epsilon \\ &= (1-\epsilon)\sigma(E) + \epsilon\delta_{\sup\beta(\sigma)}(E) = \mu'(E) \leq 1 = \delta_{\sup\beta(\sigma)}(E). \end{aligned}$$

Proving that $\sigma(E) \leq \mu'(E) \leq \delta_{\sup\beta(\sigma)}(E)$ for every increasing set E, thus $\sigma \leq_{st} \mu' \leq_{st} \delta_{\sup\beta(\sigma)}$. Now let $E = [\sup\beta(\sigma), \sup S]$. Since σ is a properly mixed equilibrium, $\sigma(E) < 1$. Then,

$$\sigma(E) = (1-\epsilon)\sigma(E) + \epsilon\sigma(E) < (1-\epsilon)\sigma(E) + \epsilon = \mu'(E) < 1 = \delta_{\sup\beta(\sigma)}(E).$$

So, $\sigma <_{st} \mu' <_{st} \delta_{\sup\beta(\sigma)}$.

Let (μ, ξ) be a system of myopic behavior and monotone beliefs, let $\mu_0 = \mu'$ and fix $\omega^\infty \in \Omega^\infty$. I will show by induction that, if the sequence $\{x_t\}$ with $x_t = \xi(\omega_t, \mu(h_{t-1}))$ is the realized play, then $\sup\beta(\sigma) \leq x_t$ and $\mu_0 \leq_{st} \mu_t$ for every t.

Let $T = \{\sigma, \mu_0\}$, and $f : S_i \times T \to \mathbf{R}$ be $f(s_i, t) = \int_{S_{-i}} u_i(s_i, \tilde{s}_{-i}) dt(\tilde{s}_{-i})$. I claim that f satisfies the strict single crossing property in (s_i, t). Let $s_i < s_i'$ and $f(s_i', \sigma) - f(s_i, \sigma) \geq 0$. Now, $f(s_i', \mu_0) - f(s_i, \mu_0) =$

$$(1-\epsilon)\int_{S_{-i}} [\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,]\, d\sigma(\tilde{s}_{-i}) + \epsilon \int_{S_{-i}} [\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,]\, d\delta_{\sup\beta(\sigma)}(\tilde{s}_{-i})$$
$$= (1-\epsilon)\int_{S_{-i}} [\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,]\, d\sigma(\tilde{s}_{-i}) + \epsilon\, [\, u_i(s_i', \sup\beta(\sigma)_{-i}) - u_i(s_i, \sup\beta(\sigma)_{-i}) \,].$$

Suppose, by way of contradiction, that $f(s_i', \mu_0) - f(s_i, \mu_0) \leq 0$. Then $\int_{S_{-i}} [\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,]\, d\sigma(\tilde{s}_{-i}) \geq 0$ implies that $[\, u_i(s_i', \sup\beta(\sigma)_{-i}) - u_i(s_i, \sup\beta(\sigma)_{-i}) \,] \leq 0$. But Γ is a strict supermodular game, so $\tilde{s}_{-i} \mapsto [\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,]$ is strictly increasing. Hence $[\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,] \leq 0$ for all $\tilde{s}_{-i} \in \beta(\sigma)_{-i}$. Then $\int_{S_{-i}} [\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,]\, d\sigma(\tilde{s}_{-i}) \geq 0$ implies that $[\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,] = 0$ σ-a.s. on $\beta(\sigma)$. This is a contradiction because $\tilde{s}_{-i} \mapsto [\, u_i(s_i', \tilde{s}_{-i}) - u_i(s_i, \tilde{s}_{-i}) \,]$ is strictly increasing, so there is a unique $s_{-i} \in S_{-i}$ such that $[\, u_i(s_i', s_{-i}) - u_i(s_i, s_{-i}) \,] = 0$. But since σ is a properly mixed strategy, $\sigma_i(\beta(\sigma)_{-i} \setminus \{s_{-i}\}) > 0$.

Now note that for every $t \in T$, $\beta(t) = (\mathrm{argmax}_{s_i \in S_i} f(s_i, t))_{i \in N}$. Milgrom and Shannon's (1994) Monotone Selection Theorem implies then that for any $x \in \beta(\sigma)$, $x \leq x_1 = \xi(\omega_1, \mu_0) \in \beta(\mu_0)$. Then $\sup\beta(\sigma) \leq x_1$ and, by monotonicity of beliefs, $\mu_0 \leq \mu_1 = \mu(x_1)$. Suppose that $\mu_0 \leq_{st} \mu_t = \mu(h_{t-1})$. The Monotone Selection Theorem and $\sigma <_{st} \mu_0 \leq_{st} \mu_t$ imply that $\sup\beta(\sigma) \leq \xi(\omega_t, \mu_t) = x_t$. Monotone beliefs imply then that $\mu_0 \leq_{st} \mu_{t+1} = \mu(h_t)$. By induction then, for every t, both μ_t and the distribution of $\omega \mapsto \xi(\omega, h_t)$ are larger than μ_0. Let $W = [\mu_0, \delta_{\sup S}]^c$, a weak neighborhood (see the proof of Lemma 1). ■

Theorem 3 *(Purification Setup) Let* Γ^0 *be a supermodular game and* σ *be a properly mixed-strategy equilibrium of* Γ^0*. Let best-responses be strictly increasing at* σ*. If* (σ_δ) *is a purifying sequence then there is* $\underline{\delta}$ *such that if* $\delta \leq \underline{\delta}$ *then* σ_δ *is unstable in the sense of Theorem 2.*

Proof: Let $\beta^0(p) = \times_{i \in N} \text{argmax}_{s_i \in S_i} \left\{ \int g_i(s_i, \tilde{s}_{-i}) dp_i(\tilde{s}_{-i}) + \delta \omega_i(s_i) \right\}$ be the aggregate, pure-strategy, best-response correspondence in the δ-perturbed game. I will first show that there is $\underline{\delta}$ such that if $\delta \leq \underline{\delta}$ then $x \leq y$ for any $x \in \beta^\delta(\sigma_\delta)$, $y \in \beta^\delta(\mu')$ and $\sigma_\delta <_{st} \mu'$. Suppose, by way of contradiction, that there is a sequence $\{\delta_k\}$ with $\delta_k \to 0$ and that, for each k, there are x_k, y_k, μ_k with $\sigma_{\delta_k} <_{st} \mu_k$, $x_k \in \beta^{\delta_k}(\sigma_{\delta_k})$, $y_k \in \beta^{\delta_k}(\mu_k)$ but where $x_k > y_k$. By dropping to a subsequence, by compactness of strategy spaces, let $x = \lim_{k \to \infty} x_k$, $y = \lim_{k \to \infty} y_k$ and $\mu = \lim_{k \to \infty} \mu_k$. Then $y \leq x$ and, since $\sigma_{\delta_k} \to \sigma$, σ.

Definition 4 *Beliefs are* **weakly asymptotically empirical** *if, for all* $i \in N$ *and* $\omega^\infty \in \Omega^\infty$*, whenever a sequence of play* $\{s_t\}$ *is convergent, say* $s = \lim s_t$*, and the resulting sequence of beliefs* $\{\mu_i(\omega^\infty, h_t)\}$ *is convergent then* $\mu_i(\omega^\infty, h_t) = \mu_i(\omega^\infty, (s_1, \ldots s_t)) \to \delta_{s_{-i}}$*. Beliefs* μ *are* **monotone** *if, for all* $i \in N$ *and* $\omega^\infty \in \Omega^\infty$*,* $s_{t-1} \leq s_t$ *implies that* $\mu_i(\omega^\infty, h_{t-1}) \leq_{st} \mu_i(\omega^\infty, h_{t-1}s_t)$ *and if* $h_t \leq h'_t$ *implies* $\mu_i(\omega^\infty, h_t) \leq_{st} \mu_i(\omega^\infty, h'_t)$*.*

Theorems 4 and 5 are the main results on "global convergence" of intended play. Theorem 4 says that myopic rules that respond to monotone beliefs about opponents' play are in the limit bounded by the largest and smallest pure strategy equilibria of the game. The result is an extension to randomized play of Milgrom and Roberts's (1990) results. Theorem 5 says that along any "purifying sequence" limit behavior is bounded by a sequence that converges to pure strategy equilibria of the original game.

In this setting, Fudenberg and Kreps (1993) present results on global convergence of intended play for a class of 2X2 games. Theorems 4 and 5 contain a weaker conclusion than global convergence, they only bound the limiting behavior of learning processes; but when equilibrium is unique–like in Fudenberg and Kreps–global convergence is obtained.

Lemma 5 is instrumental in proving the results on "global convergence" in the paper. In the lemma, $\mathcal{E}(\omega)$ denotes the set of pure strategy Nash equilibria of the one shot game obtained by fixing $\omega \in \Omega$ and letting ω be common knowledge. This is not something meaningful yet, Theorems 4 and 5 will give the implications of Lemma 5.

Lemma 5 *Let Γ be a supermodular game. Let (μ, ξ) be a myopic system of behavior and beliefs with monotone weakly asymptotically empirical beliefs. There are $\omega', \omega'' \in \Omega$ and $\underline{e} \in \mathcal{E}(\omega'), \overline{e} \in \mathcal{E}(\omega'')$ such that p-a.s.: $\underline{e} \leq \tilde{\xi} \leq \overline{e}$ and $\delta_{\underline{e}} \leq_{st} \tilde{\mu} \leq_{st} \delta_{\overline{e}}$ for all subsequential limits $\tilde{\mu}$ of $\{\mu(h_t)\}$ and $\tilde{\xi}$ of $\{\xi(h_t)\}$.*

Proof: For any $\omega \in \Omega$ and $p \in \times_{i\in N}\mathcal{P}(S_{-i})$, by Lemma 4, $(s_i, p_i) \mapsto \int u_i(\omega_i)(s_i, \tilde{s}_{-i})dp_i(\tilde{s}_{-i})$ has increasing differences. By Topkis's Theorem, the set of pure-strategy best responses $\beta^s(\omega, p) = \times_{i\in N}\text{argmax}_{s_i\in S_i} \int u_i(\omega_i)(s_i, \tilde{s}_{-i})dp_i(\tilde{s}_{-i})$ is increasing in the strong set order.

First, consider beliefs and behavior as follows. Let initial assessments be $\mu'_0 = (\delta_{\inf S_{-i}})_{i\in N}$ so that for any $h_1 \in S$ the corresponding beliefs $\mu'_1 = \underline{\mu}(h_1)$ satisfy $\mu_0 = (\delta_{\inf S_{-i}})_{i\in N} \leq_{st} \mu'_1$. Then, for any $\omega \in \Omega$, $\beta^s(\omega, \mu'_0)$ is smaller than $\beta^s(\omega, \mu'_1)$ in the strong set order. In particular, $\inf \{\beta^s(\omega, \mu'_0) : \omega \in \Omega\} \leq \inf \{\beta^s(\omega, \mu'_1) : \omega \in \Omega\}$. Also, these infima are achieved for some $\omega \in \Omega$ since $\omega \mapsto u_i(\omega)$ is continuous and S_i is a compact subset of $\mathbf{R}$. Now, construct a sequence of play $\{x_t\}$ and probability assessments over opponents' play $\{\mu'_t\}$ recursively by $x_t = \inf \left\{\beta^s(\omega, \mu'_{t-1}) : \omega \in \Omega\right\}$ and $\mu'_t = \mu(h_t)$. Since $x_0 \leq x_1$, $\mu'_0 \leq_{st} \mu'_1$ and the maps $x_t \mapsto \mu(h_{t-1}x_t)$ and $\mu'_t \mapsto \inf \{\beta^s(\omega, \mu'_t) : \omega \in \Omega\}$ are monotone increasing it is clear by induction that the sequences $\{x_t\}$ and $\{\mu'_t\}$ are monotone increasing.

The sequence $\{x_t\}$ is convergent because it is monotone increasing sequence on a bounded set $S \subset \mathbf{R}^n$. Say that $\underline{e} = \lim x_t$. Each component of the sequence $\{\mu'_t\}$ is monotone increasing. By Lemma 2, $\{\mu'_t\}$ is convergent. Since beliefs are weakly asymptotically empirical, $\delta_{\underline{e}} = \lim \mu'_t$. For each t, there is $\omega'_t \in \Omega$ such that $x_t \in \beta^s(\omega'_t, \mu'_t)$. Since Ω is compact we can say, after dropping to a subsequence, that $\{\omega_t\}$ is convergent and set $\omega' = \lim \omega'_t$. Now, $\{(x_t, \mu'_t, \omega'_t)\}$ is a convergent sequence in the graph of β^s. By upper-hemicontinuity of β^s, $\underline{e} \in \beta^s(\omega', \delta_{\underline{e}})$. This means that $\underline{e}$ is a Nash equilibrium for the stage game obtained when $\omega' \in \Omega$ is drawn, i.e. $\underline{e} \in \mathcal{E}(\omega')$.

Similarly, it is possible to construct a sequence of play $\{y_t\}$ and probability assessments over opponents' play $\{\mu''_t\}$ by setting $\mu''_0 = (\delta_{\sup S_{-i}})_{i \in N}$, $\mu''_t = \mu(h_t)$ and $y_t = \sup \left\{\beta^s(\omega, \mu''_{t-1}) : \omega \in \Omega\right\}$. Repeating the argument above we obtain a convergent sequence $\{(y_t, \mu''_t, \omega''_t)\}$, say $(\overline{e}, \mu'', \omega'') = \lim_t (y_t, \mu''_t, \omega''_t)$ and $\overline{e} \in \beta^s(\omega'', \delta_{\overline{e}})$ so that $\overline{e} \in \mathcal{E}(\omega'')$.

Fix $\omega^\infty \in \Omega^\infty$. I will show by induction that the sequence of play and probability assessments over opponents' play $\{\mu(\omega^\infty, h_t), \xi(\omega^\infty, \mu(h_t))\}$ satisfies $\mu'_t \leq_{st} \mu(h_t) \leq_{st} \mu''_t$ and $x_t \leq \xi(\omega^\infty, \mu(\omega, h_{t-1})) \leq y_t$. First, $\mu'_0 \leq_{st} \mu(h_0) \leq_{st} \mu''_0$ by the definitions of μ'_0 and μ''_0. By monotonicity of $p \mapsto \beta^s(\tilde{\omega}, p)$, $x_1 = \inf \{\beta^s(\tilde{\omega}, \mu'_0) : \tilde{\omega} \in \Omega\} \leq z \leq \sup \{\beta^s(\tilde{\omega}, \mu'_0) : \tilde{\omega} \in \Omega\} = y_1$ for all $z \in \beta^s(\tilde{\omega}, \mu'_0)$. In particular, since myopic behavior implies that $\xi(\omega^\infty, \mu(h_0)) \in \beta^s(\omega_1, \mu(\omega^\infty, h_0))$, $x_1 \leq \xi(\omega^\infty, \mu(h_0)) \leq y_1$. Second, suppose that $\mu'_{t-1} \leq_{st} \mu(h_{t-1}) \leq_{st} \mu''_{t-1}$ and $x_l \leq \xi(\omega^\infty, \mu(\omega, h_{l-1})) \leq y_l$ for $1 \leq l \leq t-1$. Using monotonicity of $p \mapsto \beta^s(\omega^\infty, p)$ again, $\mu'_{t-1} \leq_{st} \mu(h_{t-1}) \leq_{st} \mu''_{t-1}$ implies that $x_t \leq \xi(\omega^\infty, \mu(h_{t-1})) \leq y_t$. By monotonicity of beliefs,

$$\mu(x_1, \dots x_t) \leq_{st} \mu\left[\xi(\omega^\infty, \mu(h_0), \dots \xi(\omega^\infty, \mu(h_{t-1}))\right] \leq_{st} \mu(y_1, \dots y_t).$$

Then, $\mu'_t = \mu(x_1, \dots x_t)$ and $\mu''_t = \mu(y_1, \dots y_t)$ imply that $\mu'_t \leq_{st} \mu(h_t) \leq_{st} \mu''_t$. This proves that $\mu'_t \leq_{st} \mu(h_{t-1}) \leq_{st} \mu''_t$ and $x_t \leq \xi(\omega^\infty, \mu(h_{t-1})) \leq y_t$ for all t.

Now, $x_t \to \underline{e}$, $y_t \to \overline{e}$, since $\omega^\infty \in \Omega^\infty$ was arbitrary the first conclusion follows. Also $\mu'_t \to \delta_{\underline{e}}$ and $\mu''_t \to \delta_{\overline{e}}$ weakly and for every increasing subset A of S_{-i}, $\mu'_t(A) \leq$

$\mu(h_t)(A) \leq \mu''_t(A)$. Then, $\mu'_t(A) \to \delta_{\underline{e}}(A)$ and $\mu''_t(A) \to \delta_{\overline{e}}(A)$ implies that if $\mu(h_{t_k})$ is a subsequence converging to $\tilde{\mu} \in \mathcal{P}(S_{-i})$ then $\delta_{\underline{e}}(A) \leq \tilde{\mu}(A) \leq \delta_{\overline{e}}(A)$. Hence, $\delta_{\underline{e}} \leq_{st} \tilde{\mu} \leq_{st} \delta_{\overline{e}}$ a.s. ■

Theorem 4 *(Mixed Setup) Let Γ be a supermodular game. Let (μ, ξ) be a myopic system of behavior and beliefs with monotone, weakly asymptotically empirical beliefs. The smallest and largest pure equilibria of the stage game Γ, $\underline{e}$ and $\overline{e}$, satisfy a.s. that $\underline{e} \leq \tilde{\xi} \leq \overline{e}$ and $\delta_{\underline{e}} \leq_{st} \tilde{\mu} \leq_{st} \delta_{\overline{e}}$ for all subsequential limits $\tilde{\mu}$ of $\{\mu(h_t)\}$ and $\tilde{\xi}$ of $\{\xi(h_t)\}$.*

Proof: Immediate from Lemma 5. ■

Theorem 5 *(Purification Setup) Let Γ be a supermodular game. Let $\{\delta_k\}$ be a sequence in $\mathbf{R}_+$ with $\delta_k \to 0$. For each k let (μ_k, ξ_k) be a myopic system of behavior and beliefs with monotone, weakly asymptotically empirical beliefs. Then there is a subsequence $\{\delta_l\}$ and a sequence of bounds $(\underline{e}_l)$ and $(\overline{e}_l)$ such that: a) For all l, $\underline{e}_l \leq \tilde{\xi}_l \leq \overline{e}_l$ and $\delta_{\underline{e}_l} \leq_{st} \tilde{\mu}_l \leq_{st} \delta_{\overline{e}_l}$ a.s. for all subsequential limits $\tilde{\mu}_l$ of $\{\mu_l(h_t)\}$ and $\tilde{\xi}_l$ of $\{\xi_l(h_t)\}$. b) The limits $\underline{e} = \lim \underline{e}_l$, $\overline{e} = \lim \overline{e}_l$ exist and are pure equilibria of the game Γ^0.*

Proof: Let $\beta' : S \times \Omega \times \mathbf{R}_+ \to S$ be the best response correspondence $\beta'(s, \omega, \delta) = \times_{i \in N} \text{argmax}_{s_i \in S_i} u_i^{\delta}(\omega_i)(s_i, s_{-i})$. For each k, Lemma 5 provides $\omega'_k, \omega''_k \in \Omega$, $\underline{e}_k \in \mathcal{E}(\omega'_k)$ and $\overline{e}_k \in \mathcal{E}(\omega''_k)$ such that $\underline{e}_k, \overline{e}_k$ are in the conditions of part a). The sets S and W are compact so there are convergent subsequences $\{(\underline{e}_l, \overline{e}_l, \omega'_l, \omega''_l)\}$. Say $(\underline{e}, \overline{e}, \omega', \omega'') = \lim(\underline{e}_l, \overline{e}_l, \omega'_l, \omega''_l)$. By the Maximum Theorem, the correspondence β' is upper hemicontinuous and has thus a closed graph. The sequences $\{(\underline{e}_l, \underline{e}_l, \omega'_l, \delta_l)\}$ and $\{(\overline{e}_l, \overline{e}_l, \omega''_l, \delta_l)\}$ are convergent sequences in the graph of of β'. Then, the limits $\underline{e}$ and $\overline{e}$ are fixed points of $s \mapsto \beta'(s, \tilde{\omega}, 0)$ for $\tilde{\omega} = \omega', \omega''$, respectively. The correspondence $s \mapsto \beta'(s, \tilde{\omega}, 0)$ is the pure strategy aggregate best response correspondence of game Γ^0. ■

5 Justifying Monotone Beliefs

A first justification for monotone beliefs is that they are "self enforcing": if players have monotone beliefs then play will be monotone; so beliefs will be "right" in being monotone. A second justification is that Bayesian updating will produce weakly monotone beliefs, as long as priors are concentrated on a totally ordered set of strategies.

Consider a player making inferences about her opponents' play in a repeated game. Let X be a complete lattice, let $\Pi \subset \mathcal{P}(X)$ be a collection of (Borel) probability measures over X. An element $\pi \in \Pi$ is a possible (correlated) strategy by the player's opponents. Let our player be endowed with a prior distribution $\eta \in \mathcal{P}(\Pi)$ over the possible strategies employed by her opponents. Assume that η has full support. The resulting beliefs about the player's opponents is $\mu \in \mathcal{P}(X)$ where $\mu(B) = \int_\Pi \pi(B) d\eta(\pi)$ for all events $B \subset X$.

Assume that after an event $E \subset X$ occurs, the player updates her beliefs by Bayes' rule. The updated posterior $\eta\backslash_E \in \mathcal{P}(\Pi)$ is

$$\eta\backslash_E(B) = \frac{\int_B \pi(E) d\eta(\pi)}{\int_\Pi \pi(E) d\eta(\pi)},$$

whenever $\int_\Pi \pi(E) d\eta(\pi) > 0$. The resulting updated beliefs $\mu\backslash_E \in \mathcal{P}(X)$ are defined by $\mu\backslash_E(B) = \int_\Pi \pi(B) d\eta\backslash_E(\pi)$. When $\int_\Pi \pi(E) d\eta(\pi) = 0$, $\mu\backslash_E$ is arbitrary.

Proposition 1 *Let* $E = [\sup \mathit{supp}\, \mu, \sup X]$. *If* Π *is totally ordered by first order stochastic dominance, then* $\mu \leq_{st} \mu\backslash_E$.

Proof: Only if $\int_\Pi \pi(E) d\eta(\pi) > 0$ there is something to prove. Let $B \subset \Pi$ be an increasing, measurable set. Because Π is a chain under first order stochastic dominance, $\pi(E) \leq \hat{\pi}(E)$ for all $\pi \in B^c$, $\pi \in B$ (since $\pi \leq_{st} \hat{\pi}$). Hence $\int_B \pi(E) d\eta(\hat{\pi}) \leq \int_B \hat{\pi}(E) d\eta(\hat{\pi})$, so $\pi(E)\eta(B) \leq \int_B \hat{\pi}(E) d\eta(\hat{\pi})$ for any $\pi \in B^c$. Similarly, $\eta(B)\int_{B^c}\pi(E)d\eta(\pi) \leq \eta(B^c)\int_B \hat{\pi}(E)d\eta(\hat{\pi})$. Then,

$$\eta(B)\frac{\int_{B^c}\pi(E)d\eta(\pi)}{\int_\Pi \pi(E)d\eta(\pi) > 0} \leq \eta(B^c)\frac{\int_B \hat{\pi}(E)d\eta(\hat{\pi})}{\int_\Pi \pi(E)d\eta(\pi) > 0},$$

which implies that $\eta(B)\eta\backslash_E(B^c) \leq \eta(B^c)\eta\backslash_E(B)$.

Suppose, by way of contradiction, that $\eta\backslash_E(B) < \eta(B)$. Since then $\eta\backslash_E(B) < 1$ it must be that $\eta\backslash_E(B^C) > 0$. So, $0 < \eta(B)\eta\backslash_E(B^c) \leq \eta(B^c)\eta\backslash_E(B)$. Now, $\eta\backslash_E(B) < \eta(B)$ implies that $\eta\backslash_E(B^c) \leq \eta(B^c)$. But then, $\eta\backslash_E(B)+\eta\backslash_E(B^c) \leq \eta(B)+\eta(B^c) = 1$, a contradiction. This proves that $\eta(B) \leq \eta\backslash_E(B)$ for an arbitrary increasing B, i.e. that $\eta \leq_{st} \eta\backslash_E$.

Let F be an increasing event in X, then the map $\pi \mapsto \pi(F)$ is monotone increasing, since Π is linearly ordered by first order stochastic dominance. Then, $\eta \leq_{st} \eta\backslash_E$ implies that $\int_\Pi \pi(F)d\eta(\pi) \leq \int_\Pi \pi(F)d\eta\backslash_E(\pi)$. By the definition of the player's beliefs over X, then, $\mu(F) \leq \mu\backslash_E(F)$. So, $\mu \leq \mu\backslash_E$. ■

References

Aliprantis, C. D., and K. C. Border (1999): *Infinite Dimensional Analysis*. Springer-Verlag.

Benaim, M., and M. W. Hirsch (1999): "Mixed Equilibria and Dynamical Systems Arising from Fictitious Play in Perturbed Games," *Games and Economic Behavior*, 29(1-2), 36–72.

Ellison, G., and D. Fudenberg (1999): "Learning Purified Mixed Equilibria," Forthcoming in Journal of Economic Theory.

Fudenberg, D., and D. M. Kreps (1993): "Learning Mixed Equilibria," *Games and Economic Behavior*, 5(3), 320–67.

Fudenberg, D., and D. K. Levine (1998): *The theory of learning in games.* MIT Press Series on Economic Learning and Social Evolution, vol. 2. Cambridge and London: MIT Press.

Kamae, T., U. Krengel, and G. O'Brien (1977): "Stochastic Inequalities on Partially Ordered Spaces," *Annals of Probability*, 5(6), 899–912.

Kaniovski, Y. M., and H. P. Young (1995): "Learning Dynamics in Games with Stochastic Perturbations," *Games and Economic Behavior*, 11(2), 330–63.

Milgrom, P., and J. Roberts (1990): "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58(6), 1255–1277.

Milgrom, P., and C. Shannon (1994): "Monotone Comparative Statics," *Econometrica*, 62(1), 157–180.

Sobel, M. J. (1988): "Isotone Comparative Statics in Supermodular Games," mimeo, SUNY at Stony Brooks.

Topkis, D. M. (1978): "Minimizing a Submodular Function on a Lattice," *Operations Research*, 26(2), 305–321.

Topkis, D. M. (1979): "Equilibrium Points in Nonzero-Sum n-Person Submodular Games," *SIAM Journal of Control and Optimization*, 17(6), 773–787.

——— (1998): *Supermodularity and Complementarity.* Princeton University Press.

Vives, X. (1990): "Nash Equilibrium with Strategic Complementarities," *Journal of Mathematical Economics*, 19(3), 305–321.

Zhou, L. (1994): "The Set of Nash Equilibria of a Supermodular Game Is a Complete Lattice," *Games and Economic Behavior*, 7(2), 295–300.