Mixed Equilibria in Games of Strategic Complements are Unstable [*]

Federico Echenique[†]and Aaron Edlin[‡]
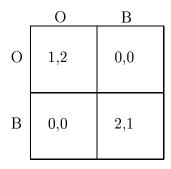
October 25, 2001

Abstract

In games with strict strategic complementarities, properly mixed Nash equilibria—equilibria that are not in pure strategies—are unstable for a broad class of learning dynamics.
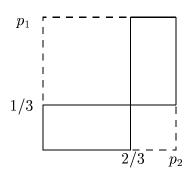
JEL Classification: C72, C73.

Figure 1: Battle of the Sexes

1 Introduction

Many economic settings can be modeled as games of strict strategic complementarities. Here we show that for such games, properly-mixed Nash equilibria (i.e., Nash equilibria that are not in pure strategies, referred to as PMNE hereafter) are unstable under a broad class of learning processes, and so are not likely to be good predictions of play.

We shall give a flavor of our results using the "Battle of the Sexes" game in Figure 1. Players 1 and 2 each simultaneously choose an element from $\{O, B\}$. Payoffs are specified in the bimatrix to the left. Let p_i be the probability with which player i selects O. The best responses are 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 $p_1 = 1$. 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 $(1, 1)$.

Battle of the Sexes is a game of strict strategic complementarities. To check for complementarities we need an order on players' strategies: say that O is "larger" than B. Then the players' best responses are monotone increasing. For example, if 1 increases her strategy from B to O then 2 increases her best response from B to O. That best-responses are increasing is only necessary for the game to have strict strategic complementarities, but for now it will suffice.

Suppose that our prediction of play for Battle of the Sexes is the PMNE $(p_1, p_2) =$

$(1/3, 2/3)$; but suppose that the players' beliefs about their opponent's play are slightly wrong. In particular, 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 perturbed beliefs are arbitrarily close to the equilibrium beliefs. Now, as can be seen from the best-response functions in Figure 1, given these beliefs both players will select O with probability 1.

We shall now argue that, if the players repeatedly play Battle of the Sexes, each time best-responding (myopically) to their beliefs, they will move away from our original prediction. Both players observed their opponent choosing O, so they might infer that they were right in giving O larger weight than what $(1/3, 2/3)$ does. So they might "update" their beliefs and give the larger action, O, more weight after observing larger play (we 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 from the Nash equilibrium beliefs $(1/3, 2/3)$—so $(1/3, 2/3)$ is unstable.

Note that there is nothing non-generic or knife-edge about the perturbations we consider, it is plausible that players would end up with perturbed beliefs like those above. As players start myopically playing the PMNE, they will play (O, O) with probability $2/9$, and any finite sequence of (O, O) play has positive probability. It is plausible that, after observing several rounds of (O, O), players change their beliefs in the direction of giving "my opponent plays O" larger probability. Our point is that deviations like these will, under our assumptions, not be corrected by subsequent play.

Games with strict strategic complementarities are common in economics—Milgrom and Roberts (1990), Topkis (1998) and Vives (1999) contain many economic examples of games with strategic complementarities, and a minor strengthening of the complementarities conditions in their models yields a game with *strict* complementarities. [1] One example is the Bertrand oligopoly model with differentiated products

---

[1] We shall be more precise further on, let us clarify for now that we do not need strict supermodularity in players' own strategies, only strictly increasing differences in players' own and opponents' strategies.

(Tirole 1988, p. 280). Pure strategies in this model are easy to analyze both using calculus-based techniques (as in Tirole), or lattice-theoretic techniques (Milgrom and Roberts 1990). Mixed strategies, on the other hand, are difficult to work with because strategy spaces are infinite. Typically, complementarities provide good tools for working with pure-strategy equilibria (see Topkis (1998) and Vives (1999) for an overview of these tools), but it is difficult to work with mixed-strategy equilibria in models without a small number of strategies; so our results provide a convenient justification for focusing on the pure-strategy equilibria of many models of interest to economists.

The literature on learning mixed-strategy equilibria has mostly studied 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 (2000)).

Fudenberg and Kreps (1993) show that, under appropriate assumptions on learning dynamics, play in the long run *can* resemble a PMNE. They show that any equilibrium is a steady state of their learning dynamics. Our point is that PMNE are *unstable* steady states, and therefore may not be good predictions.

A corollary of our results is that, in generic 2X2 games, PMNE are either unique or unstable. Fudenberg and Kreps (1993) show that if a 2X2 game has a unique equilibrium which is a PMNE, this PMNE is stable; our results complete the picture for the rest of the 2X2 games. Basically, then, 2X2 games are either isomorphic to "matching pennies" or have no stable PMNE.

The paper is organized as follows. In Section 2 we give some basic definitions. In Section 3 we describe the learning model that we will use. Section 4 contains the main result, and gives some intuition for its proof. In Section 5 we consider purified mixed-strategy equilibria. In Section 6 we justify the assumption of weakly monotone beliefs. In Section 7 we prove Theorems 1 and 2.

2 Preliminaries

2.1 Lattice-theoretic definitions

The definitions in this subsection, and the application of lattice theory to game theory and economics, is discussed at length in Topkis (1998) and Vives (1999). 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. Note that any finite lattice is complete. 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 $x \vee_X y$ 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. The **order-interval topology** on a lattice is the topology obtained by taking the closed order intervals as a sub-basis of the closed sets. In Euclidean spaces the order-interval topology coincides with the usual topology. A lattice is complete if and only if it is compact in its order-interval topology.

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

2.2 Probability measures and first-order stochastic dominance

Let X be a lattice endowed with a topology finer than its order-interval topology. [2] Let $\mathcal{P}(X)$ denote the set of (Borel) probability measures over X. A subset $A \subset X$ is **increasing** if, for all $x \in A$, $y \in X$ and $x \preceq y$ imply $y \in A$. For example, if $X \subset \mathbf{R}$, and $\mathbf{R}$ has the usual order, A is increasing if and only if it is an open or closed half-interval, i.e. either of the form $[x, \infty)$ or (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 all increasing sets $A \subset X$, $\mu(A) \leq \nu(A)$.

Let X be a complete lattice. The **support of** $\mu \in \mathcal{P}(X)$ is the intersection of all closed probability-one events; it is denoted $\operatorname{supp}(\mu)$. [3] For any $x \in X$, the

---

[2] To check that a topology τ is finer than the order-interval topology it is sufficient to prove that closed intervals are closed under τ—the order-interval topology is the coarsest topology for which order intervals are closed. For example, if $\mathcal{P}(X)$ is ordered by first-order stochastic dominance it is easy to show that its order-interval topology is coarser than its weak topology, see Remark 2 below.

[3] Defined in this way, every measure has a support, in contrast with other definitions of support,

singleton $\{x\}$ is measurable, as it is a closed order interval, and thus closed in the order-interval topology. Let $\delta_x \in \mathcal{P}(X)$ denote the degenerate probability measure that gives probability one to x. A probability measure $\mu \in \mathcal{P}(X)$ is **properly mixed** if supp (μ) is not a singleton.

Remarks:

1. If $\mu \in \mathcal{P}(X)$ is properly mixed then there is no $x \in X$ with full measure as a subset of X, i.e. there is no $x \in X$ such that $\mu(\{x\}) = 1$. To see this, let $\{x\}$ be of full measure, then supp $(\mu) \subset \{x\}$ because $\{x\}$ is closed. Since supp (μ) is nonempty, supp $(\mu) = \{x\}$, and μ is not properly mixed.

2. If $\mathcal{P}(X)$ is ordered by first-order stochastic dominance, closed order-intervals are weakly closed. That is, for any $\mu, \mu' \in \mathcal{P}(X)$, $[\mu, \mu']$ is weakly closed. Let $\mathcal{A}$ be the collection of all increasing subsets of X. Then, the order-interval $[\mu, \mu']$ is:

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

But for all x, $\{p \in \mathcal{P}(X) : \mu(E) \leq p(E)\}$ and $\{p \in \mathcal{P}(X) : p(E) \leq \mu'(E)\}$ are weakly closed sets (Aliprantis and Border (1999) Theorem 14.6). Then, order-intervals are weakly closed.

2.3   Complementarities, strategies and beliefs

A game in normal form is described by $(N, \{(S_i, u_i) : i \in N\})$, where N is a finite set of players, and each player $i \in N$ is endowed with a strategy space S_i and a payoff function $u_i : S = \times_{i \in N} S_i \to \mathbf{R}$. Let n be the number of players in N.

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

1. S_i is a complete lattice;

---

see e.g. Royden (1988).

2. u_i is bounded, $s_i \mapsto u_i(s_i, s_{-i})$ is supermodular for all $s_{-i} \in S_{-i}$, $(s_i, s_{-i}) \mapsto u_i(s_i, s_{-i})$ has increasing differences, and

3. $s_i \mapsto u_i(s_i, s_{-i})$ is upper semicontinuous for all $s_{-i} \in S_{-i}$.

Γ is a **game of strict strategic complementarities (GSSC)** if it is a GSC and, in addition, $(s_i, s_{-i}) \mapsto u_i(s_i, s_{-i})$ has strictly increasing differences.

The **mixed extension** of a game $\Gamma = (N, \{(S_i, u_i) : i \in N\})$ is the game $(N, \{(\mathcal{P}(S_i), U_i) : i \in N\})$, where each player i is allowed to choose any mixed strategy $\sigma_i \in \mathcal{P}(S_i)$, and where a mixed-strategy profile $\sigma = (\sigma_1, \ldots \sigma_n)$ gives player i the payoff $U_i(\sigma) = \int_S u_i(s) d\times_{i \in N} \sigma_i(s_i)$. We shall identify a vector of probability measures $\sigma = (\sigma_1, \ldots \sigma_n) \in \times_{i=1}^n \mathcal{P}(S_i)$ with the corresponding product measure in $\mathcal{P}(S)$. So, for example, $U_i(\sigma) = \int_S u_i(s) d\sigma(s)$.

Player i's beliefs about her opponents' play is represented by a probability distribution $\mu_i \in \mathcal{P}(S_{-i})$. Belief space is then $\Psi = \times_{i \in N} \mathcal{P}(S_{-i})$. A mixed-strategy Nash equilibrium $\sigma^* = (\sigma_1^*, \ldots \sigma_n^*)$ is a situation where i chooses the strategy σ_i^* optimally given that her beliefs about opponents' play is "right", that is i's belief is $\sigma_{-i}^* \in \mathcal{P}(S_{-i})$. So there is a natural "copy" of σ^* in belief space, the vector $(\sigma_{-1}^*, \sigma_{-2}^*, \ldots \sigma_{-n}^*) \in \Psi$. More generally, to each mixed strategy profile $\sigma = (\sigma_1, \ldots \sigma_n) \in \times_{i=1}^n \mathcal{P}(S_i)$, there corresponds beliefs $\sigma_{-i} \in \mathcal{P}(S_{-i})$ for player i. We shall denote by $\psi : \times_{i=1}^n \mathcal{P}(S_i) \to \Psi$ the map from strategy profiles to beliefs: $\psi_i \sigma \equiv \sigma_{-i}$, and $\psi \sigma \equiv (\psi_i \sigma)_{i \in N} = (\sigma_{-1}, \sigma_{-2}, \ldots \sigma_{-n})$.

The set of player i's possible beliefs $\mathcal{P}(S_{-i})$ is endowed with the weak topology and the first-order stochastic dominance order. Belief space, $\Psi = \times_{i \in N} \mathcal{P}(S_{-i})$ is endowed with the product topology and the product order.

3 Learning Model

Learning takes place through repeated play of a stage game, $\Gamma = (N, \{(S_i, u_i) : i \in N\})$. In each stage, player i observes (privately) a signal $\omega_i \in \Omega_i$, given some probability space $(\Omega_i, \mathcal{F}_i, p_i)$. These signals are not payoff-relevant; player i uses them as randomization mechanisms, making her choice of a pure strategy conditional on the

realization of the signal. If $(\Omega_i, \mathcal{F}_i, p_i)$ is rich enough this does not restrict her choice of randomization over pure strategies. The set of all signal profiles is $\Omega = \times_{i \in N} \Omega_i$.

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 repeated-game strategy $\xi_i : \Omega_i \times H \to S_i$, and is endowed with repeated-game 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 stage $t+1$ of the game. Given this assessment, and the realization of ω, she chooses a stage-game strategy $\xi_i(\omega, h^t) \in S_i$. Note that we allow player i to believe that her opponents' play is correlated—correlated beliefs arise naturally even if players mix independently, see e.g. Fudenberg and Kreps (1993).

Let $\xi = (\xi_i)_{i \in N}$ be a collection of strategies for all players and $\mu = (\mu_i)_{i \in N}$ be a collection of beliefs. The pair (ξ, μ) is a **system of behavior and beliefs**. Note that $\xi : \Omega \times H \to S$ and $\mu : H \to \Psi$.

Player i's best-response correspondence $\beta_i : \mathcal{P}(S_{-i}) \twoheadrightarrow S_i$ is defined by

$$\beta_i(\nu_i) = \operatorname{argmax}_{\tilde{s}_i \in S_i} \int_{S_{-i}} u_i(\tilde{s}_i, s_{-i}) d\nu_i(s_{-i}).$$

So, $\beta_i(\nu_i)$ is the set of best responses to beliefs $\nu_i \in \mathcal{P}(S_{-i})$ about opponents' play. The set of best responses to strategy s_{-i} is then $\beta_i(\delta_{s_{-i}})$. The players' joint best-response correspondence is $\beta : \Psi \twoheadrightarrow S$, defined as $\beta(\nu) = (\beta_i(\nu_i))_{i \in N}$.

Definition 2 *A system of behavior and beliefs (ξ, μ) is **myopic** if for all $i \in N$, $h^t \in H$ and $\omega_i \in \Omega_i$,*

$$\xi_i(\omega_i, h^t) \in \beta_i(\mu_i(h^t)) = argmax_{\tilde{s}_i \in S_i} \int_{S_{-i}} u_i(\tilde{s}_i, s_{-i}) \mu_i(h^t)(ds_{-i}).$$

The assumption of myopic behavior is very common in the literature on learning in games. It is restrictive because it implies that players do not attempt to manipulate the future behavior of their opponents—they simply maximize current payoffs based on current beliefs. Myopic behavior is usually justified by assuming that, in each

7

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).

Our results on learning will rely on an important assumption: if a player has a certain prediction about her opponents' play, and then observes play that is weakly larger than any strategy she attached a positive probability to, then she will "update" her beliefs about opponents' behavior and predict weakly larger play. This requirement on beliefs we call weak monotonicity. To be precise:

Definition 3 *Beliefs μ are **weakly monotone** if, for all $i \in N$, and $h^t \in H$,*

$$\left(\sup \left[supp \ \mu_i(h^t)\right] \leq s_{-i}^\tau \ \ for \ \ \tau = t+1, \ldots T\right) \Rightarrow \left(\mu_i(h^t) \leq_{st} \mu_i(h^T)\right).$$

Where h^T is any history that coincides with h^t in periods $\tau = 0, \ldots t$ and where i's opponents play s_{-i}^τ in periods $\tau = t+1, \ldots T$.

The idea behind this definition is that, if $\mu_i(h^t)$ gives i's beliefs at time $t+1$ and history h^t, and if play at times $t+1, t+2, \ldots T$ is weakly larger than any play i believed possible at time t, then i will have weakly larger beliefs at time T.

Weak monotonicity is the only condition we need on beliefs, and it is rather mild. Beliefs in Cournot best-response dynamics satisfy weak monotonicity. We show in Section 6 that fictitious-play beliefs, and beliefs updated by Bayes' rule, satisfy weak monotonicity (see Fudenberg and Levine (1998) for definitions and discussion of these learning models). As a simple justification for weakly monotone beliefs, note that, if beliefs are weakly monotone then play will, under our assumptions, be monotone, so beliefs are "right" in being weakly monotone. Monotonicity is then, in a sense, self enforcing. We have used the qualifier "weak" to differentiate our condition from the stronger requirement that any larger play produce larger beliefs (e.g. in Hopenhayn and Prescott (1992)). The stronger condition does not arise naturally in standard learning models.
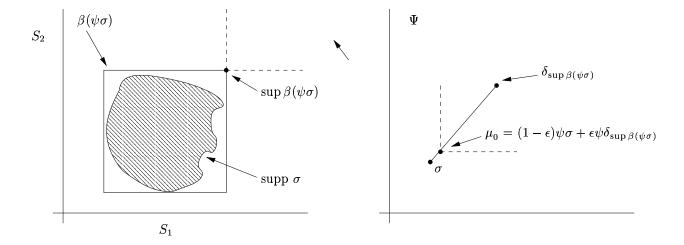
Figure 2: Unstable mixed equilibrium

4 Instability of Mixed Equilibria

We now prove that, at any PMNE σ of a GSSC, there are arbitrarily small perturbations that set off learning dynamics so that strategies are always outside of a neighborhood of σ. Ours is an instability result: these small perturbations from σ are never "corrected" by subsequent dynamics. The perturbation takes the form of slightly wrong beliefs. [4] We shall first give a heuristic argument for why the instability obtains.

4.1 Intuition for the instability result

Let $(\{1, 2\}, \{(S_i, u_i) : i = 1, 2\})$ be a two-player GSSC. If each strategy space is a subset of $\mathbf{R}$, we can represent the joint strategy spaces in $\mathbf{R}^2$—see the drawing on the left in Figure 2. Let σ be a mixed-strategy equilibrium where both players select a properly mixed strategy. The set of pure-strategy best responses to σ, $\beta(\psi\sigma)$, is a subcomplete sublattice (Topkis 1998), for example the rectangle in Figure 2. The support of σ must lie in $\beta(\psi\sigma)$, and likewise $\sup \beta(\psi\sigma) \in \beta(\psi\sigma)$.

---

[4] As an alternative, we could perturb behavior. It seems that, of the two equilibrium assumptions, that behavior is rational given beliefs and that beliefs are "correct", it is the second that most of the non-equilibrium literature seeks to weaken. In any case, we obtain the same results if we perturb behavior instead of beliefs.

Lets perturb beliefs in the direction of the largest element in $\beta(\psi\sigma)$. Consider beliefs $\mu^0 = (1-\epsilon)\psi\sigma + \epsilon\psi\delta_{\sup\beta(\psi\sigma)}$. With beliefs μ^0, player i mixes equilibrium beliefs $\psi_i\sigma$ with degenerate beliefs that i's opponent will play their largest best responses to their equilibrium beliefs. This was the perturbation we used in "Battle of the Sexes" in the Introduction.

Observe that $\psi\sigma <_{st} \mu^0$, and that, by choosing $\epsilon > 0$ small enough, μ^0 can be taken arbitrarily close to $\psi\sigma$. The support of μ^0 lies in $\beta(\psi\sigma)$, which is crucial for our results. We represent the situation in Figure 2. On the left is the strategy space, where any best response to μ^0 must lie to the north-east of $\sup\beta(\psi\sigma)$. On the right is belief space, Ψ—we represent Ψ as a subset of the plane, which is inaccurate and just a means of visualizing the ideas behind our results.

Consider any learning dynamics that starts off at the perturbed beliefs μ^0. Because complementarities are strict, any best response to perturbed beliefs μ^0 is (weakly) larger than any best response to $\psi\sigma$. In particular, then, play is weakly larger than any element in the support of μ^0. Then, if beliefs are monotone, "updated" beliefs, after observing first-period play, are weakly larger than μ^0. Now the argument follows by induction: if play in all periods $1, 2, \ldots t-1$ is larger than $\sup\beta(\psi\sigma)$, then period t beliefs must be weakly larger than μ^0, and period t play must be weakly larger than $\sup\beta(\psi\sigma)$. So, $\mu^t = \mu(h^t)$ is always to the north-east of μ^0, and therefore beliefs never approach $\psi\sigma$.

If the perturbation $\mu^0 = (1-\epsilon)\psi\sigma + \epsilon\psi\delta_{\sup\beta(\psi\sigma)}$ seems arbitrary, note that any beliefs in the interval $\left[\mu^0, \psi\delta_{\sup\beta(\psi\sigma)}\right]$ would work.

The crucial components of our argument are then:

1. Because σ is properly mixed, there is a wedge in pure strategies, so there are perturbed beliefs μ^0 that are larger than $\psi\sigma$, while still having support in $\beta(\psi\sigma)$. These perturbed beliefs can be taken arbitrarily close to $\psi\sigma$.

2. Strict complementarities between players' choices provokes an "overshooting" response to larger beliefs, so that any best response to μ^0 is larger than $\sup\beta(\psi\sigma)$.

3. By monotone beliefs, the initial deviation toward larger play is reinforced. But note that the monotonicity used is weak; since play is weakly larger than any

element in the support of μ^0, beliefs in each moment t must be weakly larger than μ^0.

4.2   Main result

Here we formalize the heuristic argument just given, but we defer the proof to section 7.

Definition 4 *Let $(N \{(S_i, u_i) : i \in N\})$ be a game. A mixed-strategy profile σ is **unstable** if, for every weak neighborhood V of $\psi\sigma$ in Ψ, there is $\mu' \in V$ such that any myopic system of behavior and beliefs (μ, ξ) with weakly monotone beliefs and $\mu^0 = \mu'$, remains outside of a neighborhood of $\psi\sigma$. That is, there is a neighborhood W of $\psi\sigma$ such that, for all $t \geq 1$, $\mu(h^t) \notin W$.*

That σ is unstable means that there are arbitrarily close perturbed beliefs μ^0 such that, if learning starts at these perturbed beliefs, then beliefs never approach σ. The definition of unstable equilibrium is an adaptation to the present context, of the definition of asymptotic instability used in the dynamical systems literature (Hirsch and Smale 1974).

Theorem 1 *Let Γ be a GSSC, and σ be a Nash Equilibrium of the mixed extension of Γ. If at least two players' strategies in σ are properly mixed, then σ is unstable.*

Proof: See section 7.

   Remark. In finite games, for generic payoffs, there are no PMNE where only one player selects a properly mixed strategy (because best responses to pure strategies are generically unique). In many non-finite games, it is not hard to rule out that only one player selects a properly mixed strategy.

   Theorem 1 has a simple consequence for 2X2 games. For generic payoffs, 2X2 games either have a unique Nash equilibrium, or two pure equilibria and one PMNE. In this last case, it is easy to order strategies so that the game is a GSSC. Then, we get

Corollary 1 *For generic payoffs, PMNE in 2X2 games are either unique or unstable.*

11

Generically, then, a 2X2 game is either isomorphic to Matching Pennies or its mixed equilibrium is unstable. [5] For 2X2 games, Fudenberg and Kreps (1993) show that, when the PMNE is the unique equilibrium, it is globally stable. This paper completes the picture for all other 2X2 games.

5 Purified Mixed Equilibria

A textbook criticism of PMNE goes like this (see e.g. Osborne and Rubinstein (1994)): In a PMNE, each player i is required to randomize in exactly the way that leaves the other players indifferent between the elements in the support of their equilibrium strategies. But i has no reason to randomize in this way, precisely because i too is indifferent between the elements in the support of her equilibrium strategies. The standard response to this criticism is Harsanyi's Purification Theorem—if we introduce a small amount of incomplete information, then pure-strategy equilibrium behavior can resemble the original PMNE.

In this section we show that "purified" PMNE in GSSC are also unstable. So, if σ is a PMNE then sufficiently close pure-strategy Nash equilibria of the game with incomplete information are unstable. [6] In other words, then, purification addresses the standard critique but not an instability critique.

Consider the Battle of the Sexes game from the Introduction, and let σ be its PMNE. We shall first introduce incomplete information and "purify" σ. Let each player i receive a payoff-relevant signal ω_i. The signals are independently uniformly distributed on $[0,1]$. The players' payoffs are in Figure 3; $\eta > 0$ is the parameter that controls the importance of the signals, we shall call the game in Figure 3 the η-augmented game.

It is easy to see that the pair of (pure) strategies $(\int_1, \int_2)$,

$$\int_1(\omega_1) = \begin{cases} O \text{ if } \omega_1 \leq 2/3 \\ B \text{ if } \omega_1 > 2/3 \end{cases} \quad \int_2(\omega_2) = \begin{cases} O \text{ if } \omega_2 > 2/3 \\ B \text{ if } \omega_2 \leq 2/3 \end{cases},$$

---

[5] For generic payoffs, if a 2X2 game has a unique equilibrium, and this equilibrium is a PMNE, then there is a re-labeling of each player's strategies into {Heads, Tails} so that the resulting preferences over {Heads, Tails} equal the Matching Pennies preferences (i.e. one player wants to match and the other wants to mis-match). This re-labeling is an isomorphism.

[6] A number of papers on learning mixed strategy equilibria have focused on purified mixed strategies, see for example Fudenberg and Kreps (1993) and Ellison and Fudenberg (2000).
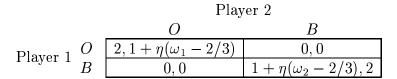
		O	B
Player 1	O	$2, 1 + \eta(\omega_1 - 2/3)$	$0, 0$
	B	$0, 0$	$1 + \eta(\omega_2 - 2/3), 2$

Figure 3: Augmented Battle of the Sexes

is a Nash equilibrium of the augmented game, no matter the value of η. Note that, for almost all ω_i, player i is selecting a strict best response to $\int_{-i}$.

The distribution of $(\int_1(\omega_1), \int_2(\omega_2))$ is the same as the PMNE distribution in the original Battle of the Sexes. This is a particularly nice example of a purification; we can be as close as we want to the original game by taking η small enough, get the same prediction as the PMNE, and avoid assuming that players select arbitrarily among a set of indifferent strategies.

Now, given $\epsilon > 0$, consider a perturbation $\mu^0 = (1 - \epsilon)\psi\sigma + \epsilon\psi\delta_{(O,O)}$. Doing some algebra, it turns out that, if

$$\eta < \min\left\{\frac{2\epsilon}{2 + \epsilon}, \frac{9\epsilon}{1 - \epsilon}\right\},$$

then, no matter the value of ω_i, player i's best response to beliefs μ_i^0 is to play O. We can then repeat the argument in the Introduction (and in Section 4.1) that play only reinforces the initial perturbed beliefs. So if behavior is myopic and beliefs are weakly monotone, the purified equilibrium is unstable.

In this example, for each of our perturbations $\mu^0 = (1 - \epsilon)\psi\sigma + \epsilon\psi\delta_{\sup\beta(\psi\sigma)}$, there is $\underline{\eta}$ such that if $\eta < \underline{\eta}$ then the purified equilibrium in the η-augmented game is unstable to the μ^0 perturbation. Note that the order of limits matters: We do not say that, for η small enough, the purified equilibrium is unstable to arbitrarily small perturbations. [7]

Now we present a general result for purified PMNE. For simplicity, we assume finite strategy spaces. The setup is from Fudenberg and Tirole's (1991) presentation of Harsanyi's Purification Theorem. Let $\Gamma^0 = (N, \{(S_i, g_i) : i \in N\})$ be a finite game. For $\eta > 0$, $\Gamma^\eta = (N, \{(\mathcal{S}_i, u_i^\eta) : i \in N\})$ is the η-**augmented game**, where each player

---

[7]This may be due to our brute-force approach to dealing with randomizations—we control best responses at *all* values of ω_i. It may be possible to do better with more sophisticated methods.

i is endowed with type space $\Omega_i = [-1,1]^{\#S_i}$ and selects a strategy $\smallint : \Omega_i \to S_i$. Types $\omega_i \in \Omega_i$ are drawn independently according to probability distribution p_i. Payoffs are $u_i^\eta(s_i, s_{-i}, \omega_i) = g_i(s_i, s_{-i}) + \eta\omega_i(s_i)$.

Harsanyi's Purification Theorem says that, in generic finite games, for any PMNE σ of Γ^0, there is a collection $(\smallint^\eta)_{\eta>0}$ such that: a) for all η, $\smallint^\eta$ is a (pure) equilibrium of Γ^η; b) the distributions of $\smallint^\eta(\omega)$ converge to σ as $\eta \to 0$. We shall call $(\Gamma^\eta, \smallint^\eta)_{\eta>0}$ a **purification sequence of** σ.

We need to control the relative sizes of the perturbations in beliefs and the augmentation of Γ^0. The following definition helps us do that.

Definition 5 *Let* $(N \{(S_i, u_i) : i \in N\})$ *be a game. A mixed-strategy profile* σ *is* **unstable to an** ϵ**-perturbation** *for* $\epsilon > 0$ *if there is a strategy profile* σ' *such that any myopic system of behavior and beliefs* (μ, ξ) *with weakly monotone beliefs and* $\mu^0 = (1 - \epsilon)\psi\sigma + \epsilon\psi\sigma'$ *remains outside of a neighborhood of* $\psi\sigma$. *That is, there is a neighborhood* W *of* $\psi\sigma$ *in* Ψ *such that, for every* t, $\mu(h^t) \notin W$.

Theorem 2 *Let* Γ^0 *be a finite GSSC, and* σ *be a Nash Equilibrium of the mixed extension of* Γ^0 *such that at least two players' strategies in* σ *are properly mixed. Let* $(\Gamma^\eta, \smallint^\eta)$ *be a purification sequence of* σ. *For any* $\epsilon > 0$, *there is* $\underline{\eta} > 0$ *such that if* $\eta < \underline{\eta}$ *then* $\smallint^\eta$ *is unstable to an* ϵ*-perturbation.*

Proof: See section 7.

6 Justifying Weakly Monotone Beliefs

We show that, in two specifications that are common in learning models, beliefs are weakly monotone. In particular, we show in (1.) that fictitious-play beliefs are weakly monotone, and in (2.) that beliefs in Bayesian learning are weakly monotone.

1. We shall first describe a model of fictitious play following Chapter 2 in Fudenberg and Levine (1998). Let $(N, \{(S_i, u_i) : i \in N\})$ be a normal-form game. Fix a player i. Let $\kappa_i^0 : S_{-i} \to \mathbf{R}+$ be an exogenous "initial weight function" that is not identically zero. Player i's weight function is updated each period $t \geq 1$. So, if play in $t - 1$ is $\hat{s}_{-i}^{t-1}$, then the period-t weight function is

$$\kappa_i^t(s_{-i}) = \kappa_i^{t-1}(s_{-i}) + \chi_{\left\{s_{-i} = \hat{s}_{-i}^{t-1}\right\}}(s_{-i}). \tag{I}$$

The function $\chi_{\left\{s_{-i}=\hat{s}_{-i}^{t-1}\right\}}$ takes the value 1 if $s_{-i} = \hat{s}_{-i}^{t-1}$ and 0 if $s_{-i} \neq \hat{s}_{-i}^{t-1}$.

Now we can define **fictitious play beliefs**. Player i's beliefs at the null history are

$$\mu_i(h^0)(s_{-i}) = \frac{\kappa_i^0(s_{-i})}{\sum_{s_{-i}\in S_{-i}} \kappa_i^0(s_{-i})};$$

and, after history $h^t = ((s_i^1, s_{-i}^1), (s_i^2, s_{-i}^2)\ldots(s_i^t, s_{-i}^t))$,

$$\mu_i(h^t)(s_{-i}) = \frac{\kappa_i^t(s_{-i})}{\sum_{s_{-i}\in S_{-i}} \kappa_i^t(s_{-i})},$$

where we get the κ_i^t weight function recursively from κ_i^0 and h^t by formula I.

Proposition 1 Let $(N, \{(S_i, u_i) : i \in N\})$ be a game where each S_i is a finite lattice. Fictitious play beliefs in this game are weakly monotone.

Proof: Let h^t be a history, and $\sup [\operatorname{supp} \mu(h^t)] \leq s^\tau$, for $\tau = t + 1, \ldots T$. Fix a player i. For each subset A of S_{-i}, let $\kappa_i^t(A) = \sum_{s_{-i}\in A} \kappa_i^t(s_{-i})$.

Let $E \subset S_{-i}$ be an increasing set. *Case 1.* If $E \cap \operatorname{supp} \mu(h^t) = \emptyset$ then $\mu(h^t)(E) = 0 \leq \mu(h^t)$. *Case 2.* If $E \cap \operatorname{supp} \mu(h^t) \neq \emptyset$ then, because E is increasing, $\sup [\operatorname{supp} \mu(h^t)] \in E$. Then, also because E is increasing, $s^\tau \in E$, for $\tau = t + 1, \ldots T$. Then $\kappa_i^T(E) = (T - t) + \kappa_i^t(E)$, so

$$
\begin{aligned}
\mu(h^T)(E) &= \frac{\kappa_i^T(E)}{\kappa_i^T(S_{-i})} \\
&= \frac{(T-t)+\kappa_i^t(E)}{(T-t)+\kappa_i^t(S_{-i})} \\
&\geq \frac{\kappa_i^t(E)}{\kappa_i^t(S_{-i})} = \mu(h^t)(E);
\end{aligned}
$$

the inequality is because $x \mapsto (a + x)/(b + x)$ is a monotone increasing function when $a \leq b$. We have shown that $\mu(h^t)(E) \leq \mu(h^T)(E)$ for every increasing set E, so $\mu(h^t) \leq_{st} \mu(h^T)$. ∎

2. Now we show that Bayesian updating respects weak monotonicity. Let $\Pi_i \subset \mathcal{P}(S_{-i})$ be a set of possible (correlated) strategies by i's opponents. Suppose that i has beliefs $\mu_i \in \mathcal{P}(S_{-i})$ that she obtains from some prior distribution η_i over Π_i. So, $\eta_i \in \mathcal{P}(\Pi_i)$ and $\mu_i(B) = \int_{\Pi_i} \pi(B)d\eta_i(\pi)$ for all events $B \subset S_{-i}$.

Suppose that, after an event $E \subset S_{-i}$ occurs, i updates her beliefs by Bayes' rule. The updated posterior $\eta_i|_E \in \mathcal{P}(\Pi_i)$ is

$$\eta|_E(B) = \frac{\int_B \pi(E)d\eta_i(\pi)}{\int_{\Pi_i} \pi(E)d\eta_i(\pi)},$$

15

whenever $\int_{\Pi_i} \pi(E)d\eta_i(\pi) > 0$. The resulting updated beliefs $\mu_i|_E \in \mathcal{P}(S_{-i})$ are defined by $\mu_i|_E(B) = \int_{\Pi_i} \pi(B)d\eta_i|_E(\pi)$. When $\int_{\Pi_i} \pi(E)d\eta_i(\pi) = 0$, $\mu_i|_E$ is arbitrary.

Theorem 3 *Let S_{-i} be a complete lattice, and Let $E \subset [\sup supp\ \mu_i, \sup S_{-i}]$. If Π_i is totally ordered by first-order stochastic dominance, then $\mu_i \leq_{st} \mu_i|_E$.*

Proof: We shall drop the i-subindexes to simplify. Only if $\int_{\Pi} \pi(E)d\eta(\pi) > 0$ is there something to prove. First we shall prove that $\eta \leq_{st} \eta|_E$.

Let $D = [\sup supp\ \mu, \sup S_{-i}]$. Let $B \subset \Pi$ be an increasing, measurable set. Because Π is a chain under first-order stochastic dominance, and D is an increasing set, $\pi(D) \leq \hat{\pi}(D)$ for all $\pi \in B^c$, $\hat{\pi} \in B$ (since $\pi \leq_{st} \hat{\pi}$). Then, integrating over $\hat{\pi} \in B$ on both sides, $\int_B \pi(D)d\eta(\hat{\pi}) = \pi(D)\int_B d\eta(\hat{\pi}) \leq \int_B \hat{\pi}(D)d\eta(\hat{\pi})$; so $\pi(D)\eta(B) \leq \int_B \hat{\pi}(D)d\eta(\hat{\pi})$ for any $\pi \in B^c$. Similarly, $\eta(B)\int_{B^c} \pi(D)d\eta(\pi) \leq \eta(B^c)\int_B \hat{\pi}(D)d\eta(\hat{\pi})$.

Now, $\mu(E) = \int_{\Pi} \pi(E)d\eta(\pi) > 0$ implies that $\int_{\Pi} \pi(E)d\eta(\pi) = \int_{\Pi} \pi(D)d\eta(\pi)$, or we would not get $\int_{\Pi} \pi([\inf S_{-i}, \sup supp\ \mu])d\eta(\pi) = \mu([\inf S_{-i}, \sup supp\ \mu]) = 1$. So, $\eta(B)\int_{B^c} \pi(E)d\eta(\pi) \leq \eta(B^c)\int_B \hat{\pi}(E)d\eta(\hat{\pi})$. Then,

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

which implies that $\eta(B)\eta|_E(B^c) \leq \eta(B^c)\eta|_E(B)$. But $\eta|_E(B) + \eta|_E(B^c) = \eta(B) + \eta(B^c) = 1$, so $\eta(B) \leq \eta|_E(B)$. The increasing event B is arbitrary, so $\eta(B) \leq_{st} \eta|_E(B)$.

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

The requirement that Π_i is totally ordered does not imply that the resulting beliefs are totally ordered, only that all priors are ranked according to the "aggressiveness" of the potential strategies.

The problem with the strong monotonicity condition "any larger play produce larger beliefs" (Hopenhayn and Prescott 1992) is that it will not hold under Bayesian updating unless priors are ordered according to monotone likelihood ratio (Milgrom 1981).

7 Proofs

The proof of Theorem 1 requires two preliminary lemmas. Lemma 2 is of some interest independent of Theorem 1. Both lemmas are used in the proof of Theorem 2 as well.

Lemma 1 *Let X be a lattice and $A \subset X$ a subcomplete sublattice. Let $p \in \mathcal{P}(X)$ be properly mixed, and such that supp $p \subset A$. If $\lambda \in (0, 1)$ then $p <_{st} (1-\lambda)p + \lambda\delta_{\sup A} <_{st} \delta_{\sup A}$.*

Proof: Let E be an increasing subset of X. Since A is a subcomplete sublattice, $\sup A \in A$. *Case 1.* If $E \cap A = \emptyset$, then $p(E) = \delta_{\sup A}(E) = 0$, so $[(1-\lambda)p + \lambda\delta_{\sup A}](E) = 0$. *Case 2.* If there is $x \in E \cap A$, then $x \leq \sup A$, so $\sup A \in E$, as E is increasing. Then $\delta_{\sup A}(E) = 1$, which implies that

$$p(E) = (1-\lambda)p(E) + \lambda p(E) \leq (1-\lambda)p(E) + \lambda$$
$$= (1-\lambda)p(E) + \lambda\delta_{\sup A}(E) \leq 1 = \delta_{\sup A}(E).$$

In either case, then, $p(E) \leq [(1-\lambda)p + \lambda\delta_{\sup A}](E) \leq \delta_{\sup A}(E)$ for every increasing set E, so $p \leq_{st} (1-\lambda)p + \lambda\delta_{\sup A} \leq_{st} \delta_{\sup A}$.

Now we show that the inequalities are strict. Let $\tilde{E} = \{x \in X : \sup A \leq x\}$. Observe that $\sup A \in \tilde{E}$. Since p is properly mixed, $p(\tilde{E}) = p\left[\tilde{E} \cap \text{supp}(p)\right] \leq p(\{\sup A\}) < 1$ (see Remark 1 in Section 2). Then,

$$p(\tilde{E}) = (1-\lambda)p(\tilde{E}) + \lambda p(\tilde{E}) < (1-\lambda)p(\tilde{E}) + \lambda < 1 = \delta_{\sup A}(\tilde{E}).$$

So, $p <_{st} (1-\lambda)p + \lambda\delta_{\sup A} <_{st} \delta_{\sup A}$. ∎

Lemma 2 *Let $\Gamma = (N, \{(S_i, u_i) : i \in N\})$ be a GSC. Suppose σ is a PMNE, and that*

$$\sup \beta(\psi\sigma) \leq \inf \beta((1 - \epsilon)\psi\sigma + \epsilon\psi\delta_{\sup \beta(\psi\sigma)}) \tag{II}$$

for some $\epsilon > 0$; then σ is unstable to an ϵ-perturbation.

Proof: Let $\mu^0 = (1 - \epsilon)\psi\sigma + \epsilon\psi\delta_{\sup \beta(\psi\sigma)}$. Since Γ is a GSC, β is non-empty-, subcomplete- and sublattice-valued. Also, the support of σ is contained in $\beta(\psi\sigma)$. Then, by Lemma 1, $\psi\sigma <_{st} \mu^0 <_{st} \psi\delta_{\sup \beta(\psi\sigma)}$.

17

Let (μ, ξ) be a system of myopic behavior and monotone beliefs, with initial beliefs μ^0. Fix a sequence of realizations of type profiles $(\omega^1, \omega^2, \ldots) = \omega^\infty \in \Omega^\infty$. We will show by induction that, if the sequence $\{s^t\}$ with $s^t = \xi(\omega^t, \mu(h^{t-1}))$ is the realized play, then $\sup \beta(\psi\sigma) \leq s^t$ and $\mu^0 \leq_{st} \mu^t$ for every t.

First, we will show that $\sup \beta(\psi\sigma) \leq s^1$ and that $\mu^0 \leq_{st} \mu^1 = \mu(h^1)$. By hypothesis, $\sup \beta(\psi\sigma) \leq \tilde{s}$ for all $\tilde{s} \in \beta(\mu^0)$. Then, $\sup \beta(\psi\sigma) \leq s^1$, as behavior is myopic, so $\sup \mathrm{supp}\, \sigma \leq s^1$. Since $h^1 = (\emptyset, s^1)$, this implies, by monotonicity of beliefs, that $\mu^0 \leq_{st} \mu^1 = \mu(h^1)$, completing the first step of the proof by induction.

Second, suppose that for a given $t \geq 1$, $\mu^0 \leq_{st} \mu^{t-1} = \mu(h^{t-1})$, and $\sup \beta(\psi\sigma) \leq s^\tau$ for $\tau = 1, 2, \ldots t-1$. Γ is a GSC, so the map $\psi\sigma \mapsto \inf \beta(\psi\sigma)$ is monotone increasing (see e.g. Topkis (1998)). Then, $\psi\sigma <_{st} \mu^0 \leq_{st} \mu^t$ implies that

$$\sup \beta(\psi\sigma) \leq \inf \beta(\mu^0) \leq \inf \beta(\mu^{t-1}) \leq \xi(\omega^t, \mu^{t-1}) = s^t.$$

This inequality, and the inductive hypothesis, imply that $\mu^0 \leq_{st} \mu^{t+1} = \mu(h^t)$, as beliefs are weakly monotone. By induction then, for every t, both μ^t and the copy in belief space of the distribution of $\omega \mapsto \xi(\omega, h^t)$ are larger than μ^0.

Let $W = [\mu^0, \psi\delta_{\sup S}]^c = \{p \in \Psi : \mu^0 \leq_{st} p\}^c$. By Remark 2 of Section 2, W is a weak open neighborhood of $\psi\sigma$ in Ψ. We have shown that W satisfies the condition in the definition of unstable to an ϵ-perturbation. ∎

Remark: If a strategy profile σ is unstable to an ϵ-perturbation for all $\epsilon > 0$, then it is unstable. To see this, note that that $(1 - \epsilon)\psi\sigma + \epsilon\psi\sigma' \to \psi\sigma$ as $\epsilon \to 0$. [8] Then, given a neighborhood V of $\psi\sigma$, there is $\epsilon \in (0, 1)$ such that $\mu^0 = (1-\epsilon)\psi\sigma + \epsilon\psi\sigma' \in V$. Hence, for any neighborhood there is a point μ^0 such that, if the dynamics start at μ^0, then the state of the system never lies in a neighborhood W of $\psi\sigma$. Therefore, learning never approaches $\psi\sigma$. [9]

Proof of Theorem 1. Let $\epsilon > 0$, and consider any PMNE σ in the conditions of the Theorem. We shall prove that the inequality II in Lemma 2 is satisfied.

Fix $i \in N$ and let $\hat{\mu}_i = (1 - \epsilon)\psi_i\sigma + \epsilon\psi_i\delta_{[\sup \beta(\psi\sigma)]}$. Since at least two players select a properly mixed strategy, $\psi_i\sigma$ is properly mixed. Γ is a GSC, so $\beta_i(\hat{\mu}_i)$ is a

---

[8]Since, for any bounded, continuous, real-valued g, $(1 - \epsilon) \int g d\sigma + \epsilon \int g d\sigma' \to \int g d\sigma$.

[9]In fact, in our results the copy of the distribution of $\omega \mapsto \xi(\omega, \mu(h^t))$ in belief space is not in W either.

subcomplete sublattice, and supp $(\psi_i\sigma) \subset [\beta(\psi\sigma)]_{-i}$ because σ is a Nash equilibrium. Then, Lemma 1 implies that $\psi_i\sigma <_{st} \hat{\mu}_i$.

Let $T = \{\psi_i\sigma, \hat{\mu}_i\}$, and $\mathcal{U}_i : S_i \times T \to \mathbf{R}$ be defined by $\mathcal{U}_i(s_i, \tau) = \int_{S_{-i}} u_i(s_i, s_{-i}) d\tau(s_{-i})$. Hence, $\beta_i(\tau) = \mathrm{argmax}_{s_i \in S_i} \mathcal{U}_i(s_i, \tau)$ for $\tau \in T$.

We claim that $\mathcal{U}_i$ satisfies the strict single crossing property in (s_i, τ). Let $s_i < s_i'$ and suppose $\mathcal{U}_i(s_i', \psi_i\sigma) - \mathcal{U}_i(s_i, \psi_i\sigma) \geq 0$. Since $\psi_i\sigma <_{st} \hat{\mu}_i$, to show that the strict single-crossing property holds, we must show that $\mathcal{U}_i(s_i', \hat{\mu}_i) - \mathcal{U}_i(s_i, \hat{\mu}_i) > 0$. Denote by $l : S_{-i} \to \mathbf{R}$ the function $s_{-i} \mapsto [u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})]$. So, for any $\tau \in T$,

$$\mathcal{U}_i(s_i', \tau) - \mathcal{U}_i(s_i, \tau) = \int_{S_{-i}} [u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] d\tau(s_{-i}) = \int_{S_{-i}} l(s_{-i}) d\tau(s_{-i}).$$

Now, $\mathcal{U}_i(s_i', \hat{\mu}_i) - \mathcal{U}_i(s_i, \hat{\mu}_i) =$

$$(1 - \epsilon) \int_{S_{-i}} l(s_{-i}) d\sigma_{-i}(s_{-i}) + \epsilon \int_{S_{-i}} l(s_{-i}) d\delta_{\sup \beta(\psi\sigma)_{-i}}(s_{-i})$$
$$= (1 - \epsilon) [\mathcal{U}_i(s_i', \psi_i\sigma) - \mathcal{U}_i(s_i, \psi_i\sigma)] + \epsilon l(\sup \beta(\psi\sigma)_{-i}).$$

Suppose, by way of contradiction, that $\mathcal{U}_i(s_i', \hat{\mu}_i) - \mathcal{U}_i(s_i, \hat{\mu}_i) \leq 0$. Then $\mathcal{U}_i(s_i', \psi_i\sigma) - \mathcal{U}_i(s_i, \psi_i\sigma) \geq 0$ implies that $l(\sup \beta(\psi\sigma)_{-i}) \leq 0$.

Note that l is strictly increasing because Γ is a GSSC. Then $l(s_{-i}) < 0$ for all $s_{-i} \in \beta(\psi\sigma)_{-i} \backslash \{\sup \beta(\psi\sigma)_{-i}\}$, as $l(\sup \beta(\psi\sigma)_{-i}) \leq 0$. We claim that

$$\sigma_{-i}(\beta(\psi\sigma)_{-i} \backslash \{\sup \beta(\psi\sigma)_{-i}\}) = 0.$$

To see this, note that $\sigma_{-i}(\beta(\psi\sigma)_{-i} \backslash \{\sup \beta(\psi\sigma)_{-i}\}) > 0$ and $l(s_{-i}) < 0$ for all $s_{-i} \in \beta(\psi\sigma)_{-i} \backslash \{\sup \beta(\psi\sigma)_{-i}\}$ imply that

$$\begin{aligned} 0 &> \int_{\beta(\psi\sigma)_{-i} \backslash \{\sup \beta(\psi\sigma)_{-i}\}} l(s_{-i}) d\sigma_{-i}(s_{-i}) \\ &\geq \int_{S_{-i}} l(s_{-i}) d\sigma_{-i}(s_{-i}) \\ &= \mathcal{U}_i(s_i', \psi_i\sigma) - \mathcal{U}_i(s_i, \psi_i\sigma). \end{aligned}$$

The first inequality follows from a simple result in integration theory, [10] the second because $l(\sup \beta(\psi\sigma)_{-i}) \leq 0$ and supp $\sigma_{-i} \subset \beta(\psi\sigma)_{-i}$. But this violates that $\mathcal{U}_i(s_i', \psi_i\sigma) - \mathcal{U}_i(s_i, \psi_i\sigma) \geq 0$, so it must be that $\sigma_{-i}(\beta(\psi\sigma)_{-i} \backslash \{\sup \beta(\psi\sigma)_{-i}\}) = 0$.

---

[10] The integral of a strictly negative function over a region of positive measure is strictly negative.

But supp $\sigma_{-i} \subset \beta(\psi\sigma)_{-i}$, as σ is a Nash equilibrium. Then $\sigma_{-i}(\sup \beta(\psi\sigma)_{-i}) = 1$, which is a contradiction because σ is properly mixed (see Remark 1 in Section 2). This shows that $\mathcal{U}_i$ satisfies the strict single-crossing property.

The strict single crossing of $\mathcal{U}_i$ implies, by Milgrom and Shannon's (1994) Monotone Selection Theorem, that $x \leq x'$ for every $x \in \beta_i(\psi\sigma)$ and $x' \in \beta_i(\hat{\mu})$. Thus $\sup \beta_i(\psi\sigma) \leq \inf \beta_i(\hat{\mu})$. This is true for all i, hence the inequality II in Lemma 2 is satisfied. ∎

Proof of Theorem 2. In Step 1 we prove a mini-lemma, which we then use in Step 2 to prove the theorem.

STEP 1. Let $(\mu_i^\eta)_{\eta>0}$ be any collection of beliefs in Ψ_i such that $\mu_i^\eta \to \mu_i$, for some $\mu_i \in \Psi_i$, as $\eta \to 0$. We shall first show that there is $\hat{\eta}$ with the property that, for all $\eta < \hat{\eta}$, if $\int_i$ is a best response to μ_i^η in the η-augmented game, then $\int_i(\omega_i) \in \beta_i(\mu_i)$ for all ω_i; where $\beta_i(\mu_i) \subset S_i$ is the set of best responses to μ_i in Γ^0.

Let

$$k = \inf \left\{ \int_{S_{-i}} g_i(\tilde{s}_i, s_{-i}) d\mu_i(s_{-i}) - \int_{S_{-i}} g_i(s_i, s_{-i}) d\mu_i(s_{-i}) : s_i \notin \beta_i(\mu_i), \tilde{s}_i \in \beta_i(\mu_i) \right\}.$$

Since S_i is finite, $k > 0$. Let $0 < \eta' < k/4$ and let $\eta'' > 0$ be such that, if $0 < \eta < \eta''$ then $|\int_{S_{-i}} g_i(s_i, s_{-i}) d\mu_i(s_{-i}) - \int_{S_{-i}} g_i(s_i, s_{-i}) d\mu_i^\eta(s_{-i})| < k/4$ for all $s_i \in S_i$, which again is possible because S_i is finite. Let $\hat{\eta} = \min\{\eta', \eta''\}$.

Let $s_i \notin \beta_i(\mu_i)$, $\tilde{s}_i \in \beta_i(\mu_i)$, and $\eta < \hat{\eta}$. The gain to playing $\tilde{s}_i$ over s_i in the η-augmented game, and in state ω_i, is $\int_{S_{-i}} u_i^\eta(\tilde{s}_i, s_{-i}, \omega_i) d\mu_i^\eta(s_{-i}) - \int_{S_{-i}} u_i^\eta(s_i, s_{-i}, \omega_i) d\mu_i^\eta(s_{-i}) =$

$$
\begin{aligned}
& \int_{S_{-i}} g_i(\tilde{s}_i, s_{-i}) d\mu_i^\eta(s_{-i}) - \int_{S_{-i}} g_i(s_i, s_{-i}) d\mu_i^\eta(s_{-i}) + \eta[\omega_i(\tilde{s}_i) - \omega_i(s_i)] \\
\geq \ & \int_{S_{-i}} g_i(\tilde{s}_i, s_{-i}) d\mu_i^\eta(s_{-i}) - \int_{S_{-i}} g_i(\tilde{s}_i, s_{-i}) d\mu_i(s_{-i}) \\
& + \int_{S_{-i}} g_i(\tilde{s}_i, s_{-i}) d\mu_i(s_{-i}) - \int_{S_{-i}} g_i(s_i, s_{-i}) d\mu_i(s_{-i}) \\
& + \int_{S_{-i}} g_i(s_i, s_{-i}) d\mu_i(s_{-i}) - \int_{S_{-i}} g_i(s_i, s_{-i}) d\mu_i^\eta(s_{-i}) - 2\eta \\
> \ & -k/4 + k - k/4 - 2k/4 = 0.
\end{aligned}
$$

The first inequality obtains because $\omega_i(s_i), \omega_i(\tilde{s}_i) \in [-1, 1]$, the second because $\eta < \hat{\eta}$.

We have shown that, for all $\eta < \hat{\eta}$, no matter the value of ω_i, any $\hat{s}_i \in \beta_i(\mu_i)$ is a better response to μ_i^η than any $s_i \notin \beta_i(\mu_i)$. As $\int_i^\eta$ is a best response to beliefs μ_i^η, this implies that $\int_i^\eta(\omega_i) \in \beta_i(\mu_i)$ for all ω_i.

STEP 2. For each η, let σ_i^η be the distribution of $\int_i^\eta(\omega_i)$. Note that, for all η, $\int_i^\eta$ is a best response to beliefs $\psi_i\sigma^\eta$, and $\sigma = \lim_{\eta\to 0}\sigma^\eta$, as $(\Gamma^\eta, \int^\eta)$ is a purification sequence of σ. By Step 1, there is $\hat\eta$ such that, if $\eta < \hat\eta$, then $\int_i^\eta(\omega_i) \in \beta_i(\psi_i\sigma)$ for all $\omega_i \in \Omega_i$. In particular, $\int_i^\eta(\omega_i) \leq \sup \beta_i(\psi_i\sigma)$ for all $\omega_i \in \Omega_i$.

Let $\tilde{\int}_i^\eta$ be a best response to beliefs $\nu = (1 - \epsilon)\psi_i\sigma + \epsilon\psi_i\delta_{\sup \beta(\psi\sigma)}$, in the η-augmented game. Beliefs ν do not depend on η, so applying Step 1 with $\mu^\eta = \nu$ for all η, there is $\tilde\eta$ such that, if $\eta < \tilde\eta$ then $\tilde{\int}_i^\eta(\omega_i) \in \beta_i(\nu)$ for all ω_i. In particular, $\inf \beta_i(\nu) \leq \tilde{\int}_i^\eta(\omega_i)$ for all ω_i.

Let $\underline{\eta} = \min\{\hat\eta, \tilde\eta\}$. Repeating the argument in the proof of Theorem 1, we have that $\sup \beta(\psi\sigma) \leq \inf \beta(\nu)$. Then, if $\eta < \underline{\eta}$, for every i and $\omega_i, \tilde\omega_i \in \Omega_i$, $\int_i^\eta(\omega_i) \leq \tilde{\int}_i^\eta(\tilde\omega_i)$. By Lemma 2, we are done. ∎

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 (2000): "Learning Purified Mixed Equilibria," *Journal of Economic Theory*, 90(1), 84–115.

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.

FUDENBERG, D., AND J. TIROLE (1991): *Game Theory.* MIT Press.

HIRSCH, M. W., AND S. SMALE (1974): *Differential Equations, Dynamical Systems, and Linear Algebra.* Academic Press.

HOPENHAYN, H. A., AND E. C. PRESCOTT (1992): "Stochastic Monotonicity and Stationary Distributions for Dynamic Economies," *Econometrica*, 60(6), 1387–1406.

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. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12(2), 380–391.

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.

OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT Press.

ROYDEN, H. L. (1988): *Real Analysis*. Prentice Hall.

TIROLE, J. (1988): *The Theory of Industrial Organization*. MIT Press.

TOPKIS, D. M. (1998): *Supermodularity and Complementarity*. Princeton University Press.

VIVES, X. (1999): *Oligopoly Pricing*. MIT Press.