

Evolution in Bayesian Games I: Theory^{*}

Jeffrey C. Ely
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
ely@bu.edu

William H. Sandholm
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706
whs@ssc.wisc.edu
www.ssc.wisc.edu/~whs

June 18, 2004

* We thank Drew Fudenberg, Josef Hofbauer, Larry Samuelson, and Jörgen Weibull, two anonymous referees and an anonymous Associate Editor, and seminar audiences at the Stockholm School of Economics and at Wisconsin for their comments. Results in this paper were previously circulated under the title "Evolution with Diverse Preferences". Financial support from NSF Grants SBR-9810787 and SES-0092145 is gratefully acknowledged.

Abstract

We introduce best response dynamics for settings where agents' preferences are diverse. Under these dynamics, which are defined on the space of Bayesian strategies, rest points and Bayesian Nash equilibria are identical. We prove the existence and uniqueness of solution trajectories to these dynamics, and provide methods of analyzing the dynamics based on aggregation.

JEL Classification: C72, C73

1. Introduction

We study best response dynamics for populations with diverse preferences. The state variables for these dynamics are Bayesian strategies: that is, maps from preferences to distributions over actions. We prove the existence, uniqueness, and continuity of solutions of these dynamics, and show that the rest points of the dynamics are the Bayesian equilibria of the underlying game. We then characterize the dynamic stability of Bayesian equilibria in terms of aggregate dynamics defined on the simplex, making it possible to evaluate stability using standard dynamical systems techniques.

We offer three motivations for this study. First, we feel that in interactions involving large populations, different individuals are unlikely to evaluate payoffs in precisely the same way. Therefore, in constructing evolutionary models, it seems realistic to explicitly allow for diversity in preferences. We shall see that doing so eliminates pathological solution trajectories that can arise under best response dynamics when preferences are common.

A second motivation for our study is to provide foundations for models of preference evolution.¹ In these models, natural selection of preferences is mediated through behavior, as the preferences that survive are those that induce the fittest behavior. Ideally, models of preference evolution should be built up from models of behavior adjustment defined for settings where preferences are diverse but fixed. By providing tools for analyzing behavior under diverse preferences, this paper provides the groundwork for studying competition among the preferences themselves.

Our third and most important motivation is to provide methods for the evolutionary analysis of Bayesian games. Nearly all work in evolutionary game theory has considered games with complete information. At the same time, the proliferation of game theory in applied economic analysis is in large part due to its deft handling of informational asymmetries; in this development, games of incomplete information have played a leading role. In offering evolutionary techniques for studying Bayesian games, we hope that the insights of evolutionary game theory can be brought to bear more broadly in applied work.

We consider a population of agents facing a repeated strategic interaction. Unlike their counterparts in standard evolutionary models, different agents in our

¹ See, for example, Güth and Yaari (1992), Ely and Yilankaya (2001), and Sandholm (2001).

model evaluate payoffs using different payoff functions. We assume that the subpopulation of agents with any given payoff function is of negligible size relative to the population as a whole. A complete description of behavior is given by a Bayesian strategy: a map that specifies the distribution of actions played in each subpopulation. The appropriate notion of equilibrium behavior is Bayesian equilibrium, which requires that each subpopulation play a best response to the aggregate behavior of the population as a whole.

Our goal is to model the evolution of behavior in a diverse population in a plausible and tractable way. To do so, we build on the work of Gilboa and Matsui (1991), who introduced the *best response dynamic* for the common preference setting. Under their dynamic, the distribution of actions in a population always adjusts toward some best response to current behavior. To define our *Bayesian best response dynamic*, we require instead that the distribution of actions within each subpopulation adjust toward that subpopulation's current best response.

To complete the definition of the Bayesian dynamic, we must specify a notion of distance between Bayesian strategies.² We utilize the L^1 norm, which measures the distance between two Bayesian strategies as the average change in the subpopulations' behaviors. We establish that the law of motion of the Bayesian dynamic is Lipschitz continuous under this norm, enabling us to prove that solutions to the dynamic exist and are unique.

This uniqueness result is of particular interest because it fails to hold when preferences are common. Under common preferences, multiple solution trajectories to the best response dynamic can originate from a single initial condition. This property is the source of surprising solution trajectories: Hofbauer (1995) offers a game in which solutions to the best response dynamic cycle in and out of a Nash equilibrium in perpetuity. Our uniqueness result implies that even slight diversity in preferences renders such solution trajectories impossible.

Since our dynamic is defined on the (L^1) space of Bayesian strategies, it is difficult to analyze directly. To contend with this, we introduce an *aggregate best response dynamic* defined directly on the simplex. We show that there is a many-to-one mapping from solutions to the Bayesian dynamic to solutions to the aggregate dynamic; the relevant mapping is the one that converts Bayesian strategies to the aggregate behavior they induce. Thus, if we run the Bayesian dynamic from two Bayesian strategies whose aggregate behaviors are the same, the two solutions to the

² By doing so, we fix the interpretation of the differential equation that defines the dynamic—see Section 2.2.

Bayesian dynamic exhibit the same aggregate behavior at all subsequent times.

Were we only interested aggregate behavior, we could focus our attention entirely on the aggregate dynamic. But in most applications of Bayesian games, the full Bayesian strategy is itself of cardinal importance. For example, in a private value auction, the distribution of bids is on its own an inadequate description of play; to determine efficiency, one must also know which bidders are placing which bids. Knowing the entire Bayesian strategy is also critical in studying preference evolution: there we must know which preferences lead players to choose the fittest actions, as these are the preferences that will thrive under natural selection.

Since the full Bayesian strategy is of central interest, it is important to be able to determine which Bayesian equilibria are dynamically stable. To accomplish this, we establish a one-to-one correspondence between the equilibria that are stable under the Bayesian dynamic and the distributions that are stable under the aggregate dynamic. Using this result, one can determine which Bayesian equilibria are stable under the original L^1 dynamic by considering a much simpler dynamic defined on the simplex.³

Of course, this simpler dynamic is still a nonlinear differential equation, so it is not immediately clear whether these aggregation results are of practical importance. Fortunately, Hofbauer and Sandholm (2002, 2004) have established global convergence results for the aggregate best response dynamic in a number of interesting classes of games. In addition, a companion to the present paper (Sandholm (2003)) uses the aggregation results developed here to prove dynamic versions of Harsanyi's (1973) purification theorem.

Ellison and Fudenberg (2000) study fictitious play in a population with diverse preferences. In fictitious play, all players choose a best response to the time average of past play. Since this time average is the model's state variable, fictitious play defines a dynamic directly on the simplex even when preferences are diverse. In fact, it is easy to show that the dynamic studied by Ellison and Fudenberg (2000) are equivalent (after a time reparameterization) to our *aggregate* best response dynamic. The connections between these processes are considered in greater detail in the final section of the paper.

³ Were the mapping between solution trajectories one-to-one as well, the stability results would follow as an immediate consequence. However, since this mapping is actually many-to-one, these results are not obvious—see Section 6.

2. The Best Response Dynamic

A unit mass of agents recurrently plays a population game. Each agent chooses one of n actions, which we identify with basis vectors in \mathbf{R}^n : $S = \{e_1, e_2, \dots, e_n\}$.⁴ We let $\Delta = \{x \in \mathbf{R}_+^n: \sum_i x_i = 1\}$ denote the set of distributions over actions.

2.1 Common Preferences

In typical evolutionary models, all agents share the same preferences. Here, we represent these preferences by a Lipschitz continuous function $\pi: \Delta \rightarrow \mathbf{R}^n$; $\pi_i(x)$ represents the payoff to strategy i when aggregate behavior is x . An important special case is based on random matching in a symmetric normal form game with payoff matrix $A \in \mathbf{R}^{n \times n}$; in this case, the payoff function is the linear function $\pi(x) = Ax$. More generally, our setup also allows the payoffs to each action to depend nonlinearly on the population state, a feature that is essential in some applications—see Sandholm (2004).

Let $BR^\pi: \Delta \Rightarrow \Delta$ denote the best response correspondence for payoff function π :

$$BR^\pi(x) = \operatorname{argmax}_{y \in \Delta} y \cdot \pi(x)$$

Action distribution $x^* \in \Delta$ is a *Nash equilibrium* under π if $x^* \in BR^\pi(x^*)$: that is, if each agent chooses an action that is optimal given the behavior of the others.

The *best response dynamic* on Δ is defined by

$$(BR) \quad \dot{x} \in BR^\pi(x) - x.$$

The usual interpretation of this dynamic is that agents occasionally consider switching actions, choosing a best response whenever they do so. The $-x$ term arises because at each moment in time, all agents are equally likely to consider a switch.

Gilboa and Matsui (1991), Matsui (1992), and Hofbauer (1995) study the best response dynamic in the context of random matching in normal form games. For most payoff matrices, there are action distributions that admit multiple best responses, and hence many possible directions of motion under (BR); hence,

⁴ All results in this paper are easily extended to allow multiple player roles (i.e., to allow different subsets of the population to choose from different sets of actions).

solutions to (BR) need not be unique. For example, if the population begins at a Nash equilibrium x^* , agents who switch to best responses can do so in proportions x^* , resulting in a stationary solution trajectory at x^* . But if the agents who switch to a best response do so in proportions other than x^* , the population may move away from the equilibrium. This can lead to complicated solution trajectories: Hofbauer (1995) presents a game in which the population continually travels through cycles of varying lengths, passing through a Nash equilibrium at the start of each circuit.

We show that the existence of solutions that leave Nash equilibria is a consequence of the assumption that all agents' preferences are identical. The source of the nonuniqueness of solutions to (BR) is the fact that for most payoff matrices, there is a set of action distributions admitting multiple best responses. Indeed, Hofbauer's (1995) example is generic, in that all payoff matrices close to the one he considers yield qualitatively similar dynamics.

Our analysis shows that there is another sense in which Hofbauer's (1995) example is not generic. The analysis relies on the following observation: if we fix a distribution over actions, the set of payoff matrices that generate indifference at that distribution is negligible. Therefore, in a population with diverse preferences, best responses are “essentially unique”, and the function that defines the best response dynamic in this context is single valued. To establish the uniqueness of solutions, and thus the equivalence of rest points and Bayesian equilibria, we must establish that this function is not only single valued, but also Lipschitz continuous. We show below that this is true if distances between Bayesian strategies are measured in an appropriate way.

2.2 Diverse Preferences

To incorporate diverse preferences, we suppose that the distribution of payoff functions in the population is described by a probability measure μ on the set of payoff functions $\Pi = \{\pi: \Delta \rightarrow \mathbf{R}^n \mid \pi \text{ is Lipschitz continuous}\}$. In the language of Bayesian games, μ represents the distribution of types, which in the current context are simply the agents' payoff functions. The common preferences model corresponds to the case in which μ places all mass on a single point in Π . We rule out such cases below, focusing instead on settings with genuine diversity.

We suppose that there are a continuum of agents with each preference $\pi \in \Pi$ in the support of μ . Each agent chooses a pure action in S . The behavior of the subpopulation with preference π is described by a distribution in Δ . A *Bayesian*

strategy is a map $\sigma: \Pi \rightarrow \Delta$, where $\sigma(\pi)$ is the distribution over pure actions chosen in aggregate by the agents of type π . Each Bayesian strategy σ can be viewed as a random vector on the probability space (Π, μ) that takes values in Δ . The set $\Sigma = \{\sigma: \Pi \rightarrow \Delta\}$ contains all (Borel measurable) Bayesian strategies. We consider a pair of Bayesian strategies $\sigma, \rho \in \Sigma$ equivalent if $\sigma(\pi) = \rho(\pi)$ for μ -almost every π . In other words, we do not distinguish between Bayesian strategies that indicate the same action distribution for almost every type.

Let E denote expectation taken with respect to the probability measure μ . The proportion of agents who play action i under the Bayesian strategy σ is then given by $E\sigma_i = \int_{\Pi} \sigma_i(\pi) d\mu$, and the *aggregate behavior* induced by $\sigma \in \Sigma$ is $E\sigma \equiv (E\sigma_1, \dots, E\sigma_n) \in \Delta$. That is, the operator E takes both random variables and random vectors as arguments, handling each in the appropriate way. We sometimes call $E\sigma$ the *distribution* induced by σ . Our notion of distance between distributions is the summation norm on \mathbf{R}^n : for $x \in \mathbf{R}^n$, let

$$|x| = \sum_{i=1}^n |x_i|.$$

Each agent's best responses are defined with respect to current aggregate behavior $x = E\sigma \in \Delta$. We let $B: \Delta \Rightarrow \Sigma$ denote the best response correspondence, which we define by

$$B(x)(\pi) \equiv BR^\pi(x) = \arg \max_{y \in \Delta} y \cdot \pi(x).$$

The best response $B(x) \in \Sigma$ is a Bayesian strategy; for each $\pi \in \Pi$, $B(x)(\pi)$ is the set of distributions in Δ that are best responses against aggregate behavior x for agents with preference π .

We state some weak but useful conditions on the preference distribution μ in terms of the best response correspondence B ; classes of preference distributions that satisfy these conditions are introduced below. Condition (C1) requires that for all aggregate behaviors $x \in \Delta$, the set of agents with multiple best responses has measure zero.

(C1) B is single valued.

Under condition (C1), all selections from $B(x)$ are equivalent, allowing us to regard $B: \Delta \rightarrow \Sigma$ as a function rather than as a correspondence.

Each Bayesian strategy $\sigma \in \Sigma$ induces some distribution $E\sigma \in \Delta$; the best response to this distribution is $B(E(\sigma))$. We say that the Bayesian strategy σ^* is a *Bayesian equilibrium* if it is a best response to itself: that is, if $\sigma^* = B(E(\sigma^*))$. We let $\Sigma^* \subseteq \Sigma$ denote the set of Bayesian equilibria. Observe that under condition (C1), all aggregate behaviors induce a unique, *pure* best response: for all x , $\mu\{\pi: B(x)(\pi) \in \{e_1, \dots, e_n\}\} = 1$. Hence, all Bayesian equilibria must also be pure.⁵

The *Bayesian best response dynamic* is described by the law of motion

$$(B) \quad \dot{\sigma} = B(E(\sigma)) - \sigma$$

on Σ , the space of Bayesian strategies. The right hand side of this equation is a map from Σ to $\hat{\Sigma} = \{\sigma: \Pi \rightarrow \mathbf{R}^n\}$, a linear space containing all directions of motion through Σ .

To complete the definition of the dynamic, we must specify the norm used to measure distances between points in $\hat{\Sigma}$. To interpret equation (B) preference by preference, one would employ the L^∞ norm,

$$\|\sigma\|_{L^\infty} = \text{esssup}_{\pi \in \Pi} |\sigma(\pi)|.$$

This norm defines too strong a topology for our purposes. To see why, consider two strategy distributions $x, y \in \Delta$ that lie close to one another. As long as there is a non-null set of preferences whose best responses to x and y differ, the Bayesian best responses $B(x)$ and $B(y)$ are far apart in the L^∞ norm: $\|B(x) - B(y)\|_{L^\infty} = 2$. For this reason, the law of motion (B) is discontinuous in this norm, so standard methods of establishing the existence and uniqueness of solution trajectories fail.

To create a tractable model, we need to use a norm on $\hat{\Sigma}$ that makes it easier for two points to be close to one another, so that under this norm equation (B) defines a continuous law of motion. In particular, we want pairs of Bayesian strategies that only differ in the choices of agents whose preferences lie in some set of small measure to be regarded as close together. One appropriate choice of norm is the L^1 norm, which we denote by $\|\cdot\|$:

⁵ Of course, this observation is originally due to Harsanyi (1973).

$$\|\sigma\| \equiv \sum_{i=1}^n E|\sigma_i| = E\left(\sum_{i=1}^n |\sigma_i|\right) = E|\sigma|.$$

Under this norm, the distance between a pair of Bayesian strategies is determined by the average change in behavior over all subpopulations. Hence, if the best responses to x and y differ only for a set of preferences of measure ε , then these best responses are close in L^1 norm: $\|B(x) - B(y)\| = 2\varepsilon$.⁶

In order to establish existence and uniqueness of solution trajectories to the Bayesian best response dynamic, it is enough to know that the dynamic is Lipschitz continuous. The following lemma is a first step in this direction.

Lemma 2.1: $E: \Sigma \rightarrow \Delta$ is Lipschitz continuous (with Lipschitz constant 1).

Proof: Since E is linear, it is enough to show that $|E\sigma| \leq \|\sigma\|$. And indeed,

$$|E\sigma| = \sum_{i=1}^n |E\sigma_i| \leq \sum_{i=1}^n E|\sigma_i| = \|\sigma\|. \quad \blacksquare$$

Given Lemma 2.1, Lipschitz continuity of the dynamic is a consequence of the following condition.

(C2) B is Lipschitz continuous (with respect to the L^1 norm).

Condition (C2) asks that small changes in aggregate behavior x lead to correspondingly small changes in the best response $B(x)$, where the distance between best responses is measured using the L^1 norm.

Our two conditions on the function B will hold as long as the preference distribution μ is both sufficiently diverse and sufficiently smooth. We illustrate this using two examples. Our first example concerns random matching in normal form games. In this example, every agent's payoffs are derived from some payoff matrix $A \in \mathbf{R}^{n \times n}$, but different agents have different payoff matrices.

⁶ The choice of norm is also important issue in other models of evolution with infinite dimensional state variables. For example, in Oechssler and Riedel's (2001) work on replicator dynamics for games with infinite strategy spaces, the choice of norm determines the set of payoff functions for which the dynamic is well defined.

Proposition 2.2: Let λ be a probability measure on $\mathbf{R}^{n \times n}$, and define the preference distribution μ by $\mu\{\pi: \pi(x) = Ax \text{ for some } A \in M\} = \lambda(M)$. If λ admits a bounded density function with compact support, then B satisfies conditions (C1) and (C2).

For our second example, we suppose that all agents' preferences are based on the same (possibly nonlinear) payoff function, but that each agent has idiosyncratic preferences θ_i in favor of or against each action $i \in S$.

Proposition 2.3: Let $F \in \Pi$ be a Lipschitz continuous payoff function, let ν be a probability measure on \mathbf{R}^n , and define the preference distribution μ by $\mu\{\pi: \pi(x) = F(x) + \theta \text{ for some } \theta \in \Theta\} = \nu(\Theta)$. Suppose that ν admits a bounded density function $f: \mathbf{R}^n \rightarrow \mathbf{R}$ and that either (i) ν has compact support, (ii) ν is a product measure, or (iii) for each pair (i, j) , $i \neq j$, the density $g_{ji}: \mathbf{R} \rightarrow \mathbf{R}$ for the difference $\theta_j - \theta_i$ is bounded. Then B satisfies conditions (C1) and (C2).

3. Basic Properties

We now establish some basic properties of solutions to the Bayesian best response dynamic (B). Since we will interpret equation (B) in the L^1 sense, we begin by reviewing the notions of continuity and differentiability for trajectories through the L^1 space $(\hat{\Sigma}, \|\cdot\|)$; see Lang (1997) for additional details.

Let $\{\sigma_t\} = \{\sigma_t\}_{t \geq 0}$ be a trajectory through $\hat{\Sigma}$. We say that $\bar{\sigma} \in \hat{\Sigma}$ is the L^1 limit of σ_s as s approaches t , denoted $\bar{\sigma} = L^1 \lim_{s \rightarrow t} \sigma_s$, if

$$\lim_{s \rightarrow t} \|\sigma_s - \bar{\sigma}\| = \lim_{s \rightarrow t} E|\sigma_s - \bar{\sigma}| = 0.$$

The trajectory $\{\sigma_t\}$ is L^1 continuous if $\sigma_t = L^1 \lim_{s \rightarrow t} \sigma_s$ for all t . If there exists a $\dot{\sigma}_t \in \hat{\Sigma}$ such that

$$\dot{\sigma}_t = L^1 \lim_{\varepsilon \rightarrow 0} \left(\frac{\sigma_{t+\varepsilon} - \sigma_t}{\varepsilon} \right),$$

we call $\dot{\sigma}_t$ the L^1 derivative of trajectory $\{\sigma_t\}$ at time t .

As usual, the L^1 derivative $\dot{\sigma}_t$ describes the direction of motion of the trajectory $\{\sigma_t\} \subset \hat{\Sigma}$ at time t . But even when this derivative exists, the (standard) derivative

$\frac{d}{dt}(\sigma_t(\pi))$ of the distribution trajectory $\{\sigma_t(\pi)\} \subset \mathbf{R}^n$ of any particular preference π need not exist: the slope $\frac{1}{\varepsilon}(\sigma_{t+\varepsilon}(\pi) - \sigma_t(\pi)) \in \mathbf{R}^n$ of the line segment from $(t, \sigma_t(\pi))$ to $(t + \varepsilon, \sigma_{t+\varepsilon}(\pi))$ may not converge as ε approaches zero. For the L^1 derivative to exist, the measure of the set of preferences π for which this slope is not close to $\dot{\sigma}_t(\pi) \in \mathbf{R}^n$ must become arbitrarily small as ε vanishes.

A Lipschitz continuous function $f: \hat{\Sigma} \rightarrow \hat{\Sigma}$ defines a law of motion

$$(D) \quad \dot{\sigma} = f(\sigma)$$

on $\hat{\Sigma}$. A trajectory $\sigma: \mathbf{R}_+ \rightarrow \hat{\Sigma}$ is an L^1 solution to equation (D) if $\dot{\sigma}_t = f(\sigma_t)$ μ -almost surely for all t , where $\dot{\sigma}_t$ is interpreted as an L^1 derivative.⁷

Theorem 3.1 sets out the basic properties of solutions of the Bayesian dynamic. Its proof is provided in the Appendix.

Theorem 3.1: (*Basic properties of solutions to (B)*)

(i) *There exists an L^1 solution to (B) starting from each $\sigma_0 \in \Sigma$. This solution is unique in the L^1 sense: if $\{\sigma_t\}$ and $\{\rho_t\}$ are L^1 solutions to (B) such that $\rho_0 = \sigma_0$ μ -a.s., then $\rho_t = \sigma_t$ μ -a.s. for all t .*

(ii) *If $\{\sigma_t\}$ and $\{\rho_t\}$ are L^1 solutions to (B), then*

$$\|\sigma_t - \rho_t\| \leq \|\sigma_0 - \rho_0\| e^{Kt},$$

where K is the Lipschitz constant of $f(\sigma) = B(E(\sigma)) - \sigma$.

(iii) *Solutions to (B) remain in Σ at all times $t \in [0, \infty)$.*

(iv) *From each $\sigma_0 \in \Sigma$ there is an L^1 solution to (B) with the property that*

$$\mu(\pi: \sigma_t(\pi) \text{ is continuous in } t) = 1.$$

(v) *σ^* is a rest point of (B) if and only if it is a Bayesian equilibrium.*

Part (i) guarantees the existence and uniqueness of solutions to (B), while parts (ii) and (iii) establish continuity in initial conditions and forward invariance of Σ . Since (B) is Lipschitz, these results are nearly standard; the main technicality that must be addressed is the fact that the domain Σ of the dynamic is closed.

⁷ The definition of an L^1 solution requires that the derivative $\dot{\sigma}_t$ exist at all times $t \geq 0$. In contrast, since the standard best response dynamic (BR) has a discontinuous law of motion, to ensure that solutions to (BR) exist one must allow differentiability to fail at a zero measure set of times.

If $\{\sigma_t\}$ is an L^1 solution to (B), then so is any trajectory $\{\hat{\sigma}_t\}$ that differs from $\{\sigma_t\}$ on some measure zero set $\Pi_t \subset \Pi$ at each instant t . Thus, while part (i) of the theorem guarantees the existence of a unique L^1 solution to (B), this result imposes no restrictions on the distribution trajectory $\{\sigma_t(\pi)\}$ of an individual preference π : as time passes, it is possible for the behavior of the subpopulation with preference π to jump haphazardly about the simplex. Fortunately, part (iv) of the theorem shows that we can always find an L^1 solution with the property that the behavior associated with almost every preference changes continuously over time. Finally, part (v) of the theorem observes that the rest points of (B) are precisely the Bayesian equilibria of the underlying game.

4. Aggregation and Equilibrium

We have established that solution trajectories of the best response dynamic (B) exist and are unique. However, since this dynamic operates on an L^1 space, working with it directly is rather difficult. In the coming sections, we show that many important properties of the dynamic can be understood by analyzing an aggregate dynamic. The aggregate dynamic is defined on the simplex, and so can be studied using standard methods.

Before introducing the aggregate dynamic, we reconsider the Bayesian equilibria $\sigma^* \in \Sigma^*$, which are the rest points of (B). Since the Bayesian strategy σ induces the distribution $E(\sigma) \in \Delta$, Bayesian equilibria satisfy $\sigma^* = B(E(\sigma^*))$.

If the current distribution is $x \in \Delta$, the Bayesian strategy that is a best response to this distribution is $B(x)$, which in turn induces the distribution $E(B(x))$. We therefore call $x^* \in \Delta$ an *equilibrium distribution* if $x^* = E(B(x^*))$, and let $\Delta^* \subseteq \Delta$ denote the set of equilibrium distributions.

The connection between Bayesian equilibria and equilibrium distributions is established in the following result.

Theorem 4.1: (*Characterization of equilibria*)

The map $E: \Sigma^ \rightarrow \Delta^*$ is a homeomorphism whose inverse is $B: \Delta^* \rightarrow \Sigma^*$.*

Proof: First, we show that E maps Σ^* into Δ^* . Let $\sigma \in \Sigma^*$ be a Bayesian equilibrium: $\sigma = B(E(\sigma))$. Then $E(\sigma) = E(B(E(\sigma)))$, so $E(\sigma) \in \Delta^*$.

Second, we show that E is onto. Fix a distribution $x \in \Delta^*$, so that $x = E(B(x))$; we need to show that there is a Bayesian strategy $\sigma \in \Sigma^*$ such that $E(\sigma) = x$. Let $\sigma = B(x)$.

Then since $x \in \Delta^*$, $E(\sigma) = E(B(x)) = x$. Furthermore, this equality implies that $B(E(\sigma)) = B(x) = \sigma$, so $\sigma \in \Sigma^*$. Thus, E is onto, and $B(x) \in E^{-1}(x)$.

Third, we show that E is one-to-one, which implies that $B(x) = E^{-1}(x)$. Fix two Bayesian equilibria $\sigma, \sigma' \in \Sigma^*$, and suppose that $E(\sigma) = E(\sigma')$. Then $\sigma = B(E(\sigma)) = B(E(\sigma')) = \sigma'$.

Finally, the continuity of E and B follows from Lemma 2.1 and condition (C2). ■

The space Σ of Bayesian strategies is considerably more complicated than the space of distributions Δ . Nevertheless, Theorem 4.1 shows that if we are only concerned with Bayesian equilibria $\sigma^* \in \Sigma^*$, it is sufficient to consider the equilibrium distributions $x^* \in \Delta^*$. We can move between the two representations of equilibria using the maps E and B , whose restrictions to the equilibrium sets are inverses of one another.

If we are concerned with disequilibrium behavior, then the one-to-one link between Bayesian strategies and distributions no longer exists: E maps many Bayesian strategies to the same distribution over actions, and if the Bayesian strategy σ is not an equilibrium, B does not invert E : that is, $B(E(\sigma)) \neq \sigma$.

Fortunately, we are able to prove analogues of Theorem 4.1 for solutions to the Bayesian best response dynamic (B). To do so, we introduce the aggregate best response dynamic (AB), which is defined on the simplex. In the next section, we show that the expectation operator E is a many-to-one map from solutions to (B) to solutions to (AB). In Section 6, we establish a one-to-one correspondence between stable rest points of (B) and stable rest points of (AB). Therefore, while the Bayesian dynamic operates on the complicated space Σ , the answers to many important questions about this dynamic can be obtained by applying standard tools to dynamics on the simplex.

5. Aggregation of Solution Trajectories

Under the dynamic (B), the Bayesian strategy σ_t always moves toward its best response $B(E(\sigma_t))$. Hence, the target point only depends on σ_t through its distribution $E(\sigma_t)$. This "bottleneck" provides the basis for our aggregation results.

We define the *aggregate best response dynamic* by

$$(AB) \quad \dot{x}_t = E(B(x_t)) - x_t.$$

Under this law of motion, the distribution x_t moves toward the distribution induced by the best response to x_t . Lemma 2.1 and condition (C2) imply that this dynamic is Lipschitz continuous. Therefore, solutions to (AB) exist, are unique, are Lipschitz continuous in their initial conditions, and leave Δ forward invariant (see Theorem A.1 in the Appendix). Moreover, the rest points of (AB) are easily characterized.

Observation 5.1: *The set of rest points of (AB) is Δ^* , the set of equilibrium distributions.*

Let $f: \Sigma \rightarrow \hat{\Sigma}$ and $g: \Delta \rightarrow \mathbf{R}^n$ be Lipschitz continuous functions, and consider the following laws of motion on Σ and Δ .

$$(D) \quad \dot{\sigma} = f(\sigma);$$

$$(AD) \quad \dot{x} = g(x).$$

We say that the dynamic (D) *aggregates* to the dynamic (AD) if whenever $\{\sigma_t\}$ is an L^1 solution to (D), $\{E\sigma_t\}$ is a solution to (AD).

Theorem 5.2: *(Aggregation of solution trajectories)*

The Bayesian best response dynamic (B) aggregates to the aggregate best response dynamic (AB).

Theorem 5.2 tells us that the dynamic (AB) completely describes the evolution of aggregate behavior under the dynamic (B). If $\{\sigma_t\}$ is a solution to (B), then the distribution it induces at time t , $E\sigma_t$, is equal to x_t , where $\{x_t\}$ is the solution to (AB) starting from $x_0 = E\sigma_0$. Since aggregate behavior at time t under (B) is fully determined by aggregate behavior at time 0, Bayesian strategies that induce the same aggregate behavior also induce the same aggregate behavior trajectories.

It is important to note that this mapping between solution trajectories is many-to-one. For example, consider a solution $\{\sigma_t\}$ to (B) whose initial Bayesian strategy aggregates to an equilibrium distribution: $E\sigma_0 = x^* \in \Delta^*$. Observation 5.1 and Theorem 5.2 imply that the distribution trajectory $\{E\sigma_t\}$ induced by $\{\sigma_t\}$ is degenerate: $E\sigma_t = x^*$ for all t . However, $\{\sigma_t\}$ is itself degenerate only if σ_0 is a Bayesian equilibrium; there are many Bayesian strategies $\sigma \in E^{-1}(x^*)$ that aggregate to x^* but are not Bayesian equilibria, and hence are not rest points of the dynamic (B). As we shall see in Section 6, the fact that the mapping between solutions is many-to-one

rather than one-to-one makes relating stability under (B) and (AB) more difficult than it may first appear to be.

Theorem 5.2 is an immediate consequence of Theorem 5.4, which characterizes the dynamics on Σ that can be aggregated. The proof of Theorem 5.4 requires the following lemma.

Lemma 5.3: *If $\{\sigma_t\} \subset \hat{\Sigma}$ is an L^1 differentiable trajectory, then $E(\dot{\sigma}_t) = \frac{d}{dt}E\sigma_t$.*

Proof: Since E is continuous by Lemma 2.1,

$$E(\dot{\sigma}_t) = E\left(L^1 \lim_{\varepsilon \rightarrow 0} \frac{\sigma_{t+\varepsilon} - \sigma_t}{\varepsilon}\right) = \lim_{\varepsilon \rightarrow 0} E\left(\frac{\sigma_{t+\varepsilon} - \sigma_t}{\varepsilon}\right) = \lim_{\varepsilon \rightarrow 0} \frac{E\sigma_{t+\varepsilon} - E\sigma_t}{\varepsilon} = \frac{d}{dt}E\sigma_t. \blacksquare$$

Theorem 5.4: *The dynamic (D) aggregates to the dynamic (AD) if and only if $(E \circ f)(\sigma) = (g \circ E)(\sigma)$ for all $\sigma \in \Sigma$.*

Proof: Let $\{\sigma_t\}$ be an L^1 solution to (D). Applying Lemma 5.3, and taking expectations of both sides of equation (D), we find that

$$\frac{d}{dt}E\sigma_t = E\dot{\sigma}_t = Ef(\sigma_t).$$

Thus, if $E \circ f = g \circ E$, it follows that $g(E\sigma_t) = Ef(\sigma_t) = \frac{d}{dt}E\sigma_t$; hence, $\{E\sigma_t\}$ solves (AD), and so f aggregates to g . Conversely, if f aggregates to g , then $\{E\sigma_t\}$ solves (AD), so $g(E\sigma_t) = \frac{d}{dt}E\sigma_t = Ef(\sigma_t)$. As σ_0 was chosen arbitrarily, it follows that $E \circ f = g \circ E$. \blacksquare

Theorem 5.4 implies that given any Lipschitz continuous function $F: \Delta \rightarrow \Sigma$, the dynamic

$$\dot{\sigma} = F(E(\sigma)) - \sigma$$

aggregates to (AD) with $g(x) = E(F(x)) - x$. Thus, dynamics on Σ can be aggregated whenever the target point $F(E(\sigma))$ is only a function of aggregate behavior. In fact, the stability results in the next section extend immediately to all dynamics of this form.

Before considering the question of stability in the next section, we use Theorem 5.2 to establish an instability result. Suppose that the aggregate best response

dynamic (AB) exhibits a limit cycle,⁸ or that equilibrium distributions are avoided in some more complicated fashion. What we can we say about behavior under the Bayesian dynamic (B)?

Theorem 5.5: *Let $\sigma_0 \in \Sigma$ and $x_0 = E\sigma_0 \in \Delta$, and let $\{\sigma_t\}$ and $\{x_t\}$ be the solutions to (B) and (AB) from σ_0 and x_0 . Let $\sigma^* \in \Sigma^*$ and $x^* = E\sigma^* \in \Delta^*$. Then*

$$\|\sigma_t - \sigma^*\| \geq |x_t - x^*| \text{ and } \|\dot{\sigma}_t\| \geq |\dot{x}_t| \text{ for all } t \geq 0.$$

In particular, if $\{x_t\}$ avoids an ε -neighborhood of Δ^ and maintains a speed of at least δ , then $\{\sigma_t\}$ avoids an ε -neighborhood of Σ^* and maintains a speed of at least δ .*

Proof: Theorem 5.2 tells us that $E\sigma_t = x_t$ for all $t \geq 0$. Hence, Lemma 2.1 implies that

$$\|\sigma_t - \sigma^*\| \geq |E\sigma_t - E\sigma^*| = |x_t - x^*|,$$

while Lemmas 2.1 and 5.3 imply that

$$\|\dot{\sigma}_t\| \geq |E\dot{\sigma}_t| = \left| \frac{d}{dt} E\sigma_t \right| = |\dot{x}_t|.$$

The remaining claims follow from these inequalities and Theorem 4.1. ■

6. Aggregation and Stability

As we discussed in the introduction, it is important in many applications to be able to predict not only the behavior distribution $E\sigma$, but the full Bayesian strategy σ . This observation motivates our stability analysis of the Bayesian dynamic (B). The results in this section establish that under three standard notions of stability, stability of Bayesian strategies under (B) is equivalent to stability of distributions under (AB). Thus, to evaluate the stability of Bayesian strategies under (B) it is enough to determine the stability of the corresponding distributions in $\Delta \subset \mathbf{R}^n$, which can be accomplished using standard techniques.

We begin by reviewing the three notions of dynamic stability we will consider.

⁸ This can occur, for example, in perturbed versions of Rock-Scissors-Paper games whose unique Nash equilibrium is not an ESS.

Let Z be a subset of a Banach space $(\hat{Z}, \|\cdot\|)$, and let the function $h: Z \rightarrow \hat{Z}$ define a dynamic on Z via the equation of motion

$$(M) \quad \dot{z} = h(z).$$

Suppose that Z is forward invariant under the dynamic (M), and let $z^* \in Z$ be a rest point of the dynamic: $h(z^*) = 0$. We say that z^* is *Lyapunov stable* under (M) if for each set $A \subset Z$ containing z^* that is open (relative to Z) there is an open set $A' \subset A$ that contains z^* such that any trajectory that begins in A' always stays in A : if $\{z_t\}$ is a solution to (M) with $z_0 \in A'$, then $\{z_t\} \subset A$. We call z^* *asymptotically stable* under (M) if it is Lyapunov stable and if there is an open set A containing z^* such that any trajectory starting in A converges to z^* : if $\{z_t\}$ is a solution to (M) with $z_0 \in A$, then $\lim_{t \rightarrow \infty} z_t = z^*$. If we can choose $A = Z$, we call z^* *globally asymptotically stable*.

The following lemma underlies many of our stability results.

Lemma 6.1 *Let $\sigma \in \Sigma$, let $x = E\sigma \in \Delta$, and let $y \in \Delta$. Then there exists a $\rho \in \Sigma$ satisfying $E\rho = y$ and $\|\rho - \sigma\| = |y - x|$.*

Lemma 6.1 says that if the Bayesian strategy σ has a distribution x that is ε away from the distribution y , there is another Bayesian strategy ρ that is ε away from σ and that aggregates to y . A constructive proof of this lemma can be found in the Appendix. The result is not obvious because in constructing ρ , we must be certain that the distribution $\rho(\pi)$ played by each preference π lies in the simplex.

We first characterize Lyapunov stability under the Bayesian dynamic.

Theorem 6.2: *(Lyapunov stability)*

The distribution $x^ \in \Delta^*$ is Lyapunov stable under (AB) if and only if the Bayesian strategy $\sigma^* = B(x^*) \in \Sigma^*$ is Lyapunov stable under (B).*

To establish this connection, we need ways of moving between neighborhoods of Bayesian equilibria σ^* and equilibrium distributions x^* . Since the operators $E: \Sigma \rightarrow \Delta$ and $B: \Delta \rightarrow \Sigma$ are continuous and map equilibria to equilibria, they along with Lemma 6.1 are the tools we need.

That the Lyapunov stability of σ^* implies the Lyapunov stability of x^* follows easily from these facts. The proof of the converse requires an additional lemma.

Lemma 6.3. *Let $A \subset \hat{\Sigma}$ be an open convex set, and let $\{\sigma_t\} \subset \hat{\Sigma}$ be an L^1 differentiable trajectory with $\sigma_0 \in A$ and such that $\sigma_t + \dot{\sigma}_t \in A$ for all t . Then $\{\sigma_t\} \subset A$.*

The point $\sigma_t + \dot{\sigma}_t$ is the location of the head of the vector $\dot{\sigma}_t$ if its tail is placed at σ_t , and so represents the target toward which the trajectory is moving at time t . The lemma, which is proved in the Appendix, says that if the trajectory starts in the open, convex set A and always moves toward points in A , it never leaves A .

Now, suppose that x^* is Lyapunov stable. If V is a convex neighborhood of σ^* , then $B^{-1}(V)$ is a neighborhood of x^* . Since x^* is Lyapunov stable, trajectories that start in some open set $W \subset B^{-1}(V)$ stay in $B^{-1}(V)$. Therefore, if the Bayesian trajectory $\{\sigma_t\}$ starts at $\sigma_0 \in E^{-1}(W) \cap V$, then the distribution trajectory $\{E\sigma_t\}$ stays in $B^{-1}(V)$, and hence the Bayesian trajectory $\{B(E(\sigma_t))\}$ stays in V . Since the trajectory $\{\sigma_t\}$ begins in V and always heads toward points $B(E(\sigma_t)) \in V$, Lemma 6.3 implies that it never leaves V .

Proof of Theorem 6.2: First, suppose that $\sigma^* = B(x^*)$ is Lyapunov stable under (B). To show that x^* is Lyapunov stable under (AB), we need to show that for each open set O containing x^* there is an open set $O' \subset O$ containing x^* such that solutions to (AB) that start in O' never leave O . Since $E: \Sigma \rightarrow \Delta$ is continuous, $E^{-1}(O)$ is open; since $E\sigma^* = x^*$ by Theorem 4.1, $\sigma^* \in E^{-1}(O)$. Because σ^* is Lyapunov stable, there exists an open ball $C \subset E^{-1}(O)$ about σ^* of radius ε such that solutions to (B) that start in C stay in $E^{-1}(O)$.

Let O' be an open ball about x^* of radius less than ε that is contained in the open set $B^{-1}(C) \cap O$. Let $\{x_t\}$ be a solution to (AB) with $x_0 \in O'$. By our choice of O' , $|x_0 - x^*| < \varepsilon$. Thus, by Lemma 6.1, there exists a Bayesian strategy σ_0 such that $E\sigma_0 = x_0$ and $\|\sigma_0 - \sigma^*\| = |x_0 - x^*| < \varepsilon$; the inequality implies that $\sigma_0 \in C$. Hence, if $\{\sigma_t\}$ is the solution to (B) starting from σ_0 , then $\{\sigma_t\} \subset E^{-1}(O)$. Therefore, Theorem 5.2 implies that $\{x_t\} = \{E\sigma_t\} \subset O$.

Now suppose that x^* is Lyapunov stable under (AB). To show that $\sigma^* = B(x^*)$ is Lyapunov stable under (B), it is enough to show that for each set $U \subset \Sigma$ containing σ^* that is open relative to Σ , there is an set $U' \subset U$ containing σ^* that is open relative to Σ such that solutions to (B) that start in U' never leave U .

Let V be an open ball in $\hat{\Sigma}$ about σ^* such that $V \cap \Sigma \subset U$. Since we can view the continuous function $B: \Delta \rightarrow \Sigma$ as having range $\hat{\Sigma} \supset \Sigma$, $B^{-1}(V) \subset \Delta$ is open relative to Δ and contains x^* . Because x^* is Lyapunov stable, we can find an set $W \subset B^{-1}(V)$ that

contains x^* and that is open relative to Δ such that solutions to (AB) that start in W never leave $B^{-1}(V)$.

The set $E^{-1}(W)$ is open relative to Σ and contains σ^* ; therefore, $U' = E^{-1}(W) \cap V$ possesses both of these properties as well. Let $\{\sigma_t\}$ be a solution to (B) with $\sigma_0 \in U'$. Then $E\sigma_0 \in W$. Therefore, since $\{E\sigma_t\}$ solves (AB) by Theorem 5.2, $E\sigma_t \in B^{-1}(V)$ for all t , and so $B(E(\sigma_t)) \in V$ for all t . But since $\dot{\sigma}_t = B(E(\sigma_t)) - \sigma_t$, $\sigma_t + \dot{\sigma}_t \in V$ for all t . Thus, Lemma 6.3 implies that $\{\sigma_t\} \subset V$. Moreover, Theorem 3.1 (ii) implies that $\{\sigma_t\} \subset \Sigma$; we therefore conclude that $\{\sigma_t\} \subset V \cap \Sigma \subset U$. ■

We continue by characterizing asymptotic stability.

Theorem 6.4: (*Asymptotic stability*)

The distribution $x^ \in \Delta^*$ is asymptotically stable under (AB) if and only if the Bayesian strategy $\sigma^* = B(x^*) \in \Sigma^*$ is asymptotically stable under (B).*

That the asymptotic stability of the Bayesian strategy σ^* implies the asymptotic stability of its distribution x^* follows easily from Lemma 6.1 and Theorem 5.2. The proof of the converse also requires the following lemma.

Lemma 6.5: *Let $\{\sigma_t\}$ be the solution to (B) from some $\sigma_0 \in \Sigma$ with $E\sigma_0 = x^* \in \Delta^*$, and let $\sigma^* = B(x^*) \in \Sigma^*$. Then $\lim_{t \rightarrow \infty} \sigma_t = \sigma^*$. Indeed,*

$$\sigma_t = e^{-t} \sigma_0 + (1 - e^{-t}) \sigma^*.$$

If $\{\sigma_t\}$ is a Bayesian trajectory whose initial distribution is an equilibrium, then while σ_t may change over time, its distribution does not: $E\sigma_t = x^*$ for all t . Consequently, under the best response dynamic (B), σ_t always heads from σ_0 directly toward the point $B(E(\sigma_t)) = \sigma^*$. The proof of Lemma 6.5 can be found in the Appendix.

Now, suppose that x^* is asymptotically stable under (AB). Then if σ_0 is close enough to σ^* , $E\sigma_0$ will be close to x^* , so if $\{\sigma_t\}$ solves (B), Theorem 5.2 tells us that $\{E\sigma_t\}$ converges to x^* . Lemma 6.1 then implies that if t is large, we can find a Bayesian strategy $\hat{\sigma}$ that is close to σ_t and that aggregates to x^* ; by Lemma 6.5, the solution to (B) from $\hat{\sigma}$ converges to σ^* . That $\{\sigma_t\}$ must converge to σ^* then follows from the continuity of solutions to (B) in their initial conditions.

Proof of Theorem 6.4: Since Lyapunov stability is covered by Theorem 6.2, we need only consider convergence of nearby trajectories to σ^* and x^* . For all $\varepsilon > 0$ and any $\sigma \in \Sigma$ and $x \in \Delta$, define $N_\varepsilon(\sigma) = \{\rho \in \Sigma: \|\rho - \sigma\| \leq \varepsilon\}$ and $N_\varepsilon(x) = \{y \in \Delta: |y - x| \leq \varepsilon\}$ to be the ε neighborhoods of σ and x , respectively.

Suppose that $\sigma^* = B(x^*)$ is asymptotically stable. Then there exists an $\varepsilon > 0$ such that solutions to (B) with $\sigma_0 \in N_\varepsilon(\sigma^*)$ converge to σ^* . Now suppose that $\{x_t\}$ is a solution to (AB) with $x_0 \in N_\varepsilon(x^*)$. By Lemma 6.1, there exists a $\hat{\sigma}_0 \in N_\varepsilon(\sigma^*)$ satisfying $E\sigma_0 = x_0$; therefore, the solution $\{\hat{\sigma}_t\}$ converges to σ^* . Since $x_t = E\hat{\sigma}_t$ by Theorem 5.2, and since E is continuous by Lemma 2.1,

$$\lim_{t \rightarrow \infty} x_t = \lim_{t \rightarrow \infty} E\hat{\sigma}_t = E\left(L^1 \lim_{t \rightarrow \infty} \hat{\sigma}_t\right) = E\sigma^* = x^*.$$

Hence, all solutions to (AB) starting in $N_\varepsilon(x^*)$ converge to x^* .

Now suppose that x^* is asymptotically stable and let $\sigma^* = B(x^*)$. We can choose an $\varepsilon > 0$ such that all solutions to (AB) starting in $N_\varepsilon(x^*)$ converge to x^* . Now suppose that $\sigma_0 \in N_\varepsilon(\sigma^*)$; we will show that $\{\sigma_t\}$, the solution to (B) starting from σ_0 , must converge to σ^* . First, observe that

$$|E\sigma_0 - x^*| = |E\sigma_0 - E(B(x^*))| = |E(\sigma_0 - \sigma^*)| \leq E|\sigma_0 - \sigma^*| = \|\sigma_0 - \sigma^*\| \leq \varepsilon,$$

so $E\sigma_0 \in N_\varepsilon(x^*)$. Theorem 5.2 implies that $\{E\sigma_t\}$ is the solution to (AB) starting from $E\sigma_0$; hence, $\lim_{t \rightarrow \infty} E\sigma_t = x^*$.

Fix $\eta > 0$. It is enough to show that there exists a T such that $\|\sigma_t - \sigma^*\| < \eta$ for all $t \geq T$. Let K be the Lipschitz coefficient of $f(\sigma) = B(E(\sigma)) - \sigma$, and let $\delta = n^{-K} \left(\frac{\eta}{2}\right)^{K+1}$. Since $\lim_{t \rightarrow \infty} E\sigma_t = x^*$, there is a τ_1 such that $|E\sigma_t - x^*| < \delta$ whenever $t \geq \tau_1$. Let $\tau_2 = \ln \frac{2n}{\eta}$, and choose $T = \tau_1 + \tau_2$.

Fix $t > T$. Then since $t - \tau_2 > T - \tau_2 = \tau_1$, Lemma 6.1 implies that there is a $\hat{\sigma}_0$ such that $E\hat{\sigma}_0 = x^*$ and

$$\|\sigma_{t-\tau_2} - \hat{\sigma}_0\| = |E\sigma_{t-\tau_2} - x^*| < \delta.$$

Let $\{\hat{\sigma}_t\}$ be the solution to (B) with initial condition $\hat{\sigma}_0$. Since no two points in Σ are further than distance n apart, Lemma 6.5 implies that

$$\|\hat{\sigma}_{\tau_2} - \sigma^*\| = e^{-\tau_2} \|\hat{\sigma}_0 - \sigma^*\| \leq ne^{-\tau_2}$$

Moreover, it follows from Theorem 3.1 (ii) that

$$\|\sigma_t - \hat{\sigma}_{\tau_2}\| \leq \|\sigma_{t-\tau_2} - \hat{\sigma}_0\| e^{K\tau_2}.$$

Therefore,

$$\begin{aligned} \|\sigma_t - \sigma^*\| &\leq \|\sigma_t - \hat{\sigma}_{\tau_2}\| + \|\hat{\sigma}_{\tau_2} - \sigma^*\| \\ &\leq \|\sigma_{t-\tau_2} - \hat{\sigma}_0\| e^{K\tau_2} + ne^{-\tau_2} \\ &< \delta e^{K\tau_2} + ne^{-\tau_2} \\ &= \frac{\eta}{2} + \frac{\eta}{2} = \eta. \blacksquare \end{aligned}$$

We conclude this section by characterizing global asymptotic stability. The proof of this result is analogous to that of Theorem 6.4.

Theorem 6.6: (*Global asymptotic stability*)

The distribution $x^ \in \Delta^*$ is globally asymptotically stable under (AB) if and only if the Bayesian strategy $\sigma^* = B(x^*) \in \Sigma^*$ is globally asymptotically stable under (B).*

Remark 6.7: While we have stated our stability results for isolated equilibria $x^* \in \Delta^*$ and $\sigma^* = B(x^*) \in \Sigma^*$, one can extend these results to allow for closed connected sets of equilibria $X^* \subset \Delta^*$ and $B(X^*) \subset \Sigma^*$.

7. Best Response Dynamics and Fictitious Play

Under common preferences, the close connections between the best response dynamic and fictitious play are well known. In fictitious play, players always choose a best response to their beliefs, which are given by the time average of past play. In the continuous time formulation, if we let c_t denote current behavior, then $a_t = \frac{1}{t} \int_0^t c_s ds$ represents beliefs. The requirement that current behavior is a best response to beliefs is stated as $c_t \in BR^\pi(a_t)$. By differentiating the definition of a_t and substituting, we obtain the law of motion for beliefs under fictitious play:

$$(FP) \quad \dot{a}_t = \frac{1}{t} c_t - \frac{1}{t^2} \int_0^t c_s ds$$

$$= \frac{1}{t}(BR^\pi(a_t) - a_t)$$

Therefore, after a reparameterization of time, the evolution of beliefs under (FP) is identical to the evolution of behavior under the best response dynamic (BR).

Ellison and Fudenberg (2000) study fictitious play in a population with diverse preferences. As in the standard case, players choose a best response to the time average $a_t = \frac{1}{t} \int_0^t c_s ds$ of past behavior. Since players are matched with opponents drawn from the population as a whole, the object that is averaged to determine beliefs is $c_t = E(B(a_t))$, the distribution of behavior at time t . This yields the law of motion

$$(PFP) \quad \dot{a}_t = \frac{1}{t}(E(B(a_t)) - a_t),$$

which is a reparameterization of our aggregate best response dynamic (AB).

Observe that the state variable under population fictitious play is the average *distribution* of past behavior, $a_t \in \Delta$. If one keeps track of this, one can always compute the best response $B(a_t) \in \Sigma$ as well as the best response distribution $E(B(a_t)) \in \Delta$. The latter object determines the direction in which the time average a_t evolves. In contrast, the Bayesian best response dynamic must specify how behavior in every subpopulation evolves, so the relevant state variable is not the distribution of behavior $x_t \in \Delta$, but the Bayesian strategy $\sigma_t \in \Sigma_t$. Thus, while the dynamics (PFP) and (AB) are nearly identical, the evolution of Bayesian strategies under population fictitious play and under the Bayesian best response dynamic are quite different.

As an illustration, suppose that (PFP) and (AB) are currently at state $a_t = x_t$. Under population fictitious play, the current Bayesian strategy must be $B(a_t)$, the *best response* to beliefs a_t ; in particular, it is always pure. Under the best response dynamic, the current Bayesian strategy σ_t must be one that *aggregates* to a_t : in other words, $\sigma_t \in E^{-1}(a_t) = \{\sigma \in \Sigma: E(\sigma) = a_t\}$. In fact, $B(a_t)$ is contained in $E^{-1}(a_t)$ only if a_t is an equilibrium distribution and $\sigma_t = B(a_t)$ is the corresponding Bayesian equilibrium.

On the other hand, if the solution to (AB) from $a_t = x_t$ converges to some equilibrium distribution x^* , then one can show that under both population fictitious play and the Bayesian best response dynamic, the Bayesian strategy σ_t converges to the Bayesian equilibrium $\sigma^* = B(x^*)$. Indeed, by proving that equilibrium and stability analyses for the Bayesian best response dynamic (B) can be performed directly in terms of the aggregate dynamic (AB), we have demonstrated

that the close connections between fictitious play and the best response dynamic from the common preference case persist when preferences are diverse.

Ellison and Fudenberg (2000) use their model of population fictitious play to investigate the evolutionary stability of purified equilibria (Harsanyi (1973)), obtaining stability and instability results for 2×2 and 3×3 games. By building on the results in this paper and recent studies of perturbed best response dynamics, one can investigate the stability of purified equilibria in games with arbitrary numbers of strategies. This question is pursued in a companion paper, Sandholm (2003).

Appendix

A.1 Basic Properties of Dynamical Systems on Δ and Σ

We begin this appendix by establishing the existence, uniqueness, and forward invariance of solution trajectories of the aggregate best response dynamic (AB). In fact, we will establish this result for a more general class of dynamics on the simplex.

Let $g: \Delta \rightarrow \mathbf{R}^n$ be a vector field on the simplex that satisfies

- (LC) g is Lipschitz continuous on Δ .
- (FI1) $\sum_i g_i(x) = 0$ for all $x \in \Delta$.
- (FI2) For all $x \in \Delta$, $g_i(x) \geq 0$ whenever $x_i = 0$.

Condition (LC) is the usual Lipschitz condition used to prove the existence of unique solution trajectories to the differential equation $\dot{x} = g(x)$. Condition (FI1) says that $\sum_i \dot{x}_i = 0$, implying the invariance of the affine space $\tilde{\Delta} = \{x \in \mathbf{R}^n: \sum_i x_i = 1\}$. Condition (FI2) says that whenever the component x_i equals zero, its rate of change \dot{x}_i is non-negative.

Theorem A.1: *Let $g: \Delta \rightarrow \mathbf{R}^n$ satisfy (LC), (FI1), and (FI2), and let \tilde{g} be a Lipschitz continuous extension of g from Δ to $\tilde{\Delta}$. Then solutions to $\dot{x} = \tilde{g}(x)$ from each $x_0 \in \Delta$ exist, are unique, are Lipschitz continuous in x_0 , and remain in Δ at all positive times.*

Let $|x|_E = \sqrt{\sum_i x_i^2}$ denote the Euclidean norm on \mathbf{R}^n . Then the proof of Theorem A.1 relies on a geometrically obvious observation that we state without proof.

Observation A.2: Let C be a compact, convex subset of \mathbf{R}^n , and define the closest point function $c: \mathbf{R}^n \rightarrow C$ by

$$c(x) = \arg \min_{z \in C} |x - z|_E$$

Then $|c(x) - c(y)|_E \leq |x - y|_E$ for all $x, y \in \mathbf{R}^n$. Hence, by the equivalence of norms on \mathbf{R}^n , there exists a $k > 0$ such that $|c(x) - c(y)| \leq k|x - y|$ for all $x, y \in \mathbf{R}^n$.

Proof of Theorem A.1: Define $\hat{g}: \tilde{\Delta} \rightarrow \mathbf{R}^n$ by $\hat{g}(x) = g(c(x))$. Then condition (LC) and Observation A.2 imply that \hat{g} is Lipschitz. Therefore, standard results (e.g., Hirsch and Smale (1974, Chapter 8)) show that solutions to $\dot{x} = \hat{g}(x)$ exist, are unique, and are Lipschitz continuous in their initial conditions. The forward invariance of Δ under $\dot{x} = \hat{g}(x)$ follows from Theorem 5.7 of Smirnov (2002).

Now consider any Lipschitz continuous extension \tilde{g} of g to $\tilde{\Delta}$, and fix an initial condition $x_0 \in \Delta$. Since the solution $\{x_t\}_{t \geq 0}$ to $\dot{x} = \hat{g}(x)$ starting from x_0 does not leave Δ , and since \tilde{g} and \hat{g} are identical on Δ , this solution is also a solution to $\dot{x} = \tilde{g}(x)$. But since \tilde{g} is Lipschitz, this must be the only solution to $\dot{x} = \tilde{g}(x)$ from x_0 . We therefore conclude that Δ is forward invariant under \tilde{g} . Since Δ is closed, forward invariance implies that the solution is well defined at all times $t \in [0, \infty)$ (see, e.g., Hale (1969, p. 17-18)). ■

We now prove an analogue of Theorem A.1 for dynamics on Σ . Let $f: \Sigma \rightarrow \hat{\Sigma}$ satisfy

- (LC') f is L^1 Lipschitz continuous on Σ .
- (FI1') $\sum_i f_i(\sigma)(\pi) = 0$ for all $\sigma \in \Sigma$ and $\pi \in \Pi$.
- (FI2') For all $\sigma \in \Sigma$ and $\pi \in \Pi$, $f_i(\sigma)(\pi) \geq 0$ whenever $\sigma_i(\pi) = 0$.
- (UB) For all $\sigma \in \Sigma$ and $\pi \in \Pi$, $|f(\sigma)(\pi)| \leq M$

The first three conditions are analogues of the conditions considered previously. Condition (FI1') ensures that solutions stay in the affine space $\tilde{\Sigma} = \{\sigma \in \hat{\Sigma}: \sigma(\pi) \in \tilde{\Delta} \text{ for all } \pi \in \Pi\}$, while condition (FI2') ensures that whenever no one in subpopulation π uses strategy i , the growth rate of strategy i in this subpopulation is non-negative. Finally, condition (UB) places a uniform bound on $f(\sigma)(\pi)$, which is needed because $f(\sigma)$ is infinite dimensional.

Existence, uniqueness, continuity in initial conditions, and forward invariance of Σ for L^1 solutions to $\dot{\sigma} = f(\sigma)$ are established in Theorem A.3. This result implies parts (i), (ii), and (iii) of Theorem 3.1.

Theorem A.3: *Let $f: \Sigma \rightarrow \mathbf{R}^n$ satisfy (LC'), (FI1'), (FI2'), and (UB), and let \tilde{f} be a Lipschitz continuous extension of f from Σ to $\tilde{\Sigma}$. Then solutions to $\dot{\sigma} = \tilde{f}(\sigma)$ from each $\sigma_0 \in \Sigma$ exist, are unique, are Lipschitz continuous in σ_0 , and remain in Σ at all positive times.*

In addition to these properties, we would also like to establish that some L^1 solution $\{\sigma_t\}$ has continuous sample paths: in other words, that for each $\pi \in \Pi$, the behavior $\sigma_t(\pi)$ of the subpopulation with preference π changes continuously over time. While not every L^1 solution has this property, we can prove that there is always one that does. Call $\{\tilde{\sigma}_t\}$ a *modification* of $\{\sigma_t\}$ if $\mu(\pi: s_t(\pi) = \tilde{s}_t(\pi)) = 1$ for all t .

Theorem A.4: *Let $\{\sigma_t\}$ be an L^1 solution to $\dot{\sigma} = \tilde{f}(\sigma)$, where $\tilde{f}: \tilde{\Sigma} \rightarrow \hat{\Sigma}$ is L^1 continuous and pointwise bounded. Then there exists a modification $\{\tilde{\sigma}_t\}$ of $\{\sigma_t\}$ with continuous sample paths: i.e., such that $\mu(\pi: \tilde{\sigma}_t(\pi)$ is continuous in $t) = 1$.*

While of interest in its own right (in particular, because it implies Theorem 4.1(iv)), Theorem A.4 is also useful for proving Theorem A.3.

To prove these two results, we introduce the notion of an L^1 integral of a trajectory through $\hat{\Sigma}$; for a complete treatment, see Lang (1997, Chapter 10). If $\{\sigma_t\}$ is an L^1 continuous trajectory through $\hat{\Sigma}$, then its L^1 integral over the interval $[a, b]$, denoted $\int_a^b \sigma_t dt$, is the L^1 limit of the integrals of any sequence of step functions $\{\sigma_t^n\}$ satisfying $\lim_{n \rightarrow \infty} \sup_{t \in [a, b]} \|\sigma_t^n - \sigma_t\| = 0$. If $\{\sigma_t\}$ is an L^1 solution to $\dot{\sigma} = \tilde{f}(\sigma)$, then by definition we have that $\sigma_u = \sigma_0 + \int_0^u \tilde{f}(\sigma_t) dt$. Moreover, if $\tau: \Pi \rightarrow [0, u]$ is a random time and f is pointwise bounded, then a step function approximation can be used to show that $\sigma_u = \sigma_\tau + \int_0^u \tilde{f}(\sigma_t) 1_{\{t \geq \tau\}} dt$.

Proof of Theorem A.3:

Define $\hat{f}: \tilde{\Sigma} \rightarrow \hat{\Sigma}$ by $\hat{f}(\sigma) = f(c(\sigma))$, where $c(\sigma)(\pi) \equiv c(\sigma(\pi))$. Then for all $\sigma, \rho \in \tilde{\Sigma}$,

$$\|\hat{f}(\sigma) - \hat{f}(\rho)\| = \|f(c(\sigma)) - f(c(\rho))\|$$

$$\begin{aligned}
&\leq K\|c(\sigma) - c(\rho)\| \\
&= K \cdot E|c(\sigma(\pi)) - c(\rho(\pi))| \\
&\leq K \cdot Ek|\sigma(\pi) - \rho(\pi)| \\
&= Kk\|\sigma - \rho\|,
\end{aligned}$$

where K and k are the Lipschitz constants for f and c , respectively. Hence, \hat{f} is L^1 Lipschitz on $\tilde{\Sigma}$. Therefore, standard results imply that there exist unique solutions to $\dot{\sigma} = \hat{f}(\sigma)$ from each initial condition $\sigma_0 \in \tilde{\Sigma}$, and that solutions are Lipschitz continuous in σ_0 .

Let $\sigma_0 \in \Sigma$, let $\{\sigma_t\}$ be the L^1 solution to $\dot{\sigma} = \hat{f}(\sigma)$ from σ_0 , and suppose that $\sigma_u \notin \Sigma$ for some u . Then for some strategy i the set $A_i = \{\pi \in \Pi: [\sigma_u(\pi)]_i < 0\}$ has positive measure under μ . By Theorem A.4, we can suppose that $\{\sigma_t\}$ has continuous sample paths. Hence, the random time $\tau(\pi) = \max\{t \leq u: [\sigma_t(\pi)]_i \geq 0\}$ is well defined and is strictly less than u when $\pi \in A_i$.

Observe that if $\sigma \in \tilde{\Sigma}$ has $\sigma_i(\pi) \leq 0$, then $c_i(\sigma)(\pi) = 0$, and hence $\hat{f}_i(\sigma)(\pi) = f_i(c(\sigma))(\pi) \geq 0$ by condition (FI2'). We therefore have the following L^1 integral inequality:

$$[\sigma_u]_i = [\sigma_\tau]_i + \int_0^u \hat{f}_i(\sigma_t) 1_{\{t \geq \tau\}} dt \geq [\sigma_\tau]_i$$

Observe that $[\sigma_\tau(\pi)]_i = 0$ when $\pi \in A_i$. Hence, for almost every $\pi \in A_i$, $[\sigma_u(\pi)]_i \geq 0$, contradicting the definition of A_i . Therefore, the trajectory $\{\sigma_t\}$ cannot leave Σ , which is thus forward invariant under $\dot{\sigma} = \hat{f}(\sigma)$.

Forward invariance of Σ under any Lipschitz continuous extension of f to $\tilde{\Sigma}$ is proved in the same fashion as the analogous part of Theorem A.1. ■

We now prove Theorem A.4. To do so, we introduce the L^2 norm on $\hat{\Sigma}$:

$$\|\sigma\|_{L^2} = \sqrt{\sum_{i=1}^n E\sigma_i^2}.$$

If the trajectory $\{\sigma_t\}$ is L^2 continuous, then its L^2 integral, $\oint_a^b \sigma_t dt$, is defined as the L^2 limit of the integrals of any sequence of step functions $\{\sigma_t^n\}$ satisfying $\lim_{n \rightarrow \infty} \sup_{t \in [a, b]} \|\sigma_t^n - \sigma_t\|_{L^2} = 0$. The L^2 integral satisfies this standard inequality:

$$\left\| \oint_a^b \sigma_t dt \right\|_{L^2} \leq \int_a^b \|\sigma_t\|_{L^2} dt.$$

Since μ is a finite measure, the L^1 and L^2 norms define the same topology on any set of functions satisfying the uniform boundedness condition (UB). In particular, $L^1 \lim_{s \rightarrow t} \sigma_s = L^2 \lim_{s \rightarrow t} \sigma_s$ whenever either limit exists. It follows that if $\{\sigma_t\}$ is pointwise bounded and L^1 continuous, its L^1 and L^2 integrals are the same: $\int_a^b \sigma_t dt = \int_a^b \sigma_t dt$.

The proof of Theorem A.4 relies on Lemma A.5, which is a direct implication of the Kolmogorov continuity theorem (Karatzas and Shreve (1991, Theorem 2.2.8)).

Lemma A.5: *Suppose that $\{\sigma_t\}$ is L^2 Lipschitz continuous (i.e., that there is a constant $K < \infty$ such that $\|\sigma_t - \sigma_s\|_{L^2} \leq K|t - s|$ for all s and t). Then there exists a modification $\{\tilde{\sigma}_t\}$ of $\{\sigma_t\}$ such that $\mu(\pi: \tilde{\sigma}_t(\pi)$ is continuous in $t) = 1$.*

Proof of Theorem A.4:

By definition, the trajectory $\{\sigma_t\}$ satisfies the L^1 integral equation

$$\sigma_t = \sigma_0 + \int_0^t \tilde{f}(\sigma_s) ds.$$

Since the function \tilde{f} is L^1 continuous and pointwise bounded by some constant M , the trajectory $\{\tilde{f}(\sigma_t)\}$ is as well. Hence, $\int_s^t \tilde{f}(\sigma_u) du = \int_s^t \tilde{f}(\sigma_u) du$, and so

$$\|\sigma_t - \sigma_s\|_{L^2} = \left\| \int_s^t \tilde{f}(\sigma_u) du \right\|_{L^2} = \left\| \int_s^t \tilde{f}(\sigma_u) du \right\|_{L^2} \leq \int_s^t \|\tilde{f}(\sigma_u)\|_{L^2} du \leq M|t - s|.$$

That is, $\{\sigma_t\}$ is L^2 Lipschitz. The result therefore follows from Lemma A.5. ■

A.2 Other Proofs

Proof of Proposition 2.2:

Condition (C1), which requires that B is single valued, obviously holds, so we focus on the Lipschitz continuity condition (C2). In this proof, we use the Euclidean norm $|x|_E = \sqrt{\sum_i x_i^2}$ for points in \mathbf{R}^n . Since this norm is equivalent to the summation norm, our proof implies the result for the latter norm as well. It is enough to show that the Lipschitz inequality $\|B(x) - B(y)\| \leq C |x - y|_E$ holds when $|x - y|_E$ is sufficiently small.

Fix $x, y \in \Delta$ and $i \neq j$. The set of payoff matrices that choose i over j at x and j over

i at y is

$$\begin{aligned}\Pi_{ij} &= \{A: (Ax)_i > (Ax)_j \text{ and } (Ay)_i < (Ay)_j\} \\ &= \{A: (A_i - A_j) \cdot x > 0 > (A_i - A_j) \cdot y\}.\end{aligned}$$

Here, A_i and A_j denote rows of the matrix A .

We can associate with each payoff matrix A a difference vector $d_{ij} = A_i - A_j \in \mathbf{R}^n$. Let $f: \mathbf{R}^{n \times n} \rightarrow \mathbf{R}$ denote the density function of the measure λ , and let $g_{ij}: \mathbf{R}^n \rightarrow \mathbf{R}$ be the density of the measure on the difference d_{ij} that is induced by λ . If $[-c, c]^{n \times n}$ contains the support of f , and M is an upper bound on f , then by integrating out irrelevant components and changing variables, one can show that

$$g_{ij}(d) \leq (2c)^{n^2-n} M \quad \text{for all } d \in \mathbf{R}^n.$$

Moreover, the support of g_{ij} is contained in the cube $[-2c, 2c]^n$, and hence in the ball $S \subset \mathbf{R}^n$ centered at the origin with radius $r = 2c\sqrt{n}$.

Let

$$D_{ij} = \{d \in S: d \cdot x > 0 > d \cdot y\},$$

and let m represent Lebesgue measure on \mathbf{R}^n . Suppose we can show that

$$m(D_{ij}) \leq K |x - y|_E \tag{3}$$

for some K independent of x , y , i , and j . Then since a change in best response requires a reversal of preferences for at least one strategy pair, it follows that

$$\begin{aligned}\|B(x) - B(y)\| &= 2\mu(\pi: B(x)(\pi) \neq B(y)(\pi)) \\ &\leq 2 \sum_{i,j \neq i} \lambda(\Pi_{ij}) \\ &\leq 2 \sum_{i,j \neq i} (2c)^{n^2-n} M m(D_{ij}) \\ &\leq 2(n^2 - n) (2c)^{n^2-n} MK |x - y|_E.\end{aligned} \tag{4}$$

To bound $m(D_{ij})$, we first change coordinates in \mathbf{R}^n via an orthogonal transformation $T \in \mathbf{R}^{n \times n}$ so that $x' = Tx$ and $y' = Ty$ satisfy $x' = (x'_1, 0, 0, \dots, 0)$ and y'

$= (y'_1, y'_2, 0, \dots, 0)$, with $x'_1, y'_1, y'_2 \geq 0$. The orthogonal operator T is the composition of a sequence of rotations and reflections, and so preserves Euclidean distance, inner products, and Lebesgue measure (see Friedberg, Insel, and Spence (1989, Sections 6.5 and 6.10)). Hence, $D_{ij} = \{d \in S: Td \cdot Tx > 0 > Td \cdot Ty\}$, and so

$$\begin{aligned} D'_{ij} &= \{d' \in S: d' \cdot x' > 0 > d' \cdot y'\} \\ &= \{d' \in S: d' \cdot Tx > 0 > d' \cdot Ty\} \\ &= \{d' \in S: d' = Td \text{ for some } d \in D_{ij}\} \end{aligned}$$

Therefore, $m(D_{ij}) = m(D'_{ij})$.

Whether a vector is an element of D'_{ij} only depends on its first two coordinates. For $d' \in S$, let $\alpha(d') \in [0, 2\pi)$ be the amount by which the vector $(1, 0) \in \mathbf{R}^2$ must be rotated counterclockwise before it points in the same direction as (d'_1, d'_2) . Since all $d' \in D'_{ij}$ form acute angles with x' and obtuse angles with y' , we see that

$$\begin{aligned} D'_{ij} &= \{d' \in S: \alpha(d') \in [0, \frac{\pi}{2}) \cup (\frac{3\pi}{2}, 2\pi) \text{ and } \alpha(d') \in (\alpha(y') + \frac{\pi}{2}, \alpha(y') + \frac{3\pi}{2})\} \\ &= \{d' \in S: \alpha(d') \in (\frac{3\pi}{2}, \alpha(y') + \frac{3\pi}{2})\}. \end{aligned}$$

Hence, since $m(S) < (2r)^n$,

$$m(D'_{ij}) = \frac{\alpha(y')}{2\pi} m(S) < \alpha(y') (2r)^n. \quad (5)$$

Therefore, to prove inequality (3) it is enough to show that

$$\alpha(y') \leq k|x - y|_E = k|x' - y'|_E. \quad (6)$$

(To see why the equality in expression (5) holds, let (X_1, X_2, \dots, X_n) represent a random vector drawn from a uniform distribution on the ball S . Then the random angle Θ formed by the first two components (defined by $(\cos\Theta, \sin\Theta) = (X_1/\sqrt{X_1^2 + X_2^2}, X_2/\sqrt{X_1^2 + X_2^2})$) is independent of the remaining components.)

To establish inequality (6), we fix $c > \varepsilon \geq 0$ and let $Z_\varepsilon = \{z \in \mathbf{R}^2: |(c, 0) - (z_1, z_2)|_E = \varepsilon, z_2 \geq 0\}$ be the set of vectors in \mathbf{R}^2 with a positive second component that are ε away from the vector $(c, 0)$. The largest possible angle between the vector $(1, 0)$ and a vector in Z_ε is

$$\theta(\varepsilon) \equiv \max_{z \in \mathcal{Z}_\varepsilon} \alpha(z) = \cos^{-1} \left(\min_{z \in \mathcal{Z}_\varepsilon} \cos(\alpha(z)) \right) = \cos^{-1} \left(\min_{z \in \mathcal{Z}_\varepsilon} \frac{(1,0) \cdot (z_1, z_2)}{|(1,0)|_E |(z_1, z_2)|_E} \right).$$

If we let $\delta = c - z_1$, then the minimization problem becomes

$$\min_{\delta \in [0, \varepsilon]} \frac{(1,0) \cdot (c - \delta, \sqrt{\varepsilon^2 - \delta^2})}{|(c - \delta, \sqrt{\varepsilon^2 - \delta^2})|_E} = \min_{\delta \in [0, \varepsilon]} \frac{c - \delta}{\sqrt{c^2 - 2c\delta + \varepsilon^2}}.$$

Taking the derivative of this expression with respect to δ and setting it equal to zero yields $\delta = \frac{\varepsilon^2}{c}$; hence,

$$\theta(\varepsilon) = \cos^{-1} \left(\frac{\sqrt{c^2 - \varepsilon^2}}{c} \right).$$

It follows that $\theta(0) = 0$ and that $\theta'(\varepsilon) = 1/\sqrt{c^2 - \varepsilon^2}$ whenever $\varepsilon < c$. Therefore, if $c \geq 1/\sqrt{n}$ and $\varepsilon \leq 1/\sqrt{2n}$, then $\theta'(\varepsilon) \leq \sqrt{2n}$, and so

$$\theta(\varepsilon) \leq \sqrt{2n} \varepsilon.$$

Now suppose that $|x - y|_E \leq 1/\sqrt{2n}$. Then since $x'_1 = |x'|_E = |x|_E \geq 1/\sqrt{n}$, setting $c = x'_1$ and $\varepsilon = |x - y|_E = |x' - y'|_E$ yields

$$\alpha(y') \leq \theta(|x - y|_E) \leq \sqrt{2n} |x - y|_E,$$

establishing inequality (6) for all cases in which $|x - y|_E$ is sufficiently small. Thus, inequality (5) implies that

$$m(D_{ij}) = m(D'_{ij}) \leq (2r)^n \cdot \sqrt{2n} |x - y|_E,$$

and so inequalities (3) and (4) let us conclude that

$$\|B(x) - B(y)\| \leq 2(n^2 - n)(2c)^{n^2 - n} M \cdot (2r)^n \cdot \sqrt{2n} |x - y|_E. \blacksquare$$

Proof of Proposition 2.3:

Again, condition (C1) clearly holds, so we need only consider the Lipschitz continuity condition (C2). Fix $x, y \in \Delta$ and $i \neq j$. Let $\Pi_{ij} \subset \Pi$ represent the set of preferences that prefer strategy i to strategy j at distribution x but prefer j to i at y :

$$\Pi_{ij} = \{\pi: \pi_i(x) > \pi_j(x) \text{ and } \pi_i(y) < \pi_j(y)\}$$

Then by definition, $\mu(\Pi_{ij}) = \nu(D_{ij})$, where $D_{ij} \subset \mathbf{R}^n$ is given by

$$\begin{aligned} D_{ij} &= \{\theta: F_i(x) + \theta_i > F_j(x) + \theta_j \text{ and } F_i(y) + \theta_i > F_j(y) + \theta_j\} \\ &= \{\theta: F_i(x) - F_j(x) > \theta_j - \theta_i > F_i(y) - F_j(y)\}. \end{aligned}$$

Now suppose we can show that $\nu(D_{ij}) \leq K|x - y|$ for some K that is independent of x, y, i , and j . Then

$$\begin{aligned} \|B(x) - B(y)\| &= 2 \mu(\pi: B(x)(\pi) \neq B(y)(\pi)) \\ &\leq 2 \sum_{i,j \neq i} \mu(\Pi_{ij}) \\ &= 2 \sum_{i,j \neq i} \nu(D_{ij}) \\ &\leq 2(n^2 - n)K|x - y|. \end{aligned}$$

Each vector $\theta \in \mathbf{R}^n$ is associated with a single value of $\theta_j - \theta_i \in \mathbf{R}$. Let $f: \mathbf{R}^n \rightarrow \mathbf{R}$ denote the density of the measure ν , let M be the upper bound on f , and let $g_{ji}: \mathbf{R} \rightarrow \mathbf{R}$ denote the density of the measure for the difference $\theta_j - \theta_i$ that is induced by ν . In case (i), there is a compact set $[-c, c]^n$ that contains the support of f ; by integrating out irrelevant components and changing variables, one can show that

$$g_{ji}(d) \leq (2c)^{n-2} M \quad \text{for all } d \in \mathbf{R}.$$

In case (ii), $f(\theta) = \prod_i f_i(\theta_i)$ for some marginal densities f_i . Since f is bounded, there is also a constant \hat{M} that bounds all of the functions f_i . Hence, a convolution yields

$$g_{ji}(d) \leq \int_{-\infty}^{\infty} f_j(d - z) f_i(-z) dz = E f_j(d + \theta_i) \leq \hat{M} \quad \text{for all } d \in \mathbf{R}.$$

Thus, cases (i) and (ii) both imply case (iii): $g_{ji} \leq \bar{M}$ for some constant \bar{M} .

The interval of values of $\theta_j - \theta_i$ that lie in the set D_{ij} has length

$$(F_i(x) - F_j(x)) - (F_i(y) - F_j(y)) = (F_i(x) - F_i(y)) + (F_j(y) - F_j(x)) \leq 2K_F|x - y|,$$

where K_F is the Lipschitz coefficient for F . Therefore,

$$\nu(D_{ij}) \leq \bar{M} \cdot 2K_F |x - y|,$$

and we conclude that

$$\|B(x) - B(y)\| \leq 2(n^2 - n) \cdot \bar{M} \cdot 2K_F |x - y|. \blacksquare$$

The proof of Lemma 6.1 relies on the following observation.

Lemma A.6: *Let $a, b \in \mathbf{R}^n$. If a and b lie in the same orthant (i.e., if $a_i \geq 0 \Leftrightarrow b_i \geq 0$), then $|a + b| = |a| + |b|$.*

Proof of Lemma 6.1:

Let $x = E\sigma$, let $d = y - x$, and let $C = \{k: d_k < 0\}$. For all $k \in C$, define $\delta^k \in \mathbf{R}^n$ by

$$\delta_j^k = \begin{cases} d_k & \text{if } j = k, \\ 0 & \text{if } j \in C - \{k\}, \\ -\left(\frac{d_j}{\sum_{i \in C} d_i}\right) d_k & \text{if } j \notin C. \end{cases}$$

Notice that $\sum_j \delta_j^k = 0$ for each k and that $\sum_{k \in C} \delta^k = d$. Moreover, since each δ^k lies in the same orthant of \mathbf{R}^n , Lemma A.6 implies that $|\sum_{k \in C} \delta^k| = \sum_{k \in C} |\delta^k|$.

For each $k \in C$, let $\eta^k = x + \delta^k$. We want to show that $\eta^k \in \Delta$. To begin, observe that $\sum_j \eta_j^k = \sum_j x_j + \sum_j \delta_j^k = 1$. To check that $\eta_j^k \geq 0$ for all j , first note that if $j = k$, then $\eta_k^k = x_k + d_k = y_k \geq 0$. If $j \in C - \{k\}$, then $\eta_j^k = x_j \geq 0$. Finally, if $j \notin C$, then since d_k is negative, $\eta_j^k = x_j - \left(\frac{d_j}{\sum_{i \in C} d_i}\right) d_k \geq x_j \geq 0$.

For each $k \in C$, define $r_k: \Pi \rightarrow \mathbf{R}_+$ by

$$r_k(\pi) = \max \{r: \sigma(\pi) + r \delta^k \in \Delta\},$$

and define $z^k: \Pi \rightarrow \Delta$ by

$$z^k(\pi) = \sigma(\pi) + r_k(\pi) \delta^k.$$

Fix $\pi \in \Pi$; we want to show that $z_k^k(\pi) = 0$. Suppose to the contrary that $z_k^k(\pi) > 0$. Then since $z^k(\pi) \in \Delta$, $\sum_{j \neq k} z_j^k(\pi) < 1$, and so $z^k(\pi) \in \text{int}(\Delta)$; hence, $z^k(\pi) + \varepsilon \delta^k = \sigma(\pi) + (r_k(\pi) + \varepsilon) \delta^k \in \Delta$ for all small enough $\varepsilon > 0$, contradicting the definition of $r_k(\pi)$.

Next, we show that $Er_k \geq 1$. To see this, suppose to the contrary that $Er_k < 1$. Then $\eta_k^k = x_k + d_k < x_k + Er_k \delta_k^k = Ez_k^k = 0$, contradicting that $\eta^k \in \Delta$. Therefore, if we let $t_k = 1/Er_k$, then $t_k \in (0, 1]$.

Now define $\rho: \Sigma \rightarrow \Delta$ by

$$\rho(\pi) = \sigma(\pi) + \sum_{k \in C} t_k r_k(\pi) \delta^k.$$

To see that $\rho(\pi) \in \Delta$ for all $\pi \in \Pi$, observe that

$$\sum_j \rho_j(\pi) = \sum_j \sigma_j(\pi) + \sum_j \sum_{k \in C} t_k r_k(\pi) \delta_j^k = 1 + \sum_{k \in C} t_k r_k(\pi) \left(\sum_j \delta_j^k \right) = 1$$

and that $\rho_j(\pi) \leq \sigma_j(\pi)$ only if $j \in C$, in which case

$$\rho_j(\pi) = \sigma_j(\pi) + t_j r_j(\pi) \delta_j^j \geq \sigma_j(\pi) + r_j(\pi) \delta_j^j = z_j^j(\pi) = 0,$$

since $\delta_j^j < 0$. Moreover,

$$E\rho = E\sigma + E\left(\sum_{k \in C} t_k r_k \delta^k \right) = x + \sum_{k \in C} t_k \delta^k Er_k = x + \sum_{k \in C} \delta^k = x + d = y.$$

Finally, applying Lemma A.6 twice, we find that

$$\begin{aligned} \|\rho - \sigma\| &= \left\| \sum_{k \in C} t_k r_k \delta^k \right\| = E \left| \sum_{k \in C} t_k r_k \delta^k \right| = E \left(\sum_{k \in C} |t_k r_k \delta^k| \right) \\ &= \sum_{k \in C} |\delta^k| E(t_k r_k) = \sum_{k \in C} |\delta^k| = \left| \sum_{k \in C} \delta^k \right| = |d| = |y - x|. \blacksquare \end{aligned}$$

Proof of Lemma 6.3:

Let $\sigma_0 \in A$, and suppose that $\{\sigma_t\}$ leaves A in finite time. Since $\{\sigma_t\} \subset \hat{\Sigma}$ is continuous and since A is open, $\tau = \min\{t: \sigma_t \notin A\}$ exists, and $\rho \equiv \sigma_\tau$ lies on the boundary of A . To reach a contradiction, it is enough to show that $\{\sigma_t\}$ cannot reach

ρ in finite time.

The separation theorem for convex sets (Zeidler (1985, Proposition 39.4)) implies that there is a continuous linear functional $F: \hat{\Sigma} \rightarrow \mathbf{R}$ such that $F(\sigma) < F(\rho) \equiv r$ for all $\sigma \in A$. Therefore, to prove the lemma it is enough to show that if $\sigma_0 \in A$ and $F(\sigma_t + \dot{\sigma}_t) \leq r$ for all t , then $F(\sigma_t) < r$ for all t . Since F is continuous and linear, $\frac{d}{dt}F(\sigma_t) = F(\dot{\sigma}_t) \leq r - F(\sigma_t)$ (for details, see the proof of Lemma 5.3). Thus, $F(\sigma_t)$ will increase most quickly if we maximize $\frac{d}{dt}F(\sigma_t)$ by letting $\frac{d}{dt}F(\sigma_t) = r - F(\sigma_t)$ at all times t (which we can accomplish by setting $\dot{\sigma}_t \equiv \rho - \sigma_t$). In this case, $F(\sigma_t) = e^{-t}F(\sigma_0) + (1 - e^{-t})r$, which is less than r for all finite t . ■

Proof of Lemma 6.5

Let $\{\sigma_t\}$ be the solution to (B) from some $\sigma_0 \in \Sigma$ with $E\sigma_0 = x^* \in \Delta^*$, and let $\sigma^* = B(x^*)$. Since Theorem 5.2 implies that $\{E\sigma_t\}$ solves (AB), it follows from Proposition 5.1 that $E\sigma_t = x^*$ for all t . Hence, $B(E(\sigma_t)) = B(x^*) = \sigma^*$ for all t .

Since the solution to (B) from σ_0 is unique, it is enough to verify that $\sigma_t \equiv e^{-t}\sigma_0 + (1 - e^{-t})\sigma^*$ satisfies equation (B). And indeed,

$$\begin{aligned} \dot{\sigma}_t &= L^1 \lim_{\varepsilon \rightarrow 0} \left(\frac{\sigma_{t+\varepsilon} - \sigma_t}{\varepsilon} \right) \\ &= L^1 \lim_{\varepsilon \rightarrow 0} \left(\frac{1}{\varepsilon} (e^{-(t+\varepsilon)} - e^{-t})(\sigma_0 - \sigma^*) \right) \\ &= (\sigma_0 - \sigma^*) \lim_{\varepsilon \rightarrow 0} \left(\frac{e^{-(t+\varepsilon)} - e^{-t}}{\varepsilon} \right) \\ &= (\sigma_0 - \sigma^*) \frac{d}{dt} e^{-t} \\ &= (\sigma^* - \sigma_0) e^{-t} \\ &= \sigma^* - \sigma_t \\ &= B(E(\sigma_t)) - \sigma_t. \quad \blacksquare \end{aligned}$$

References

- Ellison, G., and D. Fudenberg (2000). "Learning Purified Mixed Equilibria," *J. Econ. Theory* 90, 84-115.
- Ely, J. C., and O. Yilankaya (2001). "Nash Equilibrium and the Evolution of Preferences." *J. Econ. Theory* 97, 255-272.
- Friedberg, S. H., A. J. Insel, and L. E. Spence (1989). *Linear Algebra*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall.

- Gilboa, I., and A. Matsui (1991). "Social Stability and Equilibrium," *Econometrica* 59, 859-867.
- Güth, W., and M. Yaari (1992). "Explaining Reciprocal Behavior in Simple Strategic Games: An Evolutionary Approach." In U. Witt, ed., *Explaining Process and Change: Approaches to Evolutionary Economics*, 23-34. Ann Arbor: University of Michigan Press.
- Hale, J. K. (1969). *Ordinary Differential Equations*. New York: Wiley.
- Harsanyi, J. C. (1973). "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *Int. J. Game Theory* 2, 1-23.
- Hirsch, M. W., and S. Smale (1974). *Differential Equations, Dynamical Systems, and Linear Algebra*. New York: Academic Press.
- Hofbauer, J. (1995). "Stability for the Best Response Dynamics," unpublished manuscript, Universität Wien.
- Hofbauer, J., and W. H. Sandholm (2002). "On the Global Convergence of Stochastic Fictitious Play," *Econometrica* 70, 2265-2294.
- Hofbauer, J., and W. H. Sandholm (2004). "Evolution in Games with Randomly Disturbed Payoffs," unpublished manuscript, University College London and University of Wisconsin.
- Karatzas, I., and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus*, 2nd ed. New York: Springer.
- Lang, S. (1997). *Undergraduate Analysis*, 2nd edition. New York: Springer.
- Matsui, A. (1992). "Best Response Dynamics and Socially Stable Strategies," *J. Econ. Theory* 57, 343-362.
- Oechssler, J., and F. Riedel (2001). "Evolutionary Dynamics on Infinite Strategy Spaces," *Econ. Theory* 17: 141-162.
- Sandholm, W. H. (2001). "Preference Evolution, Two-Speed Dynamics, and Rapid Social Change." *Rev. Econ. Dynamics* 4, 637-679.
- Sandholm, W. H. (2003). "Evolution in Bayesian Games II: Stability of Purified Equilibrium," unpublished manuscript, University of Wisconsin.
- Sandholm, W. H. (2004). "Negative Externalities and Evolutionary Implementation," forthcoming, *Rev. Econ. Stud.*
- Smirnov, G. V. (2002). *Introduction to the Theory of Differential Inclusions*. Providence, R.I.: American Mathematical Society.
- Zeidler, E. (1985). *Nonlinear Functional Analysis and its Applications III*. New York: Springer.