

Ely, Jeffery C.; Sandholm, William H.

Working Paper

Evolution with Diverse Preferences

Discussion Paper, No. 1317

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Ely, Jeffery C.; Sandholm, William H. (2001) : Evolution with Diverse Preferences, Discussion Paper, No. 1317, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221671>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Evolution with Diverse Preferences*

Jeffrey C. Ely
Department of Economics
Northwestern University
2003 Sheridan Road
Evanston, IL 60208
ely@nwu.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

First Version: March 16, 2000

This Version: April 10, 2001

* We thank Drew Fudenberg, Josef Hofbauer, Larry Samuelson, and Jörgen Weibull, as well as seminar audiences at the Stockholm School of Economics and at Wisconsin for their comments. Financial support from NSF Grants SBR-9810787 (Ely) and SES-0092145 (Sandholm) is gratefully acknowledged.

Abstract

We introduce best response dynamics for settings where players' 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 which are based on aggregation. Finally, we apply these techniques to prove a dynamic version of Harsanyi's (1973) purification theorem.

1. Introduction

We study best response dynamics for populations with diverse preferences. In this setting, a population's behavior is described by a Bayesian strategy: a map from preferences to distributions over actions. Our dynamics are defined on the space of Bayesian strategies; the rest points of the dynamics are precisely the Bayesian equilibria of the diverse preferences game. We prove the existence and uniqueness of solutions trajectories of these dynamics. We then characterize the dynamic stability of Bayesian equilibria using aggregate dynamics defined on the simplex, making it possible to determine the stability of Bayesian equilibria using standard dynamical systems techniques.

We offer three motivations for this study. First, we feel that in interactions involving large populations, one should not expect each individual to evaluate payoffs in precisely the same way. Therefore, in constructing evolutionary models, it seems more realistic to explicitly allow for diversity in preferences. We shall see that doing so eliminates pathological solution trajectories which 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. Briefly, these models address the natural selection of preferences in populations where preferences are diverse.¹ Selection of preferences is mediated through behavior, as the preferences which survive are those which induce the fittest behavior. By providing tools for analyzing behavior under diverse preferences, this paper provides the groundwork for studying competition between the preferences themselves.²

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 are hopeful that the insights of evolutionary game theory can be brought to bear more broadly in applied work.

¹ Examples include Güth and Yaari (1992), Ely and Yilankaya (1997), and Sandholm (1998).

² For further discussion of preference evolution, see Section 8.

We consider a population of individuals who are repeatedly randomly matched in pairwise interactions. Unlike in the standard evolutionary model, different individuals in our model evaluate payoffs using different payoff matrices. We assume that the subpopulation of players with any given payoff matrix is of negligible size relative to the population as a whole.

A complete description of behavior is given by a Bayesian strategy: a map which specifies the distribution of actions played in each subpopulation. The appropriate notion of equilibrium behavior is provided by 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 dynamics* for the common preference setting. Under their dynamics, the distribution of actions in a population always adjusts towards some best response to current behavior. To define our *Bayesian best response dynamics*, we require instead that the distribution of actions within each subpopulation adjust towards that subpopulation's current best response.

To complete the definition of the Bayesian dynamics, we must specify a notion of distance between Bayesian strategies.³ We find it convenient to use 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 dynamics is Lipschitz continuous under this norm. This enables us to prove that solutions to these dynamics 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 dynamics can originate from a single initial condition. This property can be the source of surprising solution trajectories: Hofbauer (1995) offers a game (presented below) in which solutions to the best response dynamics 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 dynamics are defined on the (L^1) space of Bayesian strategies, they are difficult to analyze directly. To contend with this, we define *aggregate best response dynamics* directly on the simplex. We show that there is a many-to-one mapping

³ By doing so, we provide an interpretation of the differential equation which defines the dynamics – see Section 2.2.

from solutions to the Bayesian dynamics to solutions to the aggregate dynamics. The map which accomplishes this is that which converts Bayesian strategies to the aggregate behavior which they induce. Thus, if we run the Bayesian dynamics from two strategies which yield the same aggregate behavior, the two solution trajectories will always yield the same aggregate behavior.

Were we only interested aggregate behavior, we could focus our attention entirely on these aggregate dynamics. But in most applications of Bayesian games, the full Bayesian strategy is itself of central 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 which 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 which are stable under the Bayesian dynamics and the distributions which are stable under the aggregate dynamics. Using this result, one can determine which Bayesian equilibria are stable under the original L^1 dynamics by considering much simpler dynamics defined on the simplex.⁴

Of course, these simpler aggregate dynamics are still a non-linear differential equation. To show that the characterization result above is of practical value, it must be demonstrated that this equation yields to analysis. Fortunately, Hofbauer and Sandholm (2001) are able to provide a detailed study of the aggregate dynamics derived from four classes of Bayesian games.⁵ We are hopeful that the aggregate dynamics will prove susceptible to analysis in other classes of games as well.

Our model of diverse preferences can be viewed as a large population version of Harsanyi's (1973) purification model. Harsanyi shows that each mixed equilibrium of nearly any normal form game can be approximated by a pure equilibrium of a Bayesian game; the Bayesian game is a version of the normal form game with slight perturbations in payoffs. Surprisingly, we are able to show that these perturbations

⁴ 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.

⁵ In particular, the Bayesian games are created from zero-sum games, games with an interior ESS, potential games, and supermodular games by adding idiosyncratic biases in the manner described in Section 2.2.

can always be chosen so that the pure equilibrium is stable under the Bayesian best response dynamics. Thus, not only can all mixed equilibria be purified; all can be purified in a dynamically robust fashion. Still, our purification result is weaker than Harsanyi's in one important sense: while Harsanyi proves the existence of a purified equilibrium under *any* payoff perturbation, we establish that under *some* of these perturbations, the purified equilibrium will be dynamically stable.⁶

Ellison and Fudenberg (2000) study population fictitious play under diverse preferences. In fictitious play, all players choose a best response to the time average of past play. Since this time average is the relevant state variable, fictitious play defines dynamics directly on the simplex even when preferences are diverse. In fact, it is easy to show that the dynamics studied by Ellison and Fudenberg (2000) are equivalent (after a time reparameterization) to our *aggregate* best response dynamics. Hence, our analysis shows that the connections between the best response dynamics and fictitious play, which are well known in the common preference setting,⁷ persist when preferences are diverse.

There are also close connections among our model, stochastic fictitious play (Fudenberg and Kreps (1993), Kaniovski and Young (1995), Benaïm and Hirsch (1999)), and certain models of stochastic evolution (Blume (1993, 1997), Young (1998)). In these stochastic models, players' payoffs are perturbed by random shocks which are i.i.d. over time. It can be shown that the aggregate dynamics in the present model describe the expected motion of the state variable in these stochastic models. For a complete treatment of these connections, we refer the reader to Hofbauer and Sandholm (2001).

Section 2 reviews the best response dynamics under common preferences and introduces the Bayesian best response dynamics. Section 3 establishes basic properties of the Bayesian dynamics. Sections 4, 5, and 6 provide aggregation results. Section 7 offers our dynamic version of Harsanyi's (1973) purification theorem. Finally, Section 8 relates our model to Ellison and Fudenberg's (2000) model of smoothed fictitious play and discusses directions for future research. All proofs omitted from the text can be found in the Appendix.

⁶ In related work, Binmore and Samuelson (2001) use static evolutionary stability concepts to study the tension between the instability of mixed equilibria when players may condition behavior on roles and the stabilizing effects of payoff perturbations.

⁷ See, e.g., Hofbauer (1995).

2. The Best Response Dynamics

A unit mass of players is repeatedly randomly matched to play a two player, symmetric, normal form game. Each matched player chooses one of n actions, which we identify with basis vectors in \mathbf{R}^n : $S = \{e_1, e_2, \dots, e_n\}$. 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 the standard environment, each player has the same preferences over action pairs. We represent these preferences by a matrix $\pi \in \Pi = \mathbf{R}^{n \times n}$, where π_{ij} is the payoff a player receives when he chooses action i and his opponent chooses action j . Players have expected utility preferences over mixtures.

We let $BR^\pi : \Delta \Rightarrow \Delta$ denote the best response correspondence for preference π :

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

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

The best response dynamics on Δ are defined by

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

The usual interpretation of these dynamics is that players occasionally consider switching actions, and that whenever a player does so he switches to a best response. The $-x$ term arises because at each moment in time, all players are equally likely to consider switching actions.⁸

For most payoff matrices π , there are action distributions x which admit multiple best responses, and hence many possible directions of motion under (BR). While it is still possible to prove that solutions to (BR) exist, uniqueness is not guaranteed.

It is clear that any rest point of (BR) must be a Nash equilibrium. Moreover, if the population begins at a Nash equilibrium x^* , players who switch to best responses can do so in proportions x^* , resulting in a stationary solution trajectory at x^* .

⁸ For an analysis of the foundations of evolutionary dynamics, see Sandholm (1999).

However, if the players who switch to a best response do so in proportions other than x^* , they may move away from the equilibrium.

As an illustration, we consider an example due to Hofbauer (1995). Suppose that all players' preferences are given by the game in Figure 1, whose symmetric Nash equilibria are $e_1 = (1, 0, 0)$, $x = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, and $y = (\frac{4}{5}, 0, \frac{1}{5})$. The phase diagram for the best response dynamics (BR) is presented in Figure 2. In the left region of the simplex, action 1 is the best response, so solution trajectories head directly towards the point e_1 . Similarly, the flow in the right region heads towards e_2 , and the flow in the triangle xyz heads towards e_3 . The intersections of these regions contain distributions which admit multiple best responses, and are therefore the possible sources of multiple solutions to (BR).

0, 0	6, -3	-4, -1
-3, 6	0, 0	5, 3
-1, -4	3, 5	0, 0

Figure 1

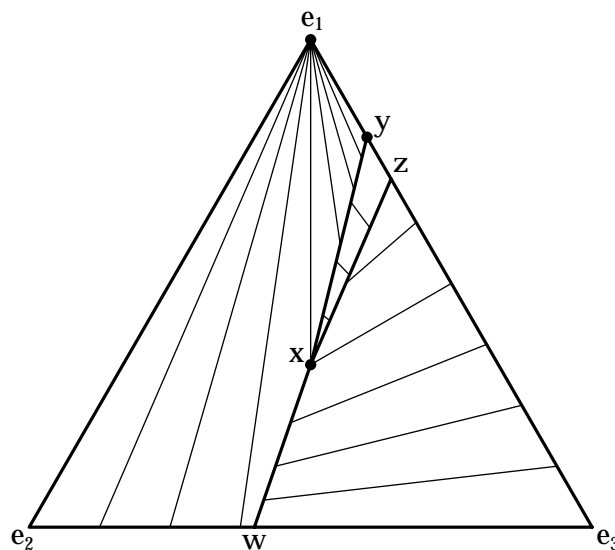


Figure 2

From the Nash equilibrium x , there are three possible courses of evolution: the population may stay put, move towards e_1 , or move towards y .⁹ From points in the interior of the segment xy , solutions may head towards e_1 , y , or e_3 . Hence, there are solution trajectories starting from the Nash equilibrium x which cycle indefinitely: they initially head towards y , but at some point break towards the segment xz ; after reaching xz , they return to x , possibly pausing there before beginning another circuit.

We show that the existence of trajectories which leave Nash equilibria is a consequence of the assumption that all players' 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 which admit multiple best responses. Indeed, Hofbauer's (1995) example is generic, in that all payoff matrices close to the one in Figure 1 yield qualitatively similar dynamics.

Our analysis shows that there is another sense in which Hofbauer's (1995) example is not generic. Our analysis relies on the following observation: if we fix a distribution over actions, the set of payoff matrices which generate indifference at that distribution is negligible. Therefore, in a population with diverse preferences, best responses are essentially unique, and hence the function which defines the best response dynamics in this context is single valued. To establish the uniqueness of solutions, and thus the equivalence of rest points and Bayesian Nash 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 matrices in the population is described by a probability measure μ on Π . In the language of Bayesian games, μ represents the distribution of types, which in the current context are simply the players' preferences. The common preferences model corresponds to the case in which μ places all mass on a single point in Π . However,

⁹ Matsui (1992) shows that there is a unique solution trajectory to (BR) from a Nash equilibrium precisely when the equilibrium is robust against equilibrium entrants (Swinkels (1992)). An interior equilibrium satisfies this property if and only if it is the unique (symmetric) Nash equilibrium of the game. However, robustness against equilibrium entrants does not imply local stability under (BR): see Section 6 of Matsui (1992) or Example 3.3 of Hofbauer (1995).

we will rule out such cases below, focusing instead on cases in which there is genuine heterogeneity in preferences.¹⁰

We suppose that there are a continuum of individuals with each preference $\pi \in \Pi$ in the support of μ . 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 of actions chosen in aggregate by the individuals of type π . Hence, each Bayesian strategy σ can be viewed as a random vector on the probability space (Π, μ) which 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 which indicate the same distribution for almost every type.

Let E denote expectation taken with respect to the probability measure μ . The proportion of players 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 will 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

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

Players are repeatedly paired with opponents who are randomly drawn from the population as a whole. Therefore, each player'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

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

¹⁰ We have assumed that there is just one population of players and that the players play a random matching game (so that payoffs are linear). These assumptions are made for simplicity; in fact, our results do not depend on either of these assumptions in a critical way.

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 which are best responses against aggregate behavior x for players with preference π .

We state our assumptions about the preference distribution μ in terms of the best response correspondence; we will describe classes of preference distributions which satisfy these assumptions below. Assumption (A1) requires that for all aggregate behaviors $x \in \Delta$, the set of players with multiple best responses has measure zero.

(A1) B is single valued.

As we shall see, Assumption (A1) holds as long as the preference distribution μ is sufficiently smooth; the assumption also allows for mass points on preferences which induce a strictly dominant action. Under Assumption (A1), all selections from $B(x)$ are equivalent, allowing us to regard $B: \Delta \rightarrow \Sigma$ as a function rather than 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 $\Sigma^* \subset \Sigma$ denote the set of Bayesian equilibria. Observe that under Assumption (A1), 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.¹¹

The *Bayesian best response dynamics* are defined 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 $\hat{\Sigma} = \{\sigma: \Pi \rightarrow \mathbf{R}^n\}$, which contains all possible directions of motion through Σ . Since Σ and $\hat{\Sigma}$ are function spaces, this law of motion must be interpreted with care.

Equation (B) specifies a direction of motion $\dot{\sigma}_t$ for the full Bayesian strategy σ_t , and hence a direction of motion for the distribution $\sigma_t(\pi)$ played by each preference $\pi \in \Pi$. The direction specified for preference π depends on $E(\sigma_t)$, the distribution played by the population as a whole. We therefore cannot consider the behavior of each preference π in isolation, but instead must regard changes in the entire

¹¹ Of course, this observation is originally due to Harsanyi (1973).

Bayesian strategy all at once. In other words, we must interpret equation (B) as a *functional* law of motion.

To do so, we must specify the norm which we use to measure distances between points in $\hat{\Sigma}$. To interpret equation (B) preference by preference, the natural norm to choose is the L^∞ norm, defined by

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

Unfortunately, under this norm equation (B) does not define a continuous law of motion. Even if two aggregate behaviors $x, y \in \Delta$ are very close together, if there is a non-null set of preferences whose best responses to x and y differ, then the best response to x and y are far apart in the L^∞ norm: $\|B(x) - B(y)\|_{L^\infty} = 2$. Therefore, the law of motion (B) is discontinuous, and the standard methods of establishing the existence and uniqueness of solution trajectories fail.

To create a tractable model, we need to use a norm on $\hat{\Sigma}$ which 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 would like any Bayesian strategies which differ only in the behavior of a small set of preferences to be regarded as close together. An appropriate choice of norm is the L^1 norm, which we denote $\|\cdot\|$:

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

Observe that if the best responses to x and y differ only on a set of measure ε , then $\|B(x) - B(y)\| = 2\varepsilon$.

In order to establish existence and uniqueness of solution trajectories to the Bayesian best response dynamics, we need to know that the dynamics are Lipschitz continuous. The following observation 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 dynamics is a consequence of the following assumption.

(A2) B is Lipschitz continuous (with respect to the L^1 norm).

Assumption (A2) 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 assumptions concerning the function B will hold as long as the preference distribution μ is both sufficiently diverse and sufficiently smooth. For example, a natural model is one of *general preferences*, in which all payoff matrices $\pi \in \Pi = \mathbf{R}^{n \times n}$ are possible. As long as the measure μ on $\mathbf{R}^{n \times n}$ is smooth and bounded, the general preferences model satisfies our assumptions.

Proposition 2.2: *In the general preferences model, if μ admits a bounded density function with compact support, then B satisfies Assumptions (A1) and (A2).*

A desirable feature of the general preferences model is that it makes no *a priori* restrictions on the set of preferences. On the other hand, our assumptions can also hold in situations in which preferences are less diverse. Consider the following model of *biases*. There is an "objective" payoff matrix $A \in \mathbf{R}^{n \times n}$ which is a component of all players' payoffs. However, each player has personal biases towards each action. These biases are represented by a vector $b \in \mathbf{R}^n$, where b_i is the extra payoff a player receives for playing action i . The preference matrix of a player with biases b is therefore $\pi = A + b\mathbf{1}^T$. If the measure ν on \mathbf{R}^n captures the distribution of biases, the measure μ on $\mathbf{R}^{n \times n}$ is defined by $\mu(C) = \nu(b: A + b\mathbf{1}^T \in C)$. As long as the bias distribution is sufficiently smooth, Assumptions (A1) and (A2) hold.

Proposition 2.3: *In the biases model, if ν admits a bounded density function $f: \mathbf{R}^n \rightarrow \mathbf{R}$ with either compact support or independent components (i.e., $f(b) = \prod_i f_i(b_i)$ for some density functions $f_i: \mathbf{R} \rightarrow \mathbf{R}$), then B satisfies Assumptions (A1) and (A2).*

3. Basic Properties of Solution Trajectories

We begin by establishing some basic properties of solutions to the Bayesian best response dynamics (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 $(\hat{\Sigma}, \|\cdot\|)$. For further details, see Lang (1983).

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 $\{\sigma_t\}$ at time t . As usual, the derivative describes the trajectory's direction of motion. However, for any individual preference π , the slope $\frac{1}{\varepsilon}(\sigma_{t+\varepsilon}(\pi) - \sigma_t(\pi)) \in \mathbf{R}^n$ of the line passing through $(t, \sigma_t(\pi))$ and $(t + \varepsilon, \sigma_{t+\varepsilon}(\pi))$ may not converge as ε approaches zero, and so the standard derivative $\frac{d}{dt}(\sigma_t(\pi))$ of the distribution trajectory $\{\sigma_t(\pi)\} \subset \mathbf{R}^n$ need not exist. For the L^1 derivative to exist, the measure of the set of preferences π for which the 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.

Theorem 3.1 sets out basic properties of solutions of the Bayesian dynamics.

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) Solution trajectories remain in Σ at all times $t \in [0, \infty)$.

(iii) 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 $f(\sigma) = B(E(\sigma)) - \sigma$.

(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 (up to sets of measure zero) of solutions to (B). Were Σ an open set, this would follow from standard results (e.g., Hirsch and Smale (1974, Sections 8.3 and 8.4)). However, the fact that Σ is closed introduces difficulties which are explained in detail and then resolved in the Appendix. Part (ii), establishing forward invariance of the space Σ under (B), is also proved there. Given these results, part (iii), continuity of solutions in their initial conditions, is standard.

If $\{\sigma_t\}$ is an L^1 solution to (B), then so is any trajectory $\{\hat{\sigma}_t\}$ which 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. The proof of this result, which relies on the Kolmogorov continuity theorem, is provided in the Appendix as well.

Clearly, only Bayesian equilibria can be rest points under (B). It is also evident that starting from every Bayesian equilibrium is a stationary solution trajectory. Since solutions are unique, these are the only trajectories starting from equilibria. Hence, part (v) of the theorem concludes that the Bayesian equilibria and the rest points of (B) are identical. This attractive property is a natural consequence of diversity in preferences.

4. Aggregation and Equilibrium

We have established that solution trajectories of the best response dynamics (B) exist and are unique. However, since these dynamics operate on an L^1 space, working with them directly is rather difficult. In the coming sections, we show that many important properties of the dynamics can be understood by analyzing aggregate dynamics. The aggregate dynamics operate directly on the simplex, and so can be studied using standard dynamical systems techniques – see Hofbauer and Sandholm (2001).

Before introducing the aggregate dynamics, 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 which 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 $\Delta^* \subset \Delta$ 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 Assumption (A2). ■

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 dynamics (B). To do so, we introduce the aggregate best response dynamics (AB), which are 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 dynamics operate on the complicated space Σ , the answers to many important questions about these dynamics can be obtained by applying standard tools to dynamics on the simplex.

5. Aggregation of Solution Trajectories

Under the dynamics (B), the Bayesian strategy σ_t always moves towards 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 dynamics* by

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

Under this law of motion, the distribution x_t moves towards the distribution induced by the best response to x_t . Some basic properties of these dynamics are noted in Theorem 5.1.

Theorem 5.1 (*Basic properties of solutions to (AB)*):

- (i) *Solutions to (AB) from every initial condition $x_0 \in \Delta$ exist and are unique.*
- (ii) *Solution trajectories remain in Δ at all times $t \in [0, \infty)$.*
- (iii) *The set of rest points of (AB) is Δ^* , the set of equilibrium distributions.*

The proof of parts (i) and (ii) can be found in the Appendix; part (iii) follows immediately from part (i) and the definition of an equilibrium distribution.

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 dynamics (D) *aggregate* to the dynamics (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 dynamics (B) aggregate to the aggregate best response dynamics (AB).

Theorem 5.2 tells us that the dynamics (AB) completely describe the evolution of aggregate behavior under the dynamics (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 $x_0 = E\sigma_0$. In other words, a sufficient statistic for aggregate behavior at time t under (B) is aggregate behavior at time 0. Bayesian strategies which 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: $E\sigma_0 = x^* \in \Delta^*$. Theorems 5.1 and 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 $\sigma \in E^{-1}(x^*)$ which aggregate to x^* but are not Bayesian equilibria, and hence are not rest points of the dynamics (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 which 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 $E(\dot{\sigma}_t) \equiv \frac{d}{dt}E\sigma_t$.*

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 dynamics (D) aggregate to the dynamics (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 $E \circ f \equiv g \circ E$, 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 $E \circ f \equiv g \circ E$. \blacksquare

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

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

aggregate to (AD) with $g(x) = E(F(x)) - x$. Thus, dynamics on Σ can be aggregated whenever the current "target point" $\sigma + \dot{\sigma} \in \Sigma$ is only a function of aggregate behavior. Indeed, the stability results in the next section extend immediately to all dynamics in this class.

6. Aggregation and Stability Analysis

We have established that the expectation operator E is both a one-to-one mapping between Bayesian equilibria and equilibrium distributions and many-to-one mapping from Bayesian trajectories which solve (B) to distribution trajectories which solve (AB). In this section, we prove that under three standard notions of stability, stability of Bayesian strategies under (B) is equivalent to stability of

distributions under (AB). Therefore, to characterize the stability of Bayesian strategies under the functional dynamics (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. Let Z be a subset of a Banach space $(\hat{Z}, \|\cdot\|)$, and let the function $h: Z \rightarrow \hat{Z}$ define dynamics on Z via the equation of motion

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

We suppose that Z is forward invariant under the dynamics (M), and let $z^* \in Z$ be a rest point of the dynamics: $h(z^*) = 0$. We say that z^* is *Lyapunov stable* under (M) if for each set $A \subset Z$ containing z^* which is open (relative to Z) there is an open set $A' \subset A$ which contains z^* such that any trajectory which 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 stable*.¹²

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 which is ε away from the distribution y , there is another Bayesian strategy ρ which is ε away from the σ and which aggregates to the distribution 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 dynamics.

¹² While all of our results are stated for single rest points, they can easily be extended to allow for closed, connected sets of rest points.

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 . Thus, $\sigma_t + \dot{\sigma}_t$ represents the point towards 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 towards 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 which 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 $\{\sigma_t\}$ always heads towards the point $B(E(\sigma_t)) \in V$, 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 which 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^* which is open relative to Σ , there is an set $U' \subset U$ containing σ^* which 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 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 $W \subset B^{-1}(V)$ which contains x^* and which is open relative to Δ such that solutions to (AB) which 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,*

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

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 dynamics (B), σ_t always heads directly towards 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 $\hat{\sigma}$ which is close to σ_t and which 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,

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

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 $\delta = \eta^{-K} \left(\frac{\eta}{2}\right)^{K+1}$. 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 $\tau_2 = \ln \frac{2\eta}{\eta}$, 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 (iii) 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 stability. The proof of this result is analogous to that of Theorem 6.4.

Theorem 6.6: (*Global stability*)

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

7. Purification

In any mixed equilibrium of a normal form game, each player is indifferent between his equilibrium strategy and all other strategies with the same support. This raises the question of why we should expect players to randomize in precisely the fashion which their equilibrium strategies dictate.

To address this issue, Harsanyi (1973) shows that almost every mixed

equilibrium of almost every normal form game can be viewed as a pure equilibrium of a Bayesian game created by slightly perturbing the payoffs of the normal form game. In fact, Harsanyi (1973) shows that these purified equilibria exist regardless of the distribution of payoff noises so long as the noises become small.

The Bayesian games which Harsanyi studies are formally equivalent to the games with diverse preferences studied here. The main difference is one of interpretation: while Harsanyi considers games played by a small group of players, each of whom has many possible payoff realizations, we consider a population of players in which the entire type distribution is realized at once. Just as in Harsanyi's model, every Bayesian equilibrium in our model has almost every preference playing a pure strategy. Thus, if all preferences in the population are close to a particular payoff matrix $\bar{\pi}$, we can think of the Bayesian equilibria of the diverse preferences game as purified equilibria of $\bar{\pi}$.

In a large population setting, it becomes natural to consider questions of stability: if behavior fluctuates slightly from a purified equilibrium, will the population return to the equilibrium? To see why this might or might not occur, consider single populations which play the games in Figures 3 and 4. The unique symmetric equilibrium of the Hawk-Dove game in Figure 3 is $(\frac{1}{2}, \frac{1}{2})$. Out of equilibrium, the less common strategy always has higher payoffs. Therefore, if the equilibrium is disturbed, we would expect the less common strategy to become more prevalent, and the population to return to the equilibrium state. In contrast, the coordination game in Figure 4 has both the $(\frac{1}{2}, \frac{1}{2})$ equilibrium and two symmetric pure equilibria. In this game, the more common strategy always has higher payoffs, so we would expect a population which starts near but not at $(\frac{1}{2}, \frac{1}{2})$ to move towards a pure equilibrium. If diverse preferences are introduced, it seems natural to expect that in the Hawk-Dove game, purified versions of the $(\frac{1}{2}, \frac{1}{2})$ equilibrium will be stable, while in the coordination game, purified versions of the $(\frac{1}{2}, \frac{1}{2})$ equilibrium will be unstable.

0, 0	1, 1
1, 1	0, 0

Figure 3

1, 1	0, 0
0, 0	1, 1

Figure 4

In fact, one can approximate *any* equilibrium of a normal form game with a stable purified equilibrium.

Theorem 7.1: (*Stable purification*)

Let x^* be a symmetric equilibrium of the game $\bar{\pi}$. For all $\varepsilon > 0$, there exists a measure μ on $\mathbf{R}^{n \times n}$ such that

- (i) $\mu\left\{\pi : \max_{i,j} |\pi_{ij} - \bar{\pi}_{ij}| \leq \varepsilon\right\} = 1$
- (ii) There is a Bayesian equilibrium σ^* satisfying $E\sigma^* = x^*$ which is locally stable under (B).

The intuition behind this result is simple. Since x^* is an equilibrium of $\bar{\pi}$, we can find slight perturbations μ of the Dirac measure $\delta_{\bar{\pi}}$ with the property that each strategy i is a strict best response for fraction x_i^* of the population when the strategy distribution is near x^* . This measure μ clearly admits a Bayesian equilibrium σ^* with distribution $E\sigma^* = x^*$. Moreover, if the Bayesian strategy σ^* is slightly disturbed, the resulting strategy distribution will be very close to x^* ; hence, no player's best response changes, and the profile returns to x^* .

The proof of the theorem relies on our aggregation results from Section 6.

Proof: Our construction of μ is based on the model of biases from Section 2.2. For each strategy i , define the set $F_i \subset \mathbf{R}^n$ as follows:

$$F_i = \{b \in \mathbf{R}^n : b_i \in [\frac{\varepsilon}{2}, \varepsilon] \text{ and } b_j \in [-\varepsilon, -\frac{\varepsilon}{2}] \text{ for all } j \neq i\}.$$

Each vector $b \in F_i$ represents biases which favor strategy i by at least ε (but no more than 2ε). Let ν be a measure on \mathbf{R}^n which admits a bounded density function and which satisfies $\nu(F_i) = x_i^*$ for all strategies i . The measure μ on $\mathbf{R}^{n \times n}$ is then defined by $\mu(C) = \nu(b : \bar{\pi} + b\mathbf{1}^T \in C)$. Proposition 2.3 implies that the best response correspondence $B : \Delta \rightarrow \Sigma$ induced by μ satisfies Assumptions (A1) and (A2).

Since x^* is a Nash equilibrium of $\bar{\pi}$, it follows that for all strategies i and $j \neq i$,

$$e_i \cdot \bar{\pi} x^* \geq e_j \cdot \bar{\pi} x^*.$$

Consequently,

$$e_i \cdot \bar{\pi} x^* + b_i \geq e_j \cdot \bar{\pi} x^* + b_j + \varepsilon \quad \text{for all } b \in F_i.$$

Thus, letting $\delta \equiv (2 \max_{k,l} |\bar{\pi}_{kl}|)^{-1}$, we see that

$$e_i \cdot \bar{\pi} x + b_i > e_j \cdot \bar{\pi} x + b_j \quad \text{for all } b \in F_i \text{ and all } x \text{ satisfying } |x - x^*| < \delta. \quad (1)$$

Let $\hat{F}_i = \{\pi : \pi = \bar{\pi} + b \mathbf{1}^T \text{ for some } b \in F_i\}$. Then if $\pi \in \hat{F}_i$ and x is within δ of x^* , equation (1) implies that $e_i \cdot \pi x > e_j \cdot \pi x$ for all $j \neq i$, and hence that $B(x)(\pi) = e_i$. Therefore, taking expectations, and observing that $\mu(\hat{F}_i) = \nu(F_i) = x_i^*$, we see that

$$E(B(x)) = x^* \text{ for all } x \text{ satisfying } |x - x^*| < \delta. \quad (2)$$

Now consider the Bayesian strategy $\sigma^* = B(x^*)$. Evaluating equation (2) at x^* , we find that $E(\sigma^*) = x^*$. Applying B to both sides of this equality yields $B(E(\sigma^*)) = \sigma^*$. Thus, σ^* is a Bayesian equilibrium under μ .

To show that σ^* is locally stable under (B), it is enough (by Theorem 6.4) to check the local stability of x^* under (AB). To do so, we observe that the derivative of $E(B(x)) - x$ evaluated at x^* is given by $D[E(B(x)) - x]|_{x=x^*} = D[E(B(x))]|_{x=x^*} - I \in \mathbf{R}^{n \times n}$. However, equation (2) shows that $E(B(x))$ is constant in a neighborhood of x^* , and so $D[E(B(x))]|_{x=x^*}$ is the null matrix. Since all eigenvalues of $-I$ equal -1 , we conclude that x^* is stable under (AB), and thus that σ^* is stable under (B). ■

To illustrate Theorem 7.1 and its proof, we describe how the mixed equilibrium of the coordination game in Figure 4 can be purified in a stable fashion. Since this game has only two actions, aggregate behavior is described by the scalar x_1 , which represents the proportion of players choosing strategy 1.

Figure 5 sketches the best response correspondence for the coordination game, as well as a phase diagram for the best response dynamics. When $BR(x_1) > x_1$, the best response dynamics increase the proportion of players choosing x_1 , leading to the pure Nash equilibrium $x_1 = 1$; when $BR(x_1) < x_1$, the proportion choosing x_1 declines, moving towards the equilibrium $x_1 = 0$. Clearly, the mixed equilibrium $x_1 = \frac{1}{2}$ is unstable.

To purify this equilibrium, we suppose that half of the population is slightly biased in favor of strategy 2, and that the other half is slightly biased in favor of strategy 1. In particular, we let the distribution of the bias difference $b_1 - b_2$ be uniformly distributed on the set $[-2\varepsilon, -\varepsilon] \cup [\varepsilon, 2\varepsilon]$. The resulting aggregate best response function is pictured in Figure 6, along with the corresponding phase

diagram for (AB). It is clear from the figure that $x_1 = \frac{1}{2}$ is an equilibrium distribution which is stable under (AB). Thus, Theorem 6.4 implies that the Bayesian equilibrium $B(x_1)$ is stable under (B). At this equilibrium, no player is within ε of indifference, and for this reason the equilibrium is robust to any small change in behavior.

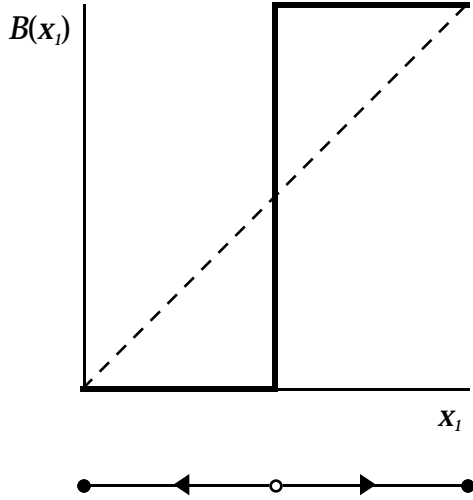


Figure 5

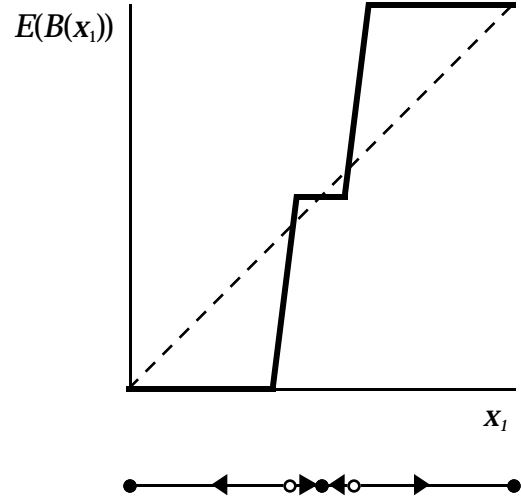


Figure 6

If we lower the magnitude ε of the largest payoff perturbation, the aggregate best response function $E(B(\cdot))$ approaches the original best response function $B(\cdot)$. Consequently, the size of the payoff perturbations which are introduced determines the size of the basin of attraction of the purified equilibrium. This observation qualifies our stability result: if the payoff perturbations used to construct μ are very small, then correspondingly small behavior disturbances will upset the purified equilibrium.

Harsanyi establishes the existence of purified equilibria under any distribution of disturbances. In contrast, our stable purification result only holds for certain choices of this distribution. However, the construction we have used so far, which ensures that no player is too close to indifference at the purified equilibrium, is more cautious than necessary. Examining Figure 6, we see that two properties are enough to ensure the stable purification of a mixed equilibrium. First, the mass of players biased towards each strategy must be close to the number playing that strategy in the mixed equilibrium. This creates a purified equilibrium near the original mixed equilibrium. Second, to guarantee stability, the density of preferences near this

purified equilibrium must not be too large. This ensures that there are not too many players who are too close to indifference. It follows that the aggregate best response function $E(B(\cdot))$ is relatively flat at the equilibrium, and so that the derivative matrix $D[E(B(x)) - x] \big|_{x=x^*}$ is stable.

Theorem 7.1 might create the impression that introducing payoff perturbations enables one to reverse the stability of any mixed equilibrium. In fact, while we can always create stable equilibria from unstable ones, we cannot always do the opposite. We illustrate this in Figures 7 and 8, which sketch the best response correspondence and an aggregate best response function for the Hawk-Dove game from Figure 3. The best response correspondence intersects the 45° line only at $x_1 = \frac{1}{2}$, which is the unique Nash equilibrium and the global attractor of the best response dynamics. If we consider *any* distribution μ consisting of preferences close to the payoff matrix of the Hawk-Dove game, the resulting aggregate best response function is a decreasing function close the original best response function. The purified equilibrium under μ must therefore be unique and globally stable.

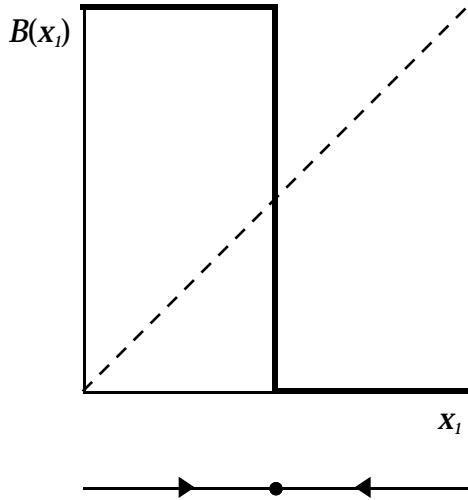


Figure 7

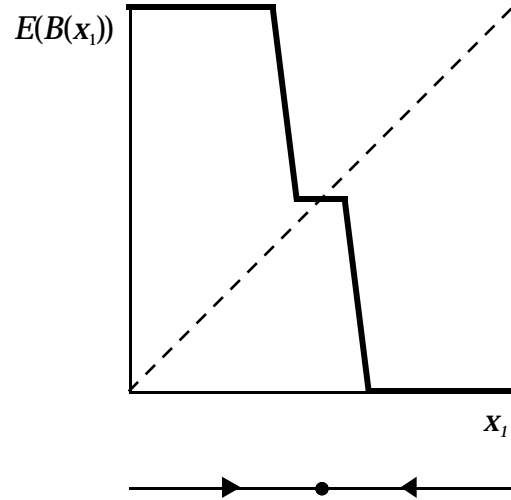


Figure 8

In fact, Hofbauer and Sandholm (2001) are able to extend this reasoning to $n \times n$ games. By combining an analysis of the aggregate best response dynamics with our Theorem 6.6, they establish the following result.

Theorem 7.2 (Hofbauer and Sandholm (2001)):

Suppose that the game $\bar{\pi}$ admits an interior ESS, and let μ be defined by $\mu(C) = v(b: \bar{\pi} + b\mathbf{1}^T \in C)$, where the bias distribution v is sufficiently smooth and has full support on \mathbf{R}^n . Then the Bayesian game defined by μ has a unique Bayesian equilibrium, which is globally stable under (B).

Theorem 7.2 shows that any perturbed version of a game with an interior ESS must have a unique and globally stable Bayesian equilibrium. Perturbations are required to take the form of biases towards each strategy, but no restriction on the size of these perturbations is required. Of course, the purified equilibrium need only closely approximate the original ESS if the perturbations are "small".

To summarize: while every mixed equilibrium admits a stable purification, there are many mixed equilibria which do not admit an unstable purification.

8. Concluding Remarks

8.1 The Best Response Dynamics and Fictitious Play

Under common preferences, the close connections between the best response dynamics 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:

$$\begin{aligned} \text{(FP)} \quad \dot{a}_t &= \frac{1}{t} c_t - \frac{1}{t^2} \int_0^t c_s ds \\ &= \frac{1}{t} (BR^\pi(a_t) - a_t) \end{aligned}$$

Since players always choose a best response to the time average of past play, this time average always moves in the direction of the best response, moving more slowly as time passes.

Notice that the expression $\frac{1}{t}$ only affects the speed of the dynamics, not the direction of motion. Therefore, after a reparameterization of time, the evolution of beliefs under (FP) is identical to the evolution of behavior under the best response

dynamics (BR).

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

Ellison and Fudenberg (2000) study fictitious play under 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 which 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

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

which is a reparameterization of our aggregate best response dynamics (AB):

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

In fictitious play with diverse preferences, the state variable is the average *distribution* of past behavior, $a_t \in \Delta$. If one keeps track of this, one can always determine the best responses $B(a_t) \in \Sigma$ and the best response distribution $E(B(a_t)) \in \Delta$. The latter object determines the direction in which the time average evolves. Since one can determine the direction of motion without ever explicitly considering a Bayesian strategy, the dynamics can be defined directly on the simplex Δ .

In contrast, under the Bayesian best response dynamics

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

the relevant state variable is current Bayesian strategy σ_t . Thus, these dynamics are defined on the L^1 space Σ , and so specify directly how behavior in every subpopulation evolves. Our results establish that equilibrium and stability analyses can be performed directly in terms of the aggregate dynamics (AB). Thus, these results show that close connections between fictitious play and the best response dynamics persist when preferences are diverse.¹³

¹³ While the dynamics (F) and (AB) are essentially the same, the evolution of Bayesian strategies under population fictitious play and under the Bayesian best response dynamics are quite different. Suppose that (F) and (AB) are currently at the 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.

8.2 Preference Evolution

The techniques developed in this paper provide foundations for models of preference evolution. Diversity in preferences typically results in diversity in behavior, even after equilibrium is reached. If we take a long term view, this variation in behavior can be a source of differences in reproductive success among the preferences themselves. Preferences which induce "evolutionarily fit" behavior prosper, to the detriment of the others.

A full account of the process of preference evolution requires dynamics which run at two speeds at once: as players quickly adjust their behavior to accord with their preferences, the preference distribution gradually evolves in response to differences in fitness. Sandholm (1998) studies such two-speed dynamics, but the analysis there is restricted to two strategy games, with diversity in preferences limited to biases in favor one of the two strategies. To understand preference evolution in more general settings, one first needs a general model of behavior adjustment under fixed preferences. The present paper provides this model. Incorporating the techniques developed here into a general theory of preference evolution is a topic for future research.

Appendix

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

The aggregate best response dynamics (AB) are defined on the closed set of strategy distributions Δ . However, the basic results concerning the existence and uniqueness of solutions to differential equations concern equations defined on open sets. In this appendix, we show how these standard results can be used to study the aggregate best response dynamics, in particular establishing the existence, uniqueness, and forward invariance of solutions in the set Δ . We then go on to prove analogous results for L^1 dynamics on the set Σ , thereby establishing the basic properties of the Bayesian best response dynamics (B).

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

Under the best response dynamics, we only know that the Bayesian strategy σ_i is an element of $E^{-1}(a_i)$. These two strategies only coincide if a_i is an equilibrium distribution and $\sigma_i = B(a_i)$ is the corresponding Bayesian equilibrium.

- (LC) g is Lipschitz continuous on Δ .
- (FI 1) $\sum_i g_i(x) = 0$ for all $x \in \Delta$.
- (FI 2) For all $x \in \Delta$, $g_i(x) \geq 0$ whenever $x_i = 0$.

Condition (LC) is the usual Lipschitz continuity condition used to prove the existence of unique solution trajectories to the differential equation $\dot{x} = g(x)$. Condition (FI 1) says that $\sum_i \dot{x}_i = 0$, implying that the affine space $\tilde{\Delta} = \{x \in \mathbf{R}^n: \sum_i x_i = 1\}$ is invariant under the differential equation. Condition (FI 2) says that whenever the component x_i equals zero, its rate of change is non-negative.

Theorem A.1 explains the precise sense in which these conditions imply the forward invariance of the simplex under $\dot{x} = g(x)$. This result implies parts (i) and (ii) of Theorem 5.1.

Theorem A.1: *Let $g: \Delta \rightarrow \mathbf{R}^n$ satisfy (LC), (FI 1), and (FI 2). Then under any Lipschitz continuous extension \tilde{g} of g from Δ to $\tilde{\Delta}$, the solution to $\dot{x} = \tilde{g}(x)$ from each $x_0 \in \Delta$ exists, is unique, and remains in Δ at all times $t \in [0, \infty)$.*

The proof relies on the following lemma. Let $|x|_E = \sqrt{\sum_i x_i^2}$ denote the Euclidean norm on \mathbf{R}^n .

Lemma 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 for some $k > 0$, $|c(x) - c(y)|_E \leq |x - y|_E$ and $|c(x) - c(y)| \leq k|x - y|$ for all $x, y \in \mathbf{R}^n$.

Proof: Fix $x, y \in \mathbf{R}^n$. If $c(x) = c(y)$ there is nothing to prove, so we assume that $c(x) \neq c(y)$. If we let $p = c(y) - c(x)$, it follows immediately that $p \cdot c(y) > p \cdot c(x)$. Let $I = \{z \in \mathbf{R}^n: p \cdot c(y) > p \cdot z > p \cdot c(x)\}$. We first show that $x \notin I$. First, observe that since C is convex, $c(x) + \lambda p = \lambda c(y) + (1 - \lambda)c(x) \in C$ for all $\lambda \in [0, 1]$. Therefore, if $w \in I$, it follows that $w = c(x) + \lambda p + v$ for some $\lambda \in (0, 1)$ and some $v \in \mathbf{R}^n$ orthogonal to p . By the Pythagorean theorem,

$$|w - c(x)|_E = |\lambda p + v|_E > |v|_E = |w - (c(x) + \lambda p)|_E$$

Therefore, w is closer to $c(x) + \lambda p \in C$ than to $c(x)$, which implies that $x \neq w$. Similarly, $y \notin I$; analogous arguments establish that $p \cdot x \leq p \cdot c(x)$ and that $p \cdot y \geq p \cdot c(y)$. We can therefore find a $\gamma \geq 1$ and a $v \in \mathbf{R}^n$ orthogonal to p such that $y = x + \gamma p + v$. Thus, the Pythagorean theorem implies that

$$\|x - y\|_E = \|\gamma p + v\|_E \geq \|p\|_E = \|c(x) - c(y)\|_E.$$

The inequality for the summation norm follows from the equivalence of norms on \mathbf{R}^n . ■

Proof of Theorem A.1: Define $\hat{g}: \tilde{\Delta} \rightarrow \mathbf{R}^n$ by $\hat{g}(x) = g(c(x))$. If K is the Lipschitz constant for g , then Lemma A.2 implies that for all x and y in $\tilde{\Delta}$,

$$\begin{aligned} \|\hat{g}(x) - \hat{g}(y)\| &= \|g(c(x)) - g(c(y))\| \\ &\leq K \|c(x) - c(y)\| \\ &\leq K k \|x - y\|. \end{aligned}$$

Hence, \hat{g} is Lipschitz. Moreover, one can check that if $x \in \tilde{\Delta}$ and $x_i \leq 0$, then $c_i(x) = 0$. Thus, condition (FI 2) implies that $\hat{g}_i(x) \geq 0$ whenever $x_i \leq 0$.

Standard results (e.g., Hirsch and Smale (1974, Sections 8.3 and 8.4)) imply that there is a unique solution to $\dot{x} = \hat{g}(x)$ from each initial condition in $\tilde{\Delta}$. Suppose that $\{x_t\}$ is the solution to this equation starting from $x_0 \in \Delta$, and suppose that $[x_u]_i < 0$. Then since $\{[x_t]_i\}$ is continuous in t , we can find a time $s \in [0, u)$ such that $[x_s]_i = 0$ and $[x_t]_i < 0$ for all $t \in (s, u)$. Consequently,

$$[x_u]_i = [x_s]_i + \int_s^u [\dot{x}_t]_i dt = \int_s^u \hat{g}_i(x_t) dt \geq 0,$$

which is a contradiction. Therefore, Δ is forward invariant under \hat{g} .

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 .
- (FI 1') $\sum_i f_i(\sigma)(\pi) = 0$ for all $\sigma \in \Sigma$ and $\pi \in \Pi$.
- (FI 2') 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 (FI 1') ensures that solutions stay in the affine space $\tilde{\Sigma} = \{\sigma \in \hat{\Sigma}: \sigma(\pi) \in \tilde{\Delta} \text{ for all } \pi \in \Pi\}$; condition (FI 2') 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, and the forward invariance of Σ for L^1 solutions to $\dot{\sigma} = f(\sigma)$ are established in Theorem A.3. This result implies parts (i) and (ii) of Theorem 3.1.

Theorem A.3: *Let $f: \Sigma \rightarrow \mathbf{R}^n$ satisfy (LC'), (FI 1'), (FI 2'), and (UB). Then under any Lipschitz continuous extension \tilde{f} of f from Σ to $\tilde{\Sigma}$, the solution to $\dot{\sigma} = \tilde{f}(\sigma)$ from each $\sigma_0 \in \Sigma$ exists, is unique, and remains in Σ at all times $t \in [0, \infty)$.*

In addition to these properties, we would also like to establish that some L^1 solution $\{\sigma_t\}$ has continuous sample paths: i.e., that $\{\sigma_t(\pi)\}$ is continuous for each preference $\pi \in \Pi$. Put differently, we would like to know that the behavior of the subpopulation with preference π changes continuously over time. While not every L^1 solution will have this property, we can prove that there is always one which 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 Lipschitz continuous and pointwise bounded. Then there exists a modification $\{\tilde{\sigma}_t\}$ of $\{\sigma_t\}$ with continuous sample paths: i.e., such that $\mu(\pi: \tilde{\sigma}_t(\pi) \text{ is continuous in } t) = 1$.*

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.

The difficulty in extending the proof of Theorem A.1 to the current setting is that if a trajectory $\{\sigma_t\}$ starts in Σ but eventually leaves, the behavior trajectories $\{\sigma_t(\pi)\}$ for different preferences π may leave Δ at different times. In fact, if we only know that $\{\sigma_t\}$ is an L^1 solution, we do not know enough about the sample paths $\{\sigma_t(\pi)\}$ to specify the time at which the path leaves Δ , or even to say whether the path "leaves" Δ at all. The first difficulty is handled by introducing a random time, and the second by appealing to Theorem A.4.

The proof of Theorem A.3 requires us to introduce the notion of L^1 integrals of trajectories through $\hat{\Sigma}$; for a complete treatment, see Lang (1983). If $\{\sigma_t\}$ is an L^1 continuous trajectory through $\hat{\Sigma}$, L^1 integrals over this trajectory, denoted $\int_a^b \sigma_t dt$, are the L^1 limits of Riemann sums of step functions $\{\sigma_t^n\}$ which approximate the trajectory $\{\sigma_t\}$ arbitrarily well in the L^1 norm. If $\{\sigma_t\}$ is an L^1 solution to $\dot{\sigma} = \tilde{f}(\sigma)$, we have by definition 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 Riemann sum 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 $\sigma, \rho \in \tilde{\Sigma}$,

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

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 $\sigma_0 \in \tilde{\Sigma}$.

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 (FI 2'). 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 turn to the proof of Theorem A.4. To do so, we need to introduce the L^2 norm on $\hat{\Sigma}$:

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

If $\{\sigma_t\}$ is L^2 continuous, then the L^2 integral, denoted $\oint_a^b \sigma_t dt$, is defined by taking approximating Riemann sums. The following standard inequality holds for the L^2 integral:

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

If $\{\rho_t\}$ is pointwise bounded (i.e., if $|(\rho_t(\pi))_i| \leq M$ for all i, t , and π), then since μ is a probability measure, $L^1 \lim_{s \rightarrow t} \rho_t = L^2 \lim_{s \rightarrow t} \rho_t$ if either limit exists. If $\{\sigma_t\}$ is an L^1 continuous trajectory which is pointwise bounded, then the Riemann sums which approximate these trajectories over any compact interval $[a, b]$ are pointwise bounded; therefore, the L^1 and L^2 integrals of the trajectory over this interval are equal: $\int_a^b \sigma_t dt = \oint_a^b \sigma_t dt$. Finally, a trajectory $\{\sigma_t\} \subset \hat{\Sigma}$ is L^2 Lipschitz continuous if there exists a constant K such that

$$\|\sigma_t - \sigma_s\|_{L^2} \leq K|t - s|$$

The proof of Theorem A.4 relies on the Kolmogorov continuity theorem

(Karatzas and Shreve (1991, Theorem 2.2.8 and Corollary 2.2.11)). Lemma A.5 is an implication of this result.

Lemma A.5: *Suppose that $\{\sigma_t\}$ is L^2 Lipschitz continuous. 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:

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,

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

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 (A1), which requires that B is single valued, obviously holds, so we focus on the Lipschitz continuity condition (A2). 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 preferences which choose i over j at x and j over i at y is

$$\begin{aligned} \Pi_{ij} &= \{\pi. (\pi x)_i > (\pi x)_j \text{ and } (\pi y)_i < (\pi y)_j\} \\ &= \{\pi. (\pi_i - \pi_j) \cdot x > 0 > (\pi_i - \pi_j) \cdot y\}. \end{aligned}$$

We can associate with each preference π a difference vector $d_{ij} = \pi_i - \pi_j \in \mathbf{R}^n$. Let

$f: \mathbf{R}^{n \times n} \rightarrow \mathbf{R}$ denote the density function of the measure μ , and let $g_{ij}: \mathbf{R}^n \rightarrow \mathbf{R}$ be the density of the measure on the difference d_{ij} which is induced by μ . 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 λ represent Lebesgue measure on \mathbf{R}^n . Suppose we can show that

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

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

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

To bound $\lambda(D_{ij})$, 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

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

$$= \{d' \in S: d' = Td \text{ for some } d \in D_{ij}\}$$

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

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 $\lambda(S) < (2r)^n$,

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

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 which are ε away from the vector $(c, 0)$. The largest possible angle between the vector $(1, 0)$ and a vector in Z_ε is

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

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

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

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 (A1) clearly holds, so we need only consider the Lipschitz continuity condition (A2). Fix $x, y \in \Delta$ and $i \neq j$. Let $\Pi_{ij} \subset \Pi$ represent the set of preferences which prefer strategy i to strategy j at distribution x but prefer j to i at y :

$$\Pi_{ij} = \{\pi. (\pi x)_i > (\pi x)_j \text{ and } (\pi y)_i < (\pi y)_j\}$$

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

$$\begin{aligned} D_{ij} &= \{b: (Ax + b)_i > (Ax + b)_j \text{ and } (Ay + b)_i < (Ay + b)_j\} \\ &= \{b: (A_i - A_j) \cdot x > b_j - b_i > (A_i - A_j) \cdot y\}. \end{aligned}$$

Here, A_i and A_j are rows of A . Now suppose we can show that $\nu(D_{ij}) \leq K|x - y|$ for some K which 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 bias $b \in \mathbf{R}^n$ is associated with a single value of $b_j - b_i \in \mathbf{R}$. Let $f: \mathbf{R}^n \rightarrow \mathbf{R}$ denote the density of the measure ν , and let $g_{ji}: \mathbf{R} \rightarrow \mathbf{R}$ denote the density of the measure for the difference $b_j - b_i$ which is induced by ν . If $[-c, c]^n$ 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

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

On the other hand, if $f(b) = \prod_i f_i(b_i)$, then since f is bounded, there is also a constant \hat{M} which bounds all of the functions f_i . Therefore, a convolution yields

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

In either case, $g_{ji} \leq J$ for some constant J .

The interval of values of $b_j - b_i$ which lie in the set D_{ij} has length

$$(A_i - A_j) \cdot x - (A_i - A_j) \cdot y = (A_i - A_j) \cdot (x - y) \leq \sum_k 2\bar{A} |x_k - y_k| = 2\bar{A} |x - y|,$$

where $\bar{A} = \max_{i,j} |A_{ij}|$. Therefore,

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

and we conclude that

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

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,

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

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

- Benaïm M., and M. W. Hirsch (1999). "Mixed Equilibria and Dynamical Systems Arising from Repeated Games," *Games Econ. Behav.* **29**, 36-72.
- Binmore, K., and L. Samuelson (2001). "Evolution and Mixed Strategies," *Games Econ. Behav.* **34**, 200-226.
- Blume, L. E. (1993). "The Statistical Mechanics of Strategic Interaction," *Games Econ. Behav.* **4**, 387-424.
- Blume, L. E. (1997). "Population games." In *The Economy as an Evolving Complex System II*, W. B. Arthur, S. N. Durlauf, and D. A. Lane, Eds. Reading, MA: Addison-Wesley.
- Ellison, G., and D. Fudenberg (2000). "Learning Purified Mixed Equilibria," *J. Econ. Theory* **90**, 84-115.

- Ely, J. C., and O. Yilankaya (1997). "Nash Equilibrium and the Evolution of Preferences." Forthcoming, *J. Econ. Theory*.
- Friedberg, S. H., A. J. Insel, and L. E. Spence (1989). *Linear Algebra*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall.
- Fudenberg D., and D. M. Kreps (1993). "Learning Mixed Equilibria," *Games Econ. Behav.* 5 (1993), 320-367.
- 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 Interscience.
- Harsanyi, J. C. (1973). "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *International Journal of 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," mimeo, Universität Wien.
- Hofbauer, J., and W. H. Sandholm (2001). "Evolution and Learning in Games with Randomly Disturbed Payoffs," mimeo, Universität Wien and University of Wisconsin.
- Kaniovski, Y. M., and H. P. Young (1995). "Learning Dynamics in Games with Stochastic Perturbations," *Games Econ. Behav.* 11, 330-363.
- Karatzas, I., and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus*, 2nd ed. New York: Springer.
- Lang, S. (1983). *Undergraduate Analysis*. New York: Springer.
- Matsui, A. (1992). "Best Response Dynamics and Socially Stable Strategies," *J. Econ. Theory* 57, 343-362.
- Sandholm, W. H. (1998). "Preference Evolution, Two-Speed Dynamics, and Rapid Social Change." Forthcoming, *Review of Economic Dynamics*.
- Sandholm, W. H. (1999). "Evolution and Equilibrium under Inexact Information," mimeo, University of Wisconsin.
- Swinkels, J. M. (1992). "Evolutionary Stability with Equilibrium Entrants," *J. Econ. Theory* 57, 306-332.

- Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton: Princeton University Press.
- Zeidler, E. (1985). *Nonlinear Functional Analysis and its Applications III*. New York: Springer.