Orders of Limits for Stationary Distributions, Stochastic Dominance, and Stochastic Stability*

William H. Sandholm†

July 1, 2009

Abstract

A population of agents recurrently plays a two-strategy population game. When an agent receives a revision opportunity, he chooses a new strategy using a noisy best response rule that satisfies mild regularity conditions; best response with mutations, logit choice, and probit choice are all permitted. We study the long run behavior of the resulting Markov process when the noise level η is small and the population size N is large. We obtain a precise characterization of the asymptotics of the stationary distributions $\mu^{N,\eta}$ as η approaches zero and N approaches infinity, and we establish that these asymptotics are the same for either order of limits and for all simultaneous limits.

In general, different noisy best response rules can generate different stochastically stable states. To obtain a robust selection result, we introduce a refinement of risk dominance called *stochastic dominance*, and we prove that coordination on a given strategy is stochastically stable under every noisy best response rule if and only if that strategy is stochastically dominant.

1. Introduction

Stochastic stability analysis provides unique predictions of long run behavior in games played by agents who employ simple, myopic choice rules. Since the early work of Foster and Young (1990), Kandori et al. (1993), and Young (1993), this approach to modeling recurring interactions has burgeoned both in abstract strategic environments and in concrete economic applications.[1]

*I thank Jeff Ely, Drew Fudenberg, and Fernando Louge for helpful comments, and three anonymous referees for very useful reports. Financial support from NSF Grant SES-0851580 is gratefully acknowledged.

†Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu; website: http://www.ssc.wisc.edu/~whs.

[1] See Sandholm (2008) for a recent survey.

Much of the appeal of stochastic stability theory lies in its ability to offer unique predictions in settings with multiple locally stable equilibria. This aspect of the analysis is most powerful when the prediction is robust to a range of choices about how to model the agents' updating process, so that confidence in the prediction is not predicated on having precise information about how agents make decisions.

The literature has identified at least two possible sources of nonrobustness. For one, the identity of the stochastically stable state may depend on how one specifies the probabilities of suboptimal choices in the agents' decision rule. An early version of this critique was offered by Bergin and Lipman (1996), who showed that allowing mistake probabilities to depend directly and arbitrarily on the current population state abrogates the possibility of general equilibrium selection results. While Bergin and Lipman (1996) allow a very wide range of choice rules in obtaining their negative result, the sensitivity of stochastic stability to the specification of choice rules can persist even if one only allows rules that admit a convincing economic justification. For instance, Blume (2003) shows that the best response with mutations (BRM) model of Kandori et al. (1993) and Young (1993) and the logit choice model from Blume (1993) can generate different stochastically stable states in two-strategy population games. Likewise, Myatt and Wallace (2003) exhibit a game in which the BRM and probit choice models select different equilibria. Ui (1998), Maruta (2002), and Dokumacı and Sandholm (2007) offer examples along similar lines. To the extent that these examples are representative, confidence in the predictions of stochastic stability theory must be contingent on precise knowledge of the agents' choice rules.

A more subtle source of nonrobustness of predictions, emphasized by Binmore et al. (1995) and Binmore and Samuelson (1997), concerns the identity and the order of limits used in defining stochastic stability. The early contributions of Kandori et al. (1993) and Young (1993) focus on the small noise limit, defining the stochastically stable states to be those that retain positive mass in the stationary distribution as noise level in agents' choices vanishes. A majority of the subsequent literature has followed this approach, which emphasize the influence of very rare mistakes on equilibrium selection relative to that of other factors.[2] A second branch of the literature follows Binmore and Samuelson (1997), Young (1998, Section 4.5), and Benaïm and Weibull (2003) in focusing on large population limits, defining the stochastically stable states to be those whose neighborhoods retain positive mass in the stationary distribution as the population size approaches infinity, and so emphasizing the influence of population size over improbability of mistakes in equilibrium selection.[3] To ease comparisons between these approaches, Binmore et al.

[2]For the furthest developments of this approach, see Ellison (2000) and Beggs (2005).

[3]Binmore and Samuelson (1997) argue that this emphasis is appropriate for most economic modeling.

(1995) and Binmore and Samuelson (1997) proposed definitions of stochastic stability that implement the small noise and large population limits sequentially; because the parameter in the outer limit is held fixed while the parameter in the inner limit taken to its extreme, it is the inner limit that governs equilibrium selection.

The distinction between the small noise and large population definitions of stochastic stability would be of little consequence if both approaches always generated the same predictions. But Binmore et al. (1995) and Binmore and Samuelson (1997) demonstrate that this distinction can matter: they find that when agents' choices are based on imitation and mutation, small noise stochastic stability always selects monomorphic states, while large population stochastic stability can select either boundary or interior states depending on the incentives in the underlying game.[4]

Still, in most work on stochastic evolution, agents' choices are governed not by imitation with mutation, but by noisy best responses. Whether the order of limits used in defining stochastic stability can affect equilibrium selection in noisy best response models seems to us a fundamental question, but as far as we know it is a question that the literature has not addressed.

The first goal of this paper is to provide a comprehensive answer to this question in a simple strategic environment: that of two-strategy population games. We consider a model of stochastic evolution in which agents employ noisy best response rules from the class introduced by Blume (2003). Let $\rho^{\eta}(a)$ denote the probability that a revising agent chooses the strategy with payoff advantage a when the noise level in his choice rule is η. We require that $\rho^{\eta}(a)$ have a well-defined exponential rate of decay in η, and that this rate of decay not decline as the payoff disadvantage of an inferior strategy becomes more severe. All of the noisy best response models noted above—BRM, logit choice, and probit choice—fall within this class.

In models of stochastic evolution, the stationary distribution $\mu^{N,\eta}$ summarizes the long-run behavior of the evolutionary process. Theorem 3.4 considers the double limit of the distributions $\mu^{N,\eta}$ as the noise level η approaches zero and the population size N.

[4]The logic behind these results can be explained as follows. In a model of imitation without mutations, the monomorphic states are absorbing states of the evolutionary process. This is so even if the underlying game is a Hawk-Dove game, whose unique symmetric Nash equilibrium is mixed. In such a game, fixing a small, positive mutation rate and taking the population size to infinity causes the stationary distribution to become concentrated in the vicinity of the mixed equilibrium. If instead we fix the population size and make the noise level sufficiently small, then the probability of escaping from a monomorphic state must become much smaller than the probability of reaching such a state from one near the mixed equilibrium; as a result, the stationary distribution becomes concentrated on (typically just one of) the monomorphic states. Börgers and Sarin (1997) make a similar point in their analysis of reinforcement learning in normal form games. For recent work on stochastic stability in models of imitation with mutations, see Fudenberg and Imhof (2006, 2008) and Sandholm (2009b).

It provides an exact characterization of the rates of decay of the stationary distribution weights in η and N, and, most importantly, establishes that these asymptotics are the identical for both orders of limits. Capturing the intermediate cases, Theorem 3.5 shows that the same asymptotics obtain when the limits in η and N are taken simultaneously. Given these descriptions of the limiting behavior of the full stationary distribution, it follows *a fortiori* that under noisy best response rules, the identity of the stochastically stable state is independent of the limits used to define stochastic stability.

With this analysis in hand, we return to the issue we raised first: that of the robustness of stochastic stability to the specification of agents' choice rules. We start with a series of examples in which different noisy best response rules generate different predictions of long run behavior. Taking these negative examples as background, we seek a condition on a game's payoffs that ensures the invariance of the stochastically stable state across all noisy best response rules. Ideally, this condition should be both simple and tight, being not only sufficient for "detail-free" equilibrium selection, but also necessary.

In the case of the BRM model, the necessary and sufficient condition for stochastic stability in two-strategy coordination games is well known. Kandori et al. (1993) and Young (1993) show that stochastic stability is determined by risk dominance, where a strategy is risk dominant if the set of states where it is optimal is larger than the set of states where the alternative strategy is optimal. The solution concept we introduce in this paper, which we call *stochastic dominance*, is a natural refinement of risk dominance: a strategy is stochastically dominant if for every level of payoff advantage $a \geq 0$, the set of states where the strategy outperforms its alternative by at least a is larger than the set where the alternative outperforms it by at least a. In population games with linear payoffs, such as those defined by random matching, risk dominance and stochastic dominance are equivalent, but in general population games stochastic dominance is the more demanding requirement of the two.

Theorem 5.3 offers our second robustness result: it establishes that stochastic dominance is a necessary and sufficient condition for a strategy to be stochastically stable under every noisy best response rule. Thus, if a game possesses a stochastically dominant strategy, our predictions of long run behavior need not be contingent on precise knowledge of the agents' decision rule. Conversely, if neither strategy is stochastically dominant, then such predictions must be made with caution, as different noisy best response rules will produce different stochastically stable states.

Section 2 introduces two-strategy population games, noisy best response rules, and our model of stochastic evolution. Section 3 provides a precise characterization of the asymptotics of the stationary distributions $\mu^{N,\eta}$, and shows that these asymptotics are unaffected by the order in which the small noise and large population limits are taken. Connections between the components of this result and earlier analyses of Blume (2003) and Sandholm (2007) are also explained here. Section 4 defines stochastic stability, and offers examples in which different noisy best response rules generate different stochastically stable states. Section 5 introduces the notion of a stochastically dominant strategy, and proves that coordination on a strategy is stochastically stable under any noisy best response rule if and only if that strategy is stochastically dominant. Section 6 concludes.

2. The Model

2.1 Two-Strategy Population Games

We consider games played by populations of N agents who choose strategies from the set $S = \{0, 1\}$. The population state $x \in \mathcal{X}^N = \{0, \frac{1}{N}, \frac{2}{N}, \ldots, 1\}$ describes the fraction of agents currently choosing strategy 1. If N is fixed, we can identify a game with its payoff function $F^N : \mathcal{X}^N \to \mathbf{R}^2$, where $F^N_i(x)$ is the payoff to strategy $i \in S$ at population state $x \in \mathcal{X}^N$.

Since we consider limits as the population size grows large, we find it useful to define a notion of convergence for sequences of finite-population games. The limit of such a sequence is a continuous-population game $F : [0, 1] \to \mathbf{R}^2$, which specifies a payoff for each strategy at each point in the unit interval. Our notion of convergence for sequences of games is uniform convergence: we say that the sequence $\{F^N\}_{N=N_0}^{\infty}$ converges to F if

$$(1) \qquad \lim_{N\to\infty} \max_{x\in\mathcal{X}^N} \left| F^N(x) - F(x) \right| = 0.$$

We assume throughout that the limit game F is bounded and piecewise continuous.

As an example, suppose that a population of size N is randomly matched without self-matching to play the two-player symmetric normal form game $A \in \mathbf{R}^{2\times 2}$. The payoffs of the resulting population game are

$$F^N_0(x) = \tfrac{N(1-x)-1}{N-1} A_{00} + \tfrac{Nx}{N-1} A_{01};$$
$$F^N_1(x) = \tfrac{N(1-x)}{N-1} A_{10} + \tfrac{Nx-1}{N-1} A_{11}.$$

As N grows large, the games $\{F^N\}$ converge uniformly to the limit game F, where

$$F_0(x) = (1 - x)A_{00} + xA_{01};$$
$$F_1(x) = (1 - x)A_{10} + xA_{11}.$$

While random matching generates games with linear payoff functions, our model allows payoffs to depend nonlinearly on the population state, as is often the case in models of congestion (Beckmann et al. (1956), Rosenthal (1973)), macroeconomic coordination (Topkis (1998), Cooper (1999)), and other sorts of multilateral externalities.

2.2 Revision Protocols and their Limits

We consider a model of stochastic evolution based on Blume (2003). Agents in this model receive revision opportunities via independent Poisson process. When a current strategy i player receives a revision opportunity, he switches to strategy $j \neq i$ with probability $\rho^\eta(a) \in (0,1)$, where $a \in \mathbf{R}$ represents the current payoff advantage of strategy j over strategy i: that is, the difference between the payoff to strategy j and the payoff to strategy i. The function $\rho^\eta : \mathbf{R} \to (0,1)$ is called a *revision protocol*, and is parameterized by a *noise level* $\eta > 0$.

We are interested in revision protocols under which agents typically select optimal strategies, but occasionally choose suboptimal ones. The protocols we allow satisfy

$$\lim_{\eta \to 0} \rho^\eta(a) = \begin{cases} 1 & \text{if } a > 0, \\ 0 & \text{if } a < 0. \end{cases}$$

But to place further structure on the probabilities of suboptimal choices, we impose restrictions on the rates at which the probabilities $\rho^\eta(a)$ of choosing a suboptimal strategy approach zero as η approaches zero.

Define the *cost* of switching to a strategy with payoff *disadvantage* $d \in \mathbf{R}$ as

$$c(d) = -\lim_{\eta \to 0} \eta \log \rho^\eta(-d). \tag{2}$$

When this limit exists, we can express the probability of switching to a strategy with payoff disadvantage d when the noise level is η as

$$\rho^\eta(-d) = \exp\left(-\eta^{-1}(c(d) + o(1))\right), \tag{3}$$

where $o(1)$ represents a term that vanishes as η approaches 0. Our assumptions on the revision protocols ρ^η and the cost function c are as follows:

(i) the limit in (2) exists for all $d \in \mathbf{R}$, with convergence uniform on compact intervals;
(ii) c is nondecreasing;
(iii) $c(d) = 0$ whenever $d < 0$;
(iv) $c(d) > 0$ whenever $d > 0$.

Conditions (ii)-(iv) impose constraints on the rates of decay of switching probabilities.[5] Condition (ii) requires the rate of decay to be nondecreasing in the payoff disadvantage of the alternative strategy. Condition (iii) requires the switching probability of an agent currently playing the suboptimal strategy to have rate of decay zero; the condition is satisfied when the probability is bounded away from zero, although this is not necessary for the condition to hold. Finally, condition (iv) requires the probability of switching from the optimal strategy to the suboptimal one to have a positive rate of decay. These conditions are consistent with having either $c(0) > 0$ or $c(0) = 0$: thus, when both strategies earn the same payoff, the probability that a revising agent opts to switch strategies can converge to zero with a positive rate of decay, as in Example 2.1 below, or can be bounded away from zero, as in Examples 2.2 and 2.3.

The examples to follow derive the cost functions for the three most prominent noisy best response models used in the literature. We note that Examples 2.1 and 2.2 appear in a slightly different form in Blume (2003), and that Examples 2.1 and 2.3 reparameterize the noise level before determining the cost functions.

Example 2.1. Best response with mutations. Suppose as in Kandori et al. (1993) and Young (1993) that the probability $\varepsilon > 0$ of abandoning an optimal strategy is independent of the payoff consequences of doing so:

$$\rho^{\varepsilon}(a) = \begin{cases} 1 - \varepsilon & \text{if } a > 0, \\ \varepsilon & \text{if } a \leq 0. \end{cases}$$

We call this the *best response with mutations* protocol, or *BRM* for short.

Let $\eta = -(\log \varepsilon)^{-1}$, so that $\varepsilon = \exp(-\eta^{-1})$. Then for $d \geq 0$, we have that $-\eta \log \rho^{\eta}(-d) = 1$, and so that $c(d) = 1$; for $d < 0$, we have that $c(d) = 0$, as required by condition (iii).[6] §

Example 2.2. Logit choice. Following Blume (1993, 1997), suppose that agents employ the

[5] It is evident from equation (3) that any function c satisfying conditions (ii)-(iv) is the cost function of some revision protocol ρ^{η}.

[6] That $c(0) = 1$ reflects our assumption that an indifferent player switches strategies only in the event of a mutation. None of our results would change if we assumed instead that an indifferent player has a fixed positive probability of switching (implying that $c(0) = 0$), or a probability of switching that decays relatively slowly (so that $c(0) \in (0, 1)$).

logit choice protocol with noise level $\eta > 0$:

$$\rho^{\eta}(a) = \frac{\exp(\eta^{-1}a)}{\exp(\eta^{-1}a) + 1}.$$

Then for $d \geq 0$, we have that $-\eta \log \rho^{\eta}(-d) = d + \eta \log(\exp(-\eta^{-1}d) + 1)$, so that $c(d) = d$. §

Example 2.3. Probit choice. The logit choice protocol from Example 2.2 can be derived from a random utility model in which the strategies' payoffs are perturbed by i.i.d., double exponentially distributed random variables.[7] The *probit choice protocol*, studied in evolutionary contexts by Myatt and Wallace (2003),[8] assumes instead that the payoff perturbations are i.i.d. normal random variables with mean 0 and variance σ^2. Thus

$$\rho^{\sigma^2}(a) = \mathbb{P}(\sigma Z + a > \sigma Z'),$$

where Z and Z' are independent and standard normal. It follows easily that

$$\rho^{\sigma^2}(a) = \Phi\left(\tfrac{a}{\sqrt{2}\sigma}\right), \tag{4}$$

where Φ is the standard normal distribution function.

It is well known (Durrett (2005, Theorem 1.1.3)) that when $z < 0$,

$$\Phi(z) = K(z)\exp(\tfrac{-z^2}{2}) \tag{5}$$

for some $K(z) \in (\frac{-1}{\sqrt{2\pi}\,z}(1 - \frac{1}{z^2}), \frac{-1}{\sqrt{2\pi}\,z})$. It follows that $K(z) \in (\frac{-1}{2\sqrt{2\pi}\,z}, \frac{-1}{\sqrt{2\pi}\,z})$ whenever $z < -\sqrt{2}$. In addition, it is easy to verify directly that (5) holds with $K(z) \in [\Phi(-\sqrt{2}), \frac{e}{2}]$ whenever $z \in [-\sqrt{2}, 0]$.

Now, letting $\eta = \sigma^2$, equations (4) and (5) imply that

$$-\eta \log \rho^{\eta}(-d) = -\eta \log \Phi\left(\tfrac{-d}{\sqrt{2\eta}}\right) = \tfrac{1}{4}d^2 - \eta \log K\left(\tfrac{-d}{\sqrt{2\eta}}\right) \tag{6}$$

when $d \geq 0$, with our earlier estimates showing that

$$\begin{aligned} &\eta \log K\left(\tfrac{-d}{\sqrt{2\eta}}\right) \in \left(\tfrac{1}{2}\eta \log \eta - \eta \log 2\sqrt{\pi}d, \tfrac{1}{2}\eta \log \eta - \eta \log \sqrt{\pi}d\right) && \text{if } d > 2\sqrt{\eta}, \text{ and} \\ &\eta \log K\left(\tfrac{-d}{\sqrt{2\eta}}\right) \in \left[\eta \log \Phi(-\sqrt{2}), \eta(1 - \log 2)\right] && \text{if } d \in [0, 2\sqrt{\eta}]. \end{aligned}$$

Thus, for any $D > 0$ and any $\delta > 0$, we have $|\eta \log K(\frac{-d}{\sqrt{2\eta}})| < \delta$ for all $d \in [0, D]$ once

[7] See Anderson et al. (1992) or Hofbauer and Sandholm (2002).

[8] See also Ui (1998) and Dokumacı and Sandholm (2007).

$\eta > 0$ is sufficiently small. We conclude from equation (6) that $-\eta \log \rho^\eta(-d)$ converges to $c(d) = \frac{1}{4}d^2$, with convergence uniform on compact intervals. §

Further examples of revision protocols and cost functions satisfying the assumptions above can be found in Dokumacı and Sandholm (2007).

2.3 The Stochastic Evolutionary Process

Let a population size N, a population game F^N, and a revision protocol ρ^η be given. To define the stochastic evolutionary process $\{X_t^{N,\eta}\}_{t\geq 0}$ on the state space $\mathcal{X}^N$, we suppose that each member of the population is equipped with an independent, rate 1 Poisson alarm clock. When an agent's clock rings, he uses the revision protocol ρ^η to decide whether to switch strategies.

Since the population size is finite, an agent who switches strategies moves the population state by an increment of $\frac{1}{N}$. Thus, an agent playing strategy 0 at population state x earns a payoff of $F_0^N(x)$, but if this agent switches to strategy 1, his payoff will not be $F_1^N(x)$, but rather $F_1^N(x + \frac{1}{N})$. An agent who accounts for this change when deciding whether to switch strategies is said to use *clever payoff evaluation* (Sandholm (1998)). By assuming that agents use clever payoff evaluation we simplify certain calculations, but all of our results remain true under the alternative assumption.

Since each of the N agents receives revision opportunities independently at rate 1, and since transitions are always to neighboring states in $\mathcal{X}^N$, the process $\{X_t^{N,\eta}\}$ is a birth and death process whose (possibly degenerate) jumps from each state $x \in \mathcal{X}^N$ occur at rate N.

For a transition from state x to state $x + \frac{1}{N}$ to occur, the agent who receives the revision opportunity must initially be playing strategy 0, and his revision protocol must tell him to switch to strategy 1. Under clever payoff evaluation, the probability that both of these events occur is

$$p_x^{N,\eta} = (1-x)\,\rho^\eta\left(F_1^N(x + \tfrac{1}{N}) - F_0^N(x)\right).$$

Similarly, for a transition from state x to state $x - \frac{1}{N}$ to occur, the agent who receives the revision opportunity must initially be playing strategy 1, and his revision protocol must tell him to switch to strategy 0. The probability that both of these events occur is

$$q_x^{N,\eta} = x\,\rho^\eta\left(F_0^N(x - \tfrac{1}{N}) - F_1^N(x)\right).$$

With the remaining probability of $1 - p_x^{N,\eta} - q_x^{N,\eta}$, the agent who receives the revision opportunity does not switch strategies, and the state does not change.

Because ρ^η is positive-valued, the process $\{X_t^{N,\eta}\}$ is irreducible, and so admits a unique stationary distribution $\mu^{N,\eta}$. This distribution describes the long run behavior of the process in two distinct ways: it is the the limiting distribution of the process, and it describes the limiting empirical distribution of the process along almost every sample path (see Durrett (2005, Sec. 5.5 and 6.2)).

3. The Limiting Stationary Distribution

Our goal in this section is to describe the asymptotics of the stationary distribution $\mu^{N,\eta}$ as the noise level η approaches zero and the population size N approaches infinity. To simplify the presentation we offer a few new definitions. First, given a continuous-population game $F : [0,1] \to \mathbf{R}^2$, we let

$$\Delta F(x) \equiv F_1(x) - F_0(x)$$

denote the payoff advantage of strategy 1 at state x. Next, given a cost function $c\colon \mathbf{R} \to [0,\infty)$ satisfying the conditions from Section 2.2, we define the *relative cost function* $\tilde{c}\colon \mathbf{R} \to \mathbf{R}$ by

$$\tilde{c}(d) = c(d) - c(-d) = \begin{cases} c(d) & \text{if } d > 0, \\ 0 & \text{if } d = 0, \\ -c(-d) & \text{if } d < 0. \end{cases} \tag{7}$$

Our assumptions on c imply that $\tilde{c}$ is nondecreasing, sign preserving ($\mathrm{sgn}(\tilde{c}(d)) = \mathrm{sgn}(d)$), and odd ($\tilde{c}(d) = -\tilde{c}(-d)$).

Now define the continuous function $I : [0,1] \to \mathbf{R}$ by

$$I(x) = \int_0^x \tilde{c}(\Delta F(y))\,\mathrm{d}y. \tag{8}$$

Observe that by marginally adjusting the state x so as to increase the mass on the optimal strategy, we increase the value of I at rate $\tilde{c}(a)$, where a is the optimal strategy's payoff advantage. In other words, I is an ordinal potential function for the game F (cf Monderer and Shapley (1996)). We now show that each of the revision protocols introduced in Section 2.2 generates a particularly simple ordinal potential.

Example 3.1. If ρ^η represents best response with mutations (Example 2.1), then (8) becomes the *signum potential function*

$$I_{\mathrm{sgn}}(x) = \int_0^x \mathrm{sgn}\,(\Delta F(y))\,\mathrm{d}y.$$

The slope of this function at state x is 1, -1, or 0, according to whether the optimal strategy at x is strategy 1, strategy 0, or both. Thus, I_{sgn} embodies the notion of "mutation counting" familiar from Kandori et al. (1993), Young (1993), and their successors. §

Example 3.2. If ρ^η represents logit choice (Example 2.2), then (8) becomes the (*linear*) *potential function*

$$I_1(x) = \int_0^x \Delta F(y)\,\mathrm{d}y,$$

whose slope at state x is given by the payoff difference at x.[9] Compared to I_{sgn}, the function I_1 accounts not only for the widths of the basins of attraction of the locally stable states, but also their "depths", as represented by payoff differences.[10] §

Example 3.3. If ρ^η represents probit choice (Example 2.3), then (8) becomes the *quadratic potential function*

$$I_2(x) = \int_0^x \tfrac{1}{4}\,\langle \Delta F(y)\rangle^2\,\mathrm{d}y,$$

where $\langle a\rangle^2 = \mathrm{sgn}(a)\,a^2$ is the signed square function. The values of I_2 again depend on payoff differences, but relative to the logit case, larger payoff differences play a more important role. This contrast can be traced to the fact that at small noise levels, the double exponential distribution (which underlies the logit protocol—see Example 2.3) has fatter tails than the normal distribution. §

As a final convenience, we define the function $\Delta I : [0,1] \to \mathbf{R}_-$ by

$$\Delta I(x) = I(x) - \max_{y\in[0,1]} I(y).$$

In words, ΔI is the ordinal potential function obtained from I by shifting its graph vertically until its maximum value is 0; since $I(0) = 0$, this shift can only be downward or null.

[9] Compare Sandholm (2009a), especially Example 4.4.

[10] The need to account for both "widths" and "depths" in stochastic stability analyses is emphasized by Fudenberg and Harris (1992), Binmore and Samuelson (1997), and Kandori (1997).

With these preliminaries in hand, we can now state our first main result, which describes the asymptotic behavior of the stationary distributions $\mu^{N,\eta}$ in the small noise and large population limits.

Theorem 3.4. *In the model of stochastic evolution above, the stationary distributions $\mu^{N,\eta}$ satisfy*

(*i*) $\lim_{N\to\infty}\lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\mu_x^{N,\eta}-\Delta I(x)\right|=0$ *and*

(*ii*) $\lim_{\eta\to 0}\lim_{N\to\infty}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\mu_x^{N,\eta}-\Delta I(x)\right|=0.$

In what follows, we offer an interpretation of this result, sketch its proof, and compare the components of the result to existing analyses.

As we explained in the introduction, analyses of long run behavior in models of stochastic evolution have proceeded along two distinct lines, with most analyses focusing on small noise limits, but with a significant minority emphasizing large population limits. When these two approaches lead to different conclusions—as happens, most notably, under decision rules based on imitation and rare mutations—modelers must exercise care in choosing which limit to use as the basis for predictions. As Binmore et al. (1995) and Binmore and Samuelson (1997) explain, the small noise limit, taken alone or taken first, puts precedence on the rareness of mutations, while taking the large population limit alone or first emphasizes population size over infrequency of mistakes. Binmore and Samuelson (1997) argue that in most economic contexts in which evolutionary models are relevant, it is the large population that is most appropriate. While we agree with this assessment in broad terms, it also seems to us that the best modeling choice may be application-dependent, and not always easy to discern. To the extent that this modeling choice affects predictions, it saps a key strength of the stochastic evolutionary approach: its ability to provide unique predictions of play in games with multiple strict equilibria.

Theorem 3.4 establishes in a strong sense that this concern about orders of limits is unnecessary when agents utilize noisy best response rules. Rather than focusing only on stochastic stability, the theorem characterizes the rates of decay of the stationary distribution weights at all population states.[11] It establishes that for either order of limits, the rates of decay of the stationary distribution weights are governed by the ordinal potential ΔI, with states that attain lower values of potential experiencing more rapid decay as η becomes small and N becomes large. The theorem thus shows that when agents employ noisy best response rules, the choice of the order of limits has essentially no effect on our predictions of long run behavior.

[11]To see that Theorem 3.4 concerns rates of decay, bear in mind that $\frac{\eta}{N}\log\mu_x^{N,\eta}=-r_x$ is equivalent to $\mu_x^{N,\eta}=\exp(-\eta^{-1}Nr_x)$.

The proof of Theorem 3.4 proceeds as follows. It is well known (see Durrett (2005, Sec. 5.4)) that the stationary distribution of an irreducible birth and death chain on $\mathcal{X}^N$ can be expressed as

$$\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} = \prod_{j=1}^{Nx} \frac{p_{(j-1)/N}}{q_{j/N}}$$

for $x \in \mathcal{X}^N - \{0\} = \{\frac{1}{N}, \ldots, 1\}$, with the value of $\mu_0^{N,\eta}$ being determined by the requirement that the probability weights sum to one. Substituting in the definitions of the transition probabilities p_x and q_x, taking logarithms, and then multiplying by $\frac{\eta}{N}$ yields

$$\text{(9)} \qquad \frac{\eta}{N} \log \frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} = \frac{1}{N} \sum_{j=1}^{Nx} \eta \log \frac{\rho^\eta \left(F_1^N(\frac{j}{N}) - F_0^N(\frac{j-1}{N})\right)}{\rho^\eta \left(F_0^N(\frac{j-1}{N}) - F_1^N(\frac{j}{N})\right)} + \frac{\eta}{N} \sum_{j=1}^{Nx} \log \frac{N-j+1}{j}.$$

The first term on the right hand side of (9) is a Riemann sum indexed by N, and when η is small, the jth summand is a discrete approximation of the relative cost $\tilde{c}(\Delta F(\frac{j}{N}))$. Repeated application of the dominated convergence theorem reveals that regardless of the order in which the limits in N and η are taken, this first term on the right hand side of (9) converges to the definite integral $I(x)$, and that the second term vanishes. These arguments characterize the rates of decay of the ratios $\mu_x^{N,\eta}/\mu_0^{N,\eta}$. To complete the proof, we use arguments building on the fact that each measure $\mu^{N,x}$ has total mass 1 to show that under either order of limits, the weights $\mu_x^{N,\eta}$ themselves have rates of decay given by $\Delta I(x)$. Since $\max_{x \in [0,1]} \Delta I(x) \equiv 0$, this confirms the intuition that the slowest rate of decay of a stationary distribution weight should equal zero. The details of the analysis summarized above can be found in the Appendix.

The two parts of Theorem 3.4 are descendants of earlier analyses from Blume (2003) and Sandholm (2007). Blume (2003) examines the small noise limit of the stationary distribution. He proves that in coordination games, when the population size is large enough, the mass in the limiting stationary distribution becomes concentrated on state 1 or state 0 according to whether $I(1)$ is greater than or less than $I(0)$. Relative to Blume's (2003) analysis, Theorem 3.4(i) explicitly introduces the population size limit, allows for arbitrary population games, and characterizes the asymptotics of the entire stationary distribution. Sandholm (2007) focuses on the large population limit, describing the rates of decay of the ratio $\mu_x^{N,\eta}/\mu_0^{N,\eta}$ in terms of the sum of $I(x)$ and an entropy term (see equation (24) in the Appendix). Theorem 3.4(ii) extends this analysis by introducing the small noise limit, and establishing the uniform convergence of the rates of decay of the weights $\mu_x^{N,\eta}$ themselves to the values of $\Delta I(x)$. Although the main reason for taking the second limit is to ease comparisons with the analysis in part (i), doing so also has the side benefit of providing a simple, closed form description of the asymptotics of the probit model.

While Theorem 3.4 takes the limits in η and N sequentially, one can also take these limits simultaneously. To do so, one can introduce a vanishing sequence of noise levels $\{\eta^N\}_{N=N_0}^{\infty}$, so that while the population size N approaches infinity, the noise level η^N approaches zero. Theorem 3.5 shows that taking simultaneous limits generates the same asymptotic behavior of the stationary distributions as taking either of the sequential limits. Since Theorem 3.5 does not control the relative speeds at which the noise level and population size approach their limits, apart from ruling out the lexicographic cases considered in Theorem 3.4, it demonstrates that the conclusions of Theorem 3.4 persist in all intermediate cases.

Theorem 3.5. *Let $\{\eta^N\}_{N=N_0}^{\infty}$ be a sequence of noise levels that converges to zero. Then the stationary distributions μ^{N,η^N} satisfy*

$$\lim_{N\to\infty} \max_{x\in\mathcal{X}^N} \left| \frac{\eta^N}{N} \log \mu_x^{N,\eta^N} - \Delta I(x) \right| = 0.$$

The proof of Theorem 3.5 is provided in the Appendix.

4. Stochastic Stability: Definition and Examples

Foster and Young (1990), Kandori et al. (1993), and Young (1993) define a stochastically stable state to be one that retains positive mass in the limiting stationary distribution as the noise level η approaches zero. Binmore et al. (1995) and Binmore and Samuelson (1997) extend this definition to allow the use of the large population limit and of multiple limits. In this section, we introduce definitions of stochastic stability for the present model, and use Theorem 3.4 to characterize stochastic stability in terms of the ordinal potential function I. We then present a series of examples illustrating that the stochastically stable state can depend on the choice of revision protocol. In Section 5, we provide a simple condition on payoffs that is necessary and sufficient for every noisy best response protocol to select the same stochastically stable state.

As one increases the population size N, the set of population states $\mathcal{X}^N$ becomes an increasingly fine grid in the unit interval. To account for this, we say that state x^* is *stochastically stable* if for every open set $O \subseteq \mathbf{R}$ containing x^*, we have

$$\lim_{N\to\infty} \lim_{\eta\to 0} \mu^{N,\eta}(O) = \lim_{\eta\to 0} \lim_{N\to\infty} \mu^{N,\eta}(O) > 0. \tag{10}$$

If there is an x^* such that both double limits in (10) equal 1, we call x^* *uniquely stochastically stable*.[12]

It follows from Theorem 3.4 that all stochastically stable states are maximizers of the ordinal potential function I.

Corollary 4.1. *All stochastically stable states elements of* $\operatorname{argmax}_{x\in[0,1]} I(x)$*. In particular, if* $\operatorname{argmax}_{x\in[0,1]} I(x) = \{x^*\}$*, then* x^* *is uniquely stochastically stable.*

The proof of this corollary is provided in the Appendix.

While Corollary 4.1 indicates that a unique maximizer of I must be stochastically stable, a larger set of maximizers of I may contain states that are not stochastically stable.[13] To account for this, we call the states in $\operatorname{argmax}_{x\in[0,1]} I(x)$ the *weakly stochastically stable states*.

Many stochastic stability analyses focus on equilibrium selection in coordination games. In the present context, we call $F : [0,1] \to \mathbf{R}^2$ a *coordination game* if there is a state $x^* \in (0,1)$ such that

$$\operatorname{sgn}(\Delta F(x)) = \operatorname{sgn}(x - x^*) \text{ for all } x \neq x^*.$$

Any ordinal potential function I for a coordination game is quasiconvex with two local maximizers: state $x = 0$, where all agents coordinate on strategy 0, and state $x = 1$, where all coordinate on strategy 1. Since $I(0) \equiv 0$, Corollary 4.1 tells us that state 1 is uniquely stochastically stable if $I(1) > 0$, and that state 0 is uniquely stochastically stable if $I(1) < 0$.

It is not difficult to construct examples in which different noisy best response protocols generate different equilibrium selections. To make our examples as simple as possible, we use games whose payoff functions are step functions, but one can easily construct similar examples using games with continuous payoffs. We sometimes use the notation e_i to refer the equilibrium in which all agents coordinate on strategy i: thus, $e_0 = 0$ and $e_1 = 1$.

Example 4.2. For each of the three protocols introduced in Examples 2.1–2.3, there are coordination games in which that protocol selects a different equilibrium than the other

[12]Some subtleties about these definitions should be noted. First, that x^* is the only stochastically stable state does not imply that x^* is uniquely stochastically stable in the sense specified above. Second, one can show that the requirement that $\lim_{N\to\infty} \lim_{\eta\to 0} \mu^{N,\eta} = \lim_{\eta\to 0} \lim_{N\to\infty} \mu^{N,\eta} = \delta_{x^*}$, where the limits refer to weak convergence of probability measures and δ_{x^*} represents a point mass at x^*, is more demanding than the requirement that x^* be uniquely stochastically stable. However, the differences among these formulations only arise in pathological cases.

[13]According to Theorem 3.4, the function ΔI describes exponential rates of decay of stationary distribution weights. It thus can hide subexponential discrepancies between rates of decay. For instance, if $\mu_0^{N,\eta} = \frac{1}{N}$ and $\mu_1^{N,\eta} = \frac{1}{\sqrt{N}}$, then $\lim_{N\to\infty} \frac{\eta}{N} \log \frac{1}{N} = 0 = \Delta I(0)$ and $\lim_{N\to\infty} \frac{\eta}{N} \log \frac{1}{\sqrt{N}} = 0 = \Delta I(1)$, but $\lim_{N\to\infty} \mu_0^{N,\eta}/\mu_1^{N,\eta} = 0$.

two. First consider a game with payoff differences

$$\Delta F(x) = \begin{cases} -1 & \text{if } x \in [0, \frac{2}{3}), \\ k & \text{if } x \in [\frac{2}{3}, 1], \end{cases}$$

where $k > 0$. We then have

$$\begin{aligned} I_{\text{sgn}}(1) &= -\tfrac{2}{3} + \tfrac{1}{3} = -\tfrac{1}{3}, \\ I_1(1) &= -\tfrac{2}{3} + \tfrac{1}{3}k = \tfrac{1}{3}(k-2), \text{ and} \\ I_2(1) &= \tfrac{1}{4}\left(-\tfrac{2}{3} + \tfrac{1}{3}k^2\right) = \tfrac{1}{12}(k^2 - 2). \end{aligned}$$

Under the BRM rule, the stochastically stable state is state 0 for any positive value of k; under the logit rule, the stochastically stable state switches from state 0 to state 1 at $k = 2$; and under the probit rule, this switch occurs at $k = \sqrt{2}$. Thus, when $k > 2$ only the BRM rule selects equilibrium e_0, and when $k \in (\sqrt{2}, 2)$ only the probit rule selects equilibrium e_1. §

Example 4.3. For an example in which the selection of the logit rule is distinct from that of the other two, suppose that

$$\Delta F(x) = \begin{cases} -7 & \text{if } x \in [0, \frac{1}{9}), \\ -1 & \text{if } x \in [\frac{1}{9}, \frac{2}{3}), \\ k & \text{if } x \in [\frac{2}{3}, 1], \end{cases}$$

where $k > 0$. Here we have

$$\begin{aligned} I_{\text{sgn}}(1) &= -\tfrac{1}{9} - \tfrac{5}{9} + \tfrac{1}{3} = -\tfrac{1}{3}, \\ I_1(1) &= -\tfrac{7}{9} - \tfrac{5}{9} + \tfrac{1}{3}k = \tfrac{1}{3}(k-4), \text{ and} \\ I_2(1) &= \tfrac{1}{4}\left(-\tfrac{49}{9} - \tfrac{5}{9} + \tfrac{1}{3}k^2\right) = \tfrac{1}{12}(k^2 - 18). \end{aligned}$$

When $k \in (4, \sqrt{18})$, only the logit rule selects equilibrium e_1. §

Example 4.4. If we move beyond coordination games, it is easy to construct examples in which each of the three choice protocols above generates a distinct equilibrium selection.

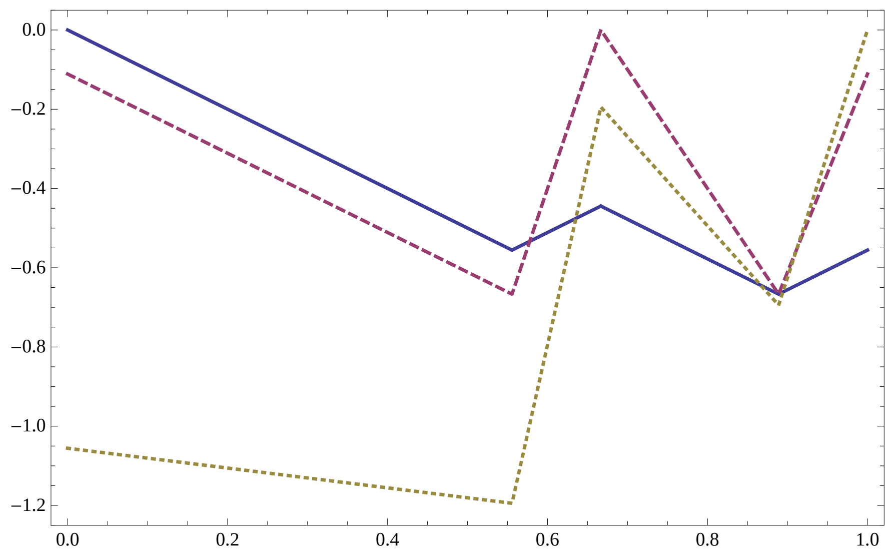

Figure 1: Ordinal potentials ΔI_{sgn} (solid), ΔI_1 (dashed), and ΔI_2 (dotted) in Example 4.4.

For instance, suppose that payoff differences in F are given by

$$\Delta F(x) = \begin{cases} -1 & \text{if } x \in [0, \frac{5}{9}), \\ 6 & \text{if } x \in [\frac{5}{9}, \frac{2}{3}), \\ -3 & \text{if } x \in [\frac{2}{3}, \frac{8}{9}), \\ 5 & \text{if } x \in [\frac{8}{9}, 1]. \end{cases}$$

The three candidates for stochastic stability are $x = 0$, $x = \frac{2}{3}$, and $x = 1$. In Figure 1, we graph the ordinal potentials ΔI_{sgn}, ΔI_1, and ΔI_2 generated by F. Evidently, state $x = 0$ is stochastically stable under the BRM rule, state $x = \frac{2}{3}$ is stochastically stable under the logit rule, and $x = 1$ is stochastically stable under the probit rule. §

5. Stochastic Dominance and Stochastic Stability

The most basic equilibrium selection result from stochastic evolutionary game theory, dating back to Kandori et al. (1993) and Young (1993), states that in two-strategy coordination games, risk dominance is a necessary and sufficient condition for stochastic stability under the BRM rule. In the coordination game $F : [0, 1] \to \mathbf{R}^2$ with mixed equilibrium $x^* \in (0, 1)$, strategy i is *strictly risk dominant* if the set of states where it is the unique best response is larger than the corresponding set for strategy $j \neq i$; thus, strategy 0 is strictly risk dominant if $x^* > \frac{1}{2}$, and strategy 1 is strictly risk dominant if $x^* < \frac{1}{2}$. If the relevant inequality holds weakly in either case, we call the strategy in question *risk dominant*.

The examples from Section 4 show that once one moves beyond the BRM rule, risk dominance is no longer a necessary or sufficient condition for stochastic stability.[14] In this section, we introduce a natural refinement of risk dominance called *stochastic dominance*, and show that it provides a necessary and sufficient condition for an equilibrium to be stochastically stable under every noisy best response rule.

To work toward our new definition, let us first observe that any function on the unit interval $[0, 1]$ can be viewed as a random variable by regarding the interval as a sample space endowed with Lebesgue measure λ. With this interpretation in mind, we define the *advantage distribution* of strategy i to be the cumulative distribution function of the payoff advantage of strategy i over the alternative strategy $j \neq i$:

$$G_i(a) = \lambda(\{x \in [0, 1] : F_i(x) - F_j(x) \leq a\}).$$

We let $\bar{G}_i$ denote the corresponding decumulative distribution function:

$$\bar{G}_i(a) = \lambda(\{x \in [0, 1] : F_i(x) - F_j(x) > a\}) = 1 - G_i(a).$$

In words, $\bar{G}_i(a)$ is the measure of the set of states at which the payoff to strategy i exceeds the payoff to strategy j by more than a.

It is easy to restate the definition of risk dominance in terms of the advantage distribution.

Observation 5.1. *Let F be a coordination game. Then strategy i is risk dominant if and only if $\bar{G}_i(0) \geq \bar{G}_j(0)$, and strategy i is strictly risk dominant if and only if $\bar{G}_i(0) > \bar{G}_j(0)$.*

To obtain our refinement of risk dominance, we require not only that strategy i be optimal at a larger set of states than strategy j, but also that strategy i have a payoff advantage of at least a at a larger set of states than strategy j for every $a \geq 0$.

Definition 5.2. *Let F be a coordination game. Then strategy i is* stochastically dominant *if $\bar{G}_i(a) \geq \bar{G}_j(a)$ for all $a \geq 0$. If in addition $\bar{G}_i(0) > \bar{G}_j(0)$, we say that strategy i is* strictly stochastically dominant.

Evidently, the notion of stochastic dominance for strategies proposed here is obtained by applying the usual definition of stochastic dominance from utility theory to the strategies' advantage distributions.

[14] Risk dominance retains its importance for equilibrium selection when payoffs are linear in the population state: see Corollary 5.4 below.

Blume (2003, Theorem 4) introduces three conditions on payoffs that are sufficient for an equilibrium of a coordination game to be stochastically stable in any noisy best response model; the conditions involve concavity, convexity, and skew-symmetry requirements on the payoff differences $\Delta F(x)$. In Theorem 5.3, we prove that stochastic dominance is both sufficient and necessary to ensure stochastic stability under every noisy best response rule.

Theorem 5.3. *Suppose that the finite-population games* $\{F^N\}_{N=N_0}^{\infty}$ *converge to the coordination game* $F : [0,1] \to \mathbf{R}^2$.

(*i*) *State* e_i *is weakly stochastically stable under every noisy best response protocol if and only if strategy* i *is stochastically dominant in* F.

(*ii*) *If strategy* i *is strictly stochastically dominant in* F, *then state* e_i *is uniquely stochastically stable under every noisy best response protocol.*

The idea behind Theorem 5.3 is simple. The definitions of I, $\tilde{c}$, c, ΔF, and G_i imply that

$$
\begin{aligned}
I(1) &= \int_0^1 \tilde{c}(\Delta F(y))\,\mathrm{d}y \\
&= \int_0^1 c(F_1(y) - F_0(y))\,\mathrm{d}y - \int_0^1 c(F_0(y) - F_1(y))\,\mathrm{d}y \\
&= \int_{-\infty}^{\infty} c(a)\,\mathrm{d}G_1(a) - \int_{-\infty}^{\infty} c(a)\,\mathrm{d}G_0(a).
\end{aligned}
\tag{11}
$$

As we have seen, whether state 1 or state 0 is stochastically stable depends on whether $I(1)$ is greater than or less than $I(0) = 0$. This in turn depends on whether the value of the first integral in the final line of (11) exceeds the value of the second integral. Once we recall that the cost function c is monotone, Theorem 5.3 reduces to a variation on the standard characterization of first-order stochastic dominance: namely, that distribution G_1 stochastically dominates distribution G_0 if and only if $\int c\,\mathrm{d}G_1 \geq \int c\,\mathrm{d}G_0$ for every nondecreasing function c.

Proof of Theorem 5.3. Again view $[0,1]$ as a sample space by endowing it with Lebesgue measure, and define $Y_i : [0,1] \to \mathbf{R}_+$ by

$$Y_i(\omega) = \sup\{a : G_i(a) < \omega\}.$$

Then it is easy to verify (or see Durrett (2005, Theorem 1.1.1)) that Y_i is a random variable with distribution G_i. It thus follows from equation (11) that

$$I(1) = \int_0^1 c(Y_1(\omega))\,\mathrm{d}\omega - \int_0^1 c(Y_0(\omega))\,\mathrm{d}\omega. \tag{12}$$

By construction we have that

$$\begin{aligned}
&Y_i(\omega) < 0 \text{ when } \omega \in [0, G_i(0-)),\\
&Y_i(\omega) = 0 \text{ when } \omega \in [G_i(0-), G_i(0)], \text{ and}\\
&Y_i(\omega) > 0 \text{ when } \omega \in (G_i(0), G_i(1)],
\end{aligned}$$

and that

$$G_1(0) - G_1(0-) = \lambda(\{x \in [0,1] : F_1(x) = F_0(x)\}) = G_0(0) - G_0(0-).$$

Thus, since c equals 0 on $(-\infty, 0)$, we can rewrite (12) as

$$\begin{aligned}
I(1) &= \int_{G_1(0-)}^1 c(Y_1(\omega))\,\mathrm{d}\omega - \int_{G_0(0-)}^1 c(Y_0(\omega))\,\mathrm{d}\omega\\
&= \int_{G_1(0)}^1 c(Y_1(\omega))\,\mathrm{d}\omega - \int_{G_0(0)}^1 c(Y_0(\omega))\,\mathrm{d}\omega.
\end{aligned} \tag{13}$$

To prove the "if" direction of part (i), suppose without loss of generality that strategy 1 is stochastically dominant in F. Then $G_1(a) \leq G_0(a)$ for all $a \geq 0$, so the definition of Y_i implies that $Y_1(\omega) \geq Y_0(\omega)$ for all $\omega \in [G_1(0), 1]$. Since c is nondecreasing and nonnegative, it follows from equation (13) that $I(1) \geq I(0)$, and hence that state 1 is weakly stochastically stable.

To prove part (ii), suppose without loss of generality that strategy 1 is strictly stochastically dominant in F. Then $G_1(a) \leq G_0(a)$ for all $a \geq 0$ and $G_1(0) < G_0(0)$. In this case, we not only have that $Y_1(\omega) \geq Y_0(\omega)$ for all $\omega \in [G_1(0), 1]$, but also that $Y_1(\omega) > 0$ when $\omega \in (G_1(0), G_0(0)]$. Since c is nondecreasing, and since it is positive on $(0, \infty)$, it follows from equation (13) that

$$I(1) \geq \int_{G_1(0)}^{G_0(0)} c(Y_1(\omega))\,\mathrm{d}\omega > 0 = I(0),$$

and hence that state 1 is uniquely stochastically stable.

Finally, to prove the "only if" direction of part (i), suppose without loss of generality that strategy 1 is not stochastically dominant in F. Then $G_1(b) > G_0(b)$ for some $b \geq 0$.

Now consider a noisy best response protocol with cost function

$$c(a) = \begin{cases} 0 & \text{if } a \leq 0, \\ 1 & \text{if } a \in (0, b], \\ C & \text{if } a > b, \end{cases}$$

where $C > G_1(b)/(G_1(b) - G_0(b))$. Then

$$\int_{-\infty}^{\infty} c(a)\, \mathrm{d}G_i(a) = (G_i(b) - G_i(0)) + C\,(1 - G_i(b))\,.$$

Therefore, equation (11) implies that

$$\begin{aligned} I(1) &= ((G_1(b) - G_1(0)) - (G_0(b) - G_0(0))) + C\,((1 - G_1(b)) - (1 - G_0(b))) \\ &\leq G_1(b) + C\,(G_0(b) - G_1(b)) \\ &< 0, \end{aligned}$$

implying that state 1 is not weakly stochastically stable. This completes the proof of the theorem. ■

Theorem 5.3 allows a simple proof of Blume's (2003) characterization of stochastic stability in coordination games with linear payoffs under noisy best response rules. The proof of the corollary below boils down to the observation that in linear coordination games, risk dominance and stochastic dominance are equivalent.

Corollary 5.4. *Let F be a coordination game with linear payoffs. Then under any noisy best response protocol, state e_i is uniquely stochastically stable if and only if strategy i is strictly risk dominant.*

Proof. Since F is a linear coordination game, $\Delta F(x) = k(x - x^*)$ for some $k > 0$, implying that $\bar{G}_0(a) = \max\{x^* - k^{-1}a, 0\}$ and that $\bar{G}_1(a) = \max\{(1 - x^*) - k^{-1}a, 0\}$ when $a \geq 0$. It follows immediately from Observation 5.1 that strategy i is strictly risk dominant if and only if it is strictly stochastically dominant. The corollary then follows from Theorem 5.3. ■

6. Conclusion

This paper considered the robustness of stochastic stability analysis to the order of limits used in defining stochastic stability, and to the specification of the agents' choice rule. We showed that in noisy best response models, the asymptotics of the stationary distribution are independent of the order in which the small noise and large population limits are taken; thus, definitions of stochastic stability based on either order of limits result yield identical predictions. We then introduced the notion of a stochastically dominant strategy, and established that coordination on a strategy is stochastically stable under any noisy best response rule if and only if that strategy is stochastically dominant.

By focusing on games with just two strategies, we were able to avail ourselves of the existence of an explicit formula for the stationary distribution $\mu^{\eta,N}$. But the questions studied here can be posed just as easily in the context of games with many strategies. Extending the foregoing analysis to this broader strategic context is an important and challenging direction for future research.

A. Appendix

The Proof of Theorem 3.4

To begin, we define $\tilde{\rho}^{\eta} : \mathbf{R} \to \mathbf{R}$ by

$$\tilde{\rho}^{\eta}(a) = \frac{\rho^{\eta}(a)}{\rho^{\eta}(-a)}.$$

Note that by the definition (7) of the relative cost function, we have that

$$\tilde{c}(a) = \lim_{\eta \to 0} \eta \log \tilde{\rho}^{\eta}(a),$$

while rewriting equation (9) in terms of $\tilde{\rho}$ yields

$$\log \frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} = \sum_{j=1}^{Nx} \left(\log \tilde{\rho}^{\eta} \left(F_1^N(\tfrac{j}{N}) - F_0^N(\tfrac{j-1}{N}) \right) + \log \tfrac{N-j+1}{j} \right). \tag{14}$$

To begin the proof of part (i), use equation (14) and the definition of $\tilde{c}$ to show that

$$\begin{aligned}
\lim_{\eta \to 0} \frac{\eta}{N} \log \frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} &= \lim_{\eta \to 0} \sum_{j=1}^{Nx} \frac{\eta}{N} \left(\log \tilde{\rho}^{\eta} \left(F_1^N(\tfrac{j}{N}) - F_0^N(\tfrac{j-1}{N}) \right) + \log \tfrac{(N-j+1)}{j} \right) \\
&= \frac{1}{N} \sum_{j=1}^{Nx} \tilde{c} \left(F_1^N(\tfrac{j}{N}) - F_0^N(\tfrac{j-1}{N}) \right).
\end{aligned}$$

for all $x \in \mathcal{X}^N$. Since $\mathcal{X}^N$ is a finite set, this limit is uniform in x: if we let

$$(15) \qquad I^N(x) = \frac{1}{N}\sum_{j=1}^{Nx} \tilde{c}\left(F_1^N(\tfrac{j}{N}) - F_0^N(\tfrac{j-1}{N})\right)$$

(which implies in particular that $I^N(0) = 0$), we have that

$$(16) \qquad \lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} - I^N(x)\right| = 0.$$

Recall that the finite-population games $F^N : \mathcal{X}^N \to \mathbf{R}^2$ converge uniformly to the limit game $F : [0,1] \to \mathbf{R}^2$, as described in equation (1), that F is bounded and piecewise continuous, and that $\tilde{c}$ is nondecreasing, and hence bounded on compact intervals. Therefore, if we define the functions $v^N : [0,1] \to \mathbf{R}$ by

$$v^N(x) = \begin{cases} \tilde{c}\left(F_1^N\left(\frac{\lceil Nx\rceil}{N}\right) - F_0^N\left(\frac{\lceil Nx\rceil - 1}{N}\right)\right) & \text{if } x \in (0,1], \\ \tilde{c}\left(F_1(0) - F_0(0)\right) & \text{if } x = 0, \end{cases}$$

then the v^N are uniformly bounded and converge almost surely to $v(x) \equiv \tilde{c}(F_1(x) - F_0(x))$. By construction, we have that $I^N(x) = \int_0^x v^N(y)\,\mathrm{d}y$ for all $x \in \mathcal{X}^N$. Thus, since $I(x) = \int_0^x v(y)\,\mathrm{d}y$ for all $x \in [0,1]$, equation (16), the triangle inequality, and the bounded convergence theorem imply that

$$(17) \qquad \lim_{N\to\infty}\lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} - I(x)\right| \leq \lim_{N\to\infty}\max_{x\in\mathcal{X}^N}\left|I^N(x) - I(x)\right| = 0.$$

For each fixed N, let x_*^N be a maximizer of I^N on $\mathcal{X}^N$. Then the uniform convergence established in (17) implies that

$$(18) \qquad \lim_{N\to\infty} I^N(x_*^N) = I(x_*), \quad \text{where } \; x_* \in \operatorname*{argmax}_{x\in[0,1]} I(x).$$

We claim that

$$(19) \qquad \lim_{\eta\to 0}\frac{\eta}{N}\log\mu_{x_*^N}^{N,\eta} = \lim_{N\to\infty}\lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\left(I^N(x) - I^N(x_*^N)\right) - \left(I(x) - I(x_*)\right)\right| = 0.$$

If this is true, then equation (16) implies that

$$\text{(20)} \qquad \lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\mu_x^{N,\eta}-\left(I^N(x)-I^N(x_*^N)\right)\right|$$
$$=\lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\left(\frac{\eta}{N}\log\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}}-I^N(x)\right)-\left(\frac{\eta}{N}\log\frac{\mu_{x_*^N}^{N,\eta}}{\mu_0^{N,\eta}}-I^N(x_*^N)\right)+\frac{\eta}{N}\log\mu_{x_*^N}^{N,\eta}\right|=0;$$

then equations (20), (17), and (18) yield

$$\text{(21)} \quad \lim_{N\to\infty}\lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\mu_x^{N,\eta}-\Delta I(x)\right|=\lim_{N\to\infty}\lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\left(I^N(x)-I^N(x_*^N)\right)-\left(I(x)-I(x_*)\right)\right|=0,$$

proving part (i) of the theorem.

To establish (19), first suppose to the contrary that there is a sequence $\{\eta^k\}$ converging to zero along which the limit in (19) is $-c<0$. In this case, the reasoning in equation (20) implies that

$$\lim_{\eta^k\to 0}\max_{x\in\mathcal{X}^N}\left|\frac{\eta^k}{N}\log\mu_x^{N,\eta^k}-\left(I^N(x)-I^N(x_*^N)-c\right)\right|=0.$$

Since $I^N(x)\le I^N(x_*^N)$ for all $x\in\mathcal{X}^N$, it follows that for η^k far enough along the sequence, we have that $\frac{\eta^k}{N}\log\mu_x^{N,\eta^k}\le-\frac{c}{2}$ for all $x\in\mathcal{X}^N$, and hence that

$$\sum_{x\in\mathcal{X}^N}\mu_x^{N,\eta^k}=\sum_{x\in\mathcal{X}^N}\exp\left(\frac{N}{\eta^k}\cdot\frac{\eta^k}{N}\log\mu_x^{N,\eta^k}\right)\le(N+1)\exp\left(-\frac{cN}{2\eta^k}\right).$$

The last expression vanishes as k grows large, contradicting the fact that μ^{N,η^k} is a probability measure.

Second, suppose contrary to (19) that there is a sequence $\{\eta^k\}$ converging to zero along which the limit in (19) is $c>0$. Then by definition, there is a sequence $\{\delta^k\}$ converging to zero such that

$$\mu_{x_*^N}^{N,\eta^k}=\exp\left(\tfrac{N}{\eta^k}(c+\delta^k)\right).$$

The right hand expression grows without bound as k grows large, contradicting the fact that μ^{N,η^k} is a probability measure. This completes the proof of part (i).

We proceed with the proof of part (ii). Equation (14) implies that

$$\frac{\eta}{N}\log\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} = \frac{\eta}{N}\sum_{j=1}^{Nx}\log\tilde{\rho}^{\eta}\left(F_1^N(\tfrac{j}{N}) - F_0^N(\tfrac{j-1}{N})\right) + \frac{\eta}{N}\sum_{j=1}^{Nx}\log\tfrac{N-j+1}{N} - \frac{\eta}{N}\sum_{j=1}^{Nx}\log\tfrac{j}{N}. \tag{22}$$

Now ρ^{η} is bounded away from zero, and $0 \geq \log\left(\frac{\lceil Nx\rceil}{N}\right) \geq \log(x)$ and $0 \geq \log\left(\frac{N-\lceil Nx\rceil+1}{N}\right) \geq \log(1-x)$ for $x \in (0,1)$. Thus, following the logic used to establish (17) (but applying the dominated convergence theorem to the second and third sums in (22)) yields

$$\begin{aligned}\lim_{N\to\infty}\frac{\eta}{N}\log\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} &= \int_0^x \eta\left(\log\tilde{\rho}^{\eta}\left(F_1(y) - F_0(y)\right) + \log(1-y) - \log(y)\right)\mathrm{d}y \\ &= \int_0^x \eta\log\tilde{\rho}^{\eta}\left(\Delta F(y)\right)\mathrm{d}y - \eta\left(x\log x + (1-x)\log(1-x)\right),\end{aligned} \tag{23}$$

where the limit is taken over those N for which $x \in \mathcal{X}^N$, and where we follow the convention that $0\log 0 = 0$. Moreover, since increasing the length of the interval of integration $[0,x]$ only worsens the bound on the speed of convergence in (23), the worst bound obtains when $x = 1$. This implies that convergence in (23) is uniform in x: if we let

$$\begin{aligned}&h(x) = -(x\log x + (1-x)\log(1-x)) \text{ and} \\ &I^{\eta}(x) = \int_0^x \eta\log\tilde{\rho}^{\eta}\left(\Delta F(y)\right)\mathrm{d}y + \eta\, h(x),\end{aligned} \tag{24}$$

we have that

$$\lim_{N\to\infty}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} - I^{\eta}(x)\right| = 0. \tag{25}$$

Since payoffs are bounded, and since convergence in (2) is uniform on compact intervals, the bounded convergence theorem implies that

$$\begin{aligned}\lim_{\eta\to 0}\int_0^x \eta\log\tilde{\rho}^{\eta}\left(\Delta F(y)\right)\mathrm{d}y &= \int_0^x\left(\lim_{\eta\to 0}\eta\log\tilde{\rho}^{\eta}\left(\Delta F(y)\right)\right)\mathrm{d}y \\ &= \int_0^x \tilde{c}\left(\Delta F(y)\right)\mathrm{d}y \\ &= I(x)\end{aligned}$$

uniformly in x. This fact, the previous two equations, and the triangle inequality yield

$$\lim_{\eta\to 0}\lim_{N\to\infty}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} - I(x)\right| = 0. \tag{26}$$

The remainder of the proof of part (ii) is similar to the second part of the proof of part (i). For each η, let x_*^η be a maximizer of I^η on $[0,1]$. Evidently, $\lim_{\eta\to 0} I^\eta(x_*^\eta) = I(x_*)$, where x_* maximizes I on $[0,1]$. Now we claim that

$$\lim_{N\to\infty}\frac{\eta}{N}\log\mu_{x_*^\eta}^{N,\eta} = 0. \tag{27}$$

If this is true, then using equations (25) and (26) to mimic the analogous argument (equations (20) and (21)) from the proof of part (i) establishes part (ii) of the theorem. But (27) can be verified through essentially the same argument used to verify (19) in the proof of part (i). Thus, part (ii) is established, and the proof of the theorem is complete. ■

The Proof of Theorem 3.5

Equation (14) implies that

$$\frac{\eta^N}{N}\log\frac{\mu_x^{N,\eta^N}}{\mu_0^{N,\eta^N}} = \frac{1}{N}\sum_{j=1}^{Nx}\eta^N\log\tilde{\rho}^{\eta^N}\left(F_1^N(\tfrac{j}{N}) - F_0^N(\tfrac{j-1}{N})\right) + \eta^N\left(\frac{1}{N}\sum_{j=1}^{Nx}\log\frac{N-j+1}{j}\right). \tag{28}$$

The proof of Theorem 3.4(ii) shows that the expression in parentheses converges to $h(x)$ uniformly in x as N approaches infinity. Since h is bounded on $[0,1]$ and since $\lim_{N\to\infty}\eta^N = 0$, the second summand in (28) converges to zero uniformly in x as N approaches infinity.

To contend with the first summand in (28), recall that (i) $F^N\colon \mathcal{X}^N \to \mathbf{R}^2$ converges uniformly to the bounded, piecewise continuous function $F\colon X \to \mathbf{R}^2$ as N approaches infinity, (ii) $\eta\log\tilde{\rho}^\eta(\pi)$ converges to $\tilde{c}(\pi)$ as η approaches zero, uniformly on compact intervals, and (iii) $\tilde{c}$ is nondecreasing, and hence bounded on compact intervals. Therefore, if we define the functions $w^N : [0,1] \to \mathbf{R}$ by

$$w^N(x) = \begin{cases} \eta^N\log\tilde{\rho}^{\eta^N}\left(F_1^N\left(\frac{\lceil Nx\rceil}{N}\right) - F_0^N\left(\frac{\lceil Nx\rceil - 1}{N}\right)\right) & \text{if } x\in(0,1], \\ \eta^N\log\tilde{\rho}^{\eta^N}\left(F_1(0) - F_0(0)\right) & \text{if } x = 0, \end{cases}$$

then the w^N are uniformly bounded and converge almost surely to $v(x) \equiv \tilde{c}(F_1(x) - F_0(x))$. Since the first summand in (28) is equal to $\int_0^x w^N(y)\,\mathrm{d}y$, the bounded convergence theorem implies that this summand converges to $I(x) = \int_0^x v(y)\,\mathrm{d}y$ uniformly in x as N approaches infinity.

Combining these arguments shows that

$$\lim_{N\to\infty} \max_{x\in\mathcal{X}^N} \left| \frac{\eta^N}{N} \log \frac{\mu_x^{N,\eta^N}}{\mu_0^{N,\eta^N}} - I(x) \right| = 0.$$

The remainder of the proof is similar to the second half of the proof of Theorem 3.4(i). ■

The Proof of Corollary 4.1

We consider only the first double limit in (10); the proof for the other double limit is virtually identical. Moreover, once we prove the first statement in the corollary, the second follows immediately.

To prove this first statement, suppose that $x \in [0,1]$ does not maximize I, so that $\Delta I(x) < 0$. Since I is continuous, we can find an open set O containing x such that for some $\varepsilon > 0$, we have $\Delta I(y) < -\varepsilon$ for all $y \in O$. Now define

$$d_x^{N,\eta} = \tfrac{\eta}{N} \log \mu_x^{N,\eta} - \Delta I(x), \;\; d^{N,\eta} = \max_{x\in\mathcal{X}^N} \left| d_x^{N,\eta} \right|, \;\; \text{and} \;\; d^N = \lim_{\eta\to 0} d^{N,\eta}.$$

Theorem 3.4 tells us that d^N exists for all large enough N, and that $\lim_{N\to\infty} d^N = 0$. It follows that there is an $\underline{N}$ such that $d^N < \frac{\varepsilon}{3}$ whenever $N \geq \underline{N}$, and thus that for each such N there is an $\underline{\eta}(N) > 0$ such that $d^{N,\eta} < \frac{2\varepsilon}{3}$ whenever $N \geq \underline{N}$ and $\eta \leq \underline{\eta}(N)$.

Using Theorem 3.4 once more, we see that for $N \geq \underline{N}$, $\eta \leq \underline{\eta}(N)$, and $y \in O$, we have

$$\mu_y^{N,\eta} = \exp\left(\eta^{-1} N \left(\Delta I(y) + d_y^{N,\eta}\right)\right) < \exp(-\eta^{-1} N \cdot \tfrac{\varepsilon}{3}).$$

Therefore, $\mu^{N,\eta}(O) < (N+1)\exp(-\eta^{-1} N \cdot \frac{\varepsilon}{3})$, which implies that

$$\lim_{N\to\infty} \lim_{\eta\to 0} \mu^{N,\eta}(O) = 0.$$

Thus, x is not stochastically stable. This competes the proof of the corollary. ■

References

Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge.

Beckmann, M., McGuire, C. B., and Winsten, C. B. (1956). *Studies in the Economics of Transportation*. Yale University Press, New Haven.

Beggs, A. W. (2005). Waiting times and equilibrium selection. *Economic Theory*, 25:599–628.

Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.

Bergin, J. and Lipman, B. L. (1996). Evolution with state-dependent mutations. *Econometrica*, 64:943–956.

Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74:235–265.

Binmore, K., Samuelson, L., and Vaughan, R. (1995). Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424.

Blume, L. E. (1997). Population games. In Arthur, W. B., Durlauf, S. N., and Lane, D. A., editors, *The Economy as an Evolving Complex System II*, pages 425–460. Addison-Wesley, Reading, MA.

Blume, L. E. (2003). How noise matters. *Games and Economic Behavior*, 44:251–271.

Börgers, T. and Sarin, R. (1997). Learning through reinforcement and the replicator dynamics. *Journal of Economic Theory*, 77:1–14.

Cooper, R. W. (1999). *Coordination Games: Complementarities and Macroeconomics*. Cambridge University Press, Cambridge.

Dokumacı E. and Sandholm, W. H. (2007). Stochastic evolution with perturbed payoffs and rapid play. Unpublished manuscript, University of Wisconsin.

Durrett, R. (2005). *Probability: Theory and Examples*. Brooks-Cole, Belmont, CA, third edition.

Ellison, G. (2000). Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–45.

Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232.

Fudenberg, D. and Harris, C. (1992). Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57:420–441.

Fudenberg, D. and Imhof, L. A. (2006). Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262.

Fudenberg, D. and Imhof, L. A. (2008). Monotone imitation dynamics in large populations. *Journal of Economic Theory*, 140:229–245.

Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70:2265–2294.

Kandori, M. (1997). Evolutionary game theory in economics. In Kreps, D. M. and Wallis, K. F., editors, *Advances in Economics and Econometrics: Theory and Applications*, volume I, pages 243–277. Cambridge University Press, Cambridge.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.

Maruta, T. (2002). Binary games with state dependent stochastic choice. *Journal of Economic Theory*, 103:351–376.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.

Myatt, D. P. and Wallace, C. C. (2003). A multinomial probit model of stochastic evolution. *Journal of Economic Theory*, 113:286–301.

Rosenthal, R. W. (1973). A class of games possessing pure strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67.

Sandholm, W. H. (1998). Simple and clever decision rules in a model of evolution. *Economics Letters*, 61:165–170.

Sandholm, W. H. (2007). Simple formulas for stationary distributions and stochastically stable states. *Games and Economic Behavior*, 59:154–162.

Sandholm, W. H. (2008). Evolutionary game theory. In Meyers, R. A., editor, *Encyclopedia of Complexity and System Science*. Forthcoming, Springer, Heidelberg.

Sandholm, W. H. (2009a). Large population potential games. *Journal of Economic Theory*, 144:1710–1725.

Sandholm, W. H. (2009b). Stochastic imitative dynamics with committed agents. Unpublished manuscript, University of Wisconsin.

Topkis, D. (1998). *Supermodularity and Complementarity*. Princeton University Press, Princeton.

Ui, T. (1998). Robustness of stochastic stability. Unpublished manuscript, Bank of Japan.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.

Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.