How Noise Matters†

Lawrence E. Blume
Department of Economics
Cornell University
Ithaca NY 14853

October 1999

† This research was supported by the National Science Foundation. It was carried out during visits to the Institute for Advanced Studies at Hebrew University and the Department of Economics at Tel Aviv University. The author is grateful to Bart Lipman and Peyton Young for good discussions, to Murali Agastya for good talk and for prodding him into writing this paper, and to Rajshri Jayaraman and Mandar Oak for helpful comments.

Abstract

Recent advances in evolutionary game theory have introduced noise into decisionmaking as an equilibrium selection device — to select equilibria robust to noisy perturbations. Noisy decisionmaking is justified on bounded rationality grounds, and consequently the sources of noise are left unmodelled. This methodological approach can only be successful if the results do not depend too much on the nature of the noise process. This paper investigates invariance to noise of the equilibrium selection results in coordination games, both for the random matching paradigm that has characterized much of the recent literature and for a larger class of two-strategy population games where payoffs may vary non-linearly with the distribution of strategies among the population. Several parametrizations of noise are investigated. The results show that a symmetry property of the noise process and, in the case of non-linear payoffs, bounds on the asymmetry of the payoff functions, suffice to preserve the known stochastic stability results.

JEL Classification: C78

Correspondent:
Professor Lawrence Blume
Department of Economics
Uris Hall
Cornell University
Ithaca, NY 14853

> Father Darling: "A little less noise there, a little less noise!"
>
> J. M. Barrie, *Peter Pan.*

1 Introduction

In recent years a dynamic analysis of the evolution of choice in strategic environments has emerged which demonstrates the importance of random noise in driving population behavior towards particular equilibria and away from others. Since the work of Foster and Young (1990), Canning (1992) and Fudenberg and Harris (1992), and of Blume (1993), Kandori, Mailath, and Rob (1993), and Young (1993), the power of small stochastic perturbations to drive equilibrium selection, and in particular the selection of risk-dominant play as the *stochastically stable* equilibrium in coordination games, has become one of the most studied topics in game theory.

Whereas the first papers in this literature simply imposed aggregate noise on deterministic population dynamics, subsequent research firmly located the source of that noise in the random behavior of individuals in the population. That is, every individual's strategic choice is a random perturbation of her best response to current beliefs.[1] There are many justifications for introducing randomness: biological mutation, experimentation, cognitive mistakes, and unobserved player-specific payoff variations. However, the point of this literature has been not to model the fine details of the noise process but instead just to assume enough properties to derive the results.

Unfortunately, the identity of stochastically stable states is not independent of the noise process. For instance, in the (Kandori, Mailath, and Rob 1993) model, it is clear that if the probability of deviating from strategy A is ϵ when A is the best response, and if the probability of deviating from strategy B is 0 when B is the best response, then B will be the only stochastically stable equilibrium regardless of risk or payoff dominance. From this example it is obvious that the rates at which perturbations go to 0 matters. This was worked out by Bergin and Lipman (1996), who showed that if perturbations were allowed to be state-dependent in an arbitrary fashion, any state can be made stochastically stable by tinkering with the rates at which perturbaions vanish. But knowing that the perturbations must behave in a certain way does not always say much about the processes that generate them. In particular this will be a problem when the perturbation magnitudes depend upon endogenously genenerated objects such as expected payoffs.

This paper takes a different tack. It introduces a somewhat generic means of constructing trembles from strategies' payoffs, called *noise models.* Noise models map payoff differences into trembles: The probability of playing strategy A depends upon the difference between A's expected utility and B's expected utility in the current state. Noise models are parametrized by a parameter β such that the variance around the best response decreases with $1/\beta$. I characterize those noise models which generate the usual stochastic stability results in two-player symmetric coordination games. As long as the noise is *symmetric* in a sense to be made precise below, the particular noise process does not influence the equilibrium selection results. Moreover, this symmetry condition turns out to be sufficient for invariance results in more general binary-choice coordination problems in which the payoffs do not depend linearly on the distribution of strategies among the population. I provide a characterization of those payoff functions for which all symmetric noise processes give the same selection result.

Section 2 contains a formal construction of the strategy revision process for pairwise random matching environments. Section 3 discusses the invariant distribution of a strategy revision process. The notion of a potential for the population game is introduced, and is related to the potential for the two-person game which describes the pairwise interactions and to the limit behavior of the strategy revision process. Theorems 1 and 2 of section 4 are the basic results of the paper. They characterize the set of noise processes for which risk-dominance is stochastically stable. One characteristic of the random matching model is that the payoff difference between the two strategies is linear with respect to the frequency of choice. Section 5 extends the analysis to other population processes which are non-linear in a sense which contrasts with the linearity of pairwise random matching interaction. Section 6 takes up one particular source of noisy choice, that of *random utility models.* Conditions are given on the distribution of the stochastic components of utility under which conclusions comparable to those of Theorems 1 and 2 hold. Wild conclusions and speculations are reserved for Section 7.

2 Player Choice Rules for Pairwise Random Matching

The population model in this paper is that of (Blume 1993), but applied to a completely connected graph of players. Consider a population of N players who are randomly paired to play a two-strategy symmetric coordination game. A *strategy revision process* records the distribution of strategies in the population through time. At randomly chosen moments, individual players have an opportunity to revise their current strategy choice. Their decision depends upon the expected utility difference between the two strategies computed according to the distribution recorded by the strategy revision process at the moment when the strategy revision opportunity arrives. For instance, players may best-respond to the current distribution of play. A stochastic alternative is the "mistakes model" of (Kandori, Mailath, and Rob 1993), in which a player chooses her best response with fixed high probability, and the alternative with the residual low probability. Another noise model, the "log-linear model" of (Blume 1993) and (Brock and Durlauf 1995), has the log of the odds of choosing a given strategy proportional to the payoff difference between the two strategies. Finally, as players revise their choices, the strategy revision process is updated to reflect the new distribution of play. The remainder of this section contains a more formal and detailed description of strategy revision processes.

A population of players must repeatedly choose between actions α and ζ. A payoff matrix G is given. Row player payoffs are

	α	ζ
α	$u_{\alpha\alpha}$	$u_{\alpha\zeta}$
ζ	$u_{\zeta\alpha}$	$u_{\zeta\zeta}$

This paper follows Blume (1993) and Kandori, Mailath, and Rob (1993) in assuming that players' beliefs at the moment of a revision opportunity are identical to the current distribution of play in the population.[2] Given the payoff matrix, let $U_\alpha(p) = pu_{\alpha\alpha} + (1-p)u_{\alpha\zeta}$ and $U_\zeta(p) = pu_{\zeta\alpha} + (1-p)u_{\zeta\zeta}$ denote payoffs for actions α and ζ, respectively, as a function of p — the likelihood that the player assigns to being matched against an opponent playing action α.[3] Thus the difference in payoff between actions α and ζ, $\Delta(p) = U_\alpha(p) - U_\zeta(p)$, is a linear function of p. (Other possibilities encompassing alternative interaction paradigms are discussed in section 6.) Throughout this paper I will assume that the underlying two-player game G is a coordination game. Consequently $\Delta(p)$ will be upward sloping. I will also assume that the strategy α is risk-dominant. Thus $\Delta(p^*) = 0$ for some $p^* < 1/2$. Finally, I will assume that no strategy is dominated (for otherwise this exercise is trivial), so $\Delta(0) < 0$; that is, $p^* > 0$.

A noise model is a mechanism for assigning choice probabilities to payoff differences. It is defined by a function $g : \mathbf{R} \to \mathbf{R}$ which defines choice probabilities in the following manner. Let $\sigma_\alpha(p)$ and $\sigma_\zeta(p)$ denote the probabilities that a player will select α and ζ, respectively, at a strategic revision opportunity when the fraction of the population choosing action α is p. These probabilities are defined by the function g such that

$$\frac{\sigma_\alpha(p)}{\sigma_\zeta(p)} = \exp g\big(\Delta(p), \beta\big) . \tag{1}$$

In other words,

$$\sigma_\alpha(p) = \frac{\exp g\big(\Delta(p), \beta\big)}{1 + \exp g\big(\Delta(p), \beta\big)} . \tag{2}$$

The function g satisfies the following assumptions:

Axiom 1. *$g(x, \beta) = \beta h(x) + r(x, \beta)$, and on every bounded interval I, $\lim_{\beta\to\infty} \sup_{x\in I} |r(x, \beta)|/\beta = 0$.*

Axiom 2. *The function h is Riemann integrable, non-decreasing and strictly increasing at 0. That is, in every open neighborhood of 0 there are numbers $x < 0 < y$ such that $h(x) < h(0) < h(y)$.*

Axiom 3. *For all β, $g(0, \beta) = 0$.*

Notice that assumptions 1 and 3 imply that $h(0) = 0$ and $r(0, \beta) = 0$ for all $\beta > 0$.

Assumptions 2 and 3 capture two ideas: First, the greater the payoff advantage to α, the more likely it is to be chosen, and second, that the noise process is unbiased in the sense that, when the two strategies have the same pay-off, each is equally likely to be chosen.

Any probability distribution on a finite state space that assigns positive probabilities to all events can be represented in terms of odds ratios, and the odds ratios can always be written in exponential form; that is, e raised to some power. What is not general about the description of choice is the parametric representation through β of the level of noise. (Section 4 explores a different parametrization of noisy choice.) Nonetheless the most popular parametrizations of noise and noise reduction fit into this scheme. For instance, the mistakes model can be represented by choosing β equal to $\ln(1-\epsilon)/\epsilon$ and $h(\Delta)$ equal to 1, 0, or -1 as Δ is positive, 0, or negative. That is, $h(\Delta) = \mathrm{sgn}(\Delta)$. The log-linear model has $h(\Delta) = \Delta$. In both cases, $r(x, \beta) \equiv 0$.

In most of the population games literature the allocation of strategy revision opportunities among players is modelled as a discrete-time stochastic processes wherein at each tick of the clock each player is awarded a strategy revision opportunity with some fixed probability independent of the situation of all other players.[4] This paper proceeds somewhat differently. Time is continuous. Each player has a rate 1 Poisson alarm clock: A player gets strategy revision opportunities at randomly chosen moments. The length of the time intervals between the arrival of successive strategy revision opportunities are independently distributed according to a mean 1 exponential distribution. Formally, to each player we associate a sequence $\{X_{ij}\}$ of random variables. $\sum_{j=1}^{J} X_{ij}$ is the time of the Jth strategy revision opportunity for individual i. The random variables X_{ij} are all independently and exponentially distributed with parameter 1. If there are N players, it follows that strategy revision oppotunities for the population as a whole arrive at independent intervals exponentially distributed with rate N. Strategy revision opportunities for any group of M players arrive at rate M. The probability of any two or more players revising at the same time is 0.

This continuous time model with a finite population of players is formally equivalent to a discrete time model in which at each tick one player is chosen at random to have a strategy revision opportunity. A technical virtue of this "one player at a time" formalism is its ease of use. The complex calculations with Z-trees are avoided because the strategy revision process is a birth-death process. Appeals to realism at the level of abstraction of the population games literature are a bit ludicrous, but we also suspect that in realistic modelling of social and economic phenomena in a population games framework, this continuous time construction will appear no less realistic than the alternatives. (This can already be seen in the literature on the microeconomics of sorting and searching, which typically starts with this assumption and then, frequently incorrectly, moves to a large-numbers limit. We hope not to misuse the SLLN here.) Moreover, the story just described is a continuous-time limit of discrete-time processes. Imagine a sequence of discrete time strategy revision processes as in Samuelson (1994), where with some probability each player (independently) gets what he calls a "learn draw" and revises her strategy. Imagine a sequence of processes. Suppose that in process n periods are of duration $1/n$, and the probability of a learn draw is approximately proportional to the period length with proportionality constant equal to 1; that is, the probability of a "learn draw" for a given player when the period length is $1/n$ equals $1/n + o(1/n)$. The limit process found by taking n to ∞ gives the continuous-time process described above with revision opportunities arriving at rate 1.

In summary, time in the model is continuous and the player population is finite. At a random moment τ, a single player is chosen to revise her choice. When chosen, she draws a new action from the probability distribution described by equation 2, where the variable p in that equation is set equal to the fraction of the population currently playing action α.

3 Strategy Revision Processes

The process $\{M_t\}_{t\geq 0}$ which records the number of players choosing action α is called a *strategy revision process.* One consequence of the preceding description of strategy revision is that the number of players choosing a given strategy, say strategy α, can change by at most 1 unit at a time. This is to say, the process $\{M_t\}_{t\geq 0}$ is a birth-death process.

3.1 Invariant Distributions

A birth-death process on the non-negative integers is characterized by a sequence of birth rates $\{\lambda_{m,m+1}\}_{m=0}^{\infty}$ and death rates $\{\mu_{m+1,m}\}_{m=0}^{\infty}$. Given that the process is in state m at time t, the probability of being in $m+1$ at time $t+h$ (a birth) is $\lambda_{m,m+1}h + o(h)$, and the probability of being in $m-1$ at time $t+h$ (a death) is $\mu_{m,m-1}h + o(h)$. It will always be the case that for all $0 \leq m < N$, $\lambda_{m,m+1} > 0$, and for all $0 < m \leq N$, $\mu_{m,m-1} > 0$. All other transition rates are 0. The process is a Markov process which has a unique stationary distribution ρ with support $\{0, \ldots, N\}$. That stationary distribution satisfies the *detailed balance conditions.*

$$\frac{\rho(M)}{\rho(M-1)} = \frac{\lambda_{M-1,M}}{\mu_{M,M-1}} ,$$

implying

$$\frac{\rho(M)}{\rho(0)} = \prod_{m=1}^{M} \frac{\lambda_{m-1,m}}{\mu_{m,m-1}} \tag{3}$$

(See Durrett (1991) and Karlin (1966) for different derivations of this rule.)

The birth and death rates for a stocastic strategy revision process are put together from the choice probabilities and the Poisson alarm clock processes. Suppose that the process is in state M. That is, M players are choosing action α and the remaining players are choosing action ζ. First consider a birth. A ζ player must switch to α. Since strategy revision opportunities come to each player at rate 1 and since players' Poisson alarm clocks are independent, strategy revision opportunities for the collection of $N-M$ players playing ζ at time t arrive at rate $N-M$. That is, in a small interval of time h beginning at time t, the probability that a strategy revision opportunity will arrive to the set of players playing ζ at time t is $(N-M)h + o(h)$ (and the probability of more than one arrival to the group is $O(h^2)$, which we can ignore). The probability that a randomly chosen ζ player will choose α is $\sigma_\alpha\big(M/(N-1)\big)$. Hence the probability that one of the $N-M$ players playing ζ at time t will switch to α before time $t+h$ is $(N-M)\sigma_\alpha\big(M/(N-1)\big)h + o(h)$ (and again the probability of more switches is $O(h^2)$). Consequently the *rate* at which the process will increase is $(N-M)\sigma_a\big(M/(N-1)\big)$. A similar calculation gives the death rates. We thereby derive the following birth and death rates:

$$
\begin{aligned}
M \to M+1 \quad \text{at rate} \quad (N-M)\frac{\exp g\left(\Delta\left(\frac{M}{N-1}\right),\beta\right)}{1+\exp g\left(\Delta\left(\frac{M}{N-1}\right),\beta\right)} \\
M+1 \to M \quad \text{at rate} \quad (M+1)\frac{1}{1+\exp g\left(\Delta\left(\frac{M}{N-1}\right),\beta\right)}
\end{aligned}
\tag{4}
$$

More details on the construction of Poisson processes can be found in Kingman (1993).

Substituting these rates into equation (3) gives the following relations which characterize the invariant distribution:

$$
\ln \frac{\rho_\beta(M)}{\rho_\beta(M-1)} = \ln \frac{N-M+1}{M} + g\left(\Delta\left(\frac{M-1}{N-1}\right),\beta\right)
$$

and

$$
\ln \frac{\rho_\beta(M)}{\rho_\beta(0)} = \ln \binom{N}{M} + \sum_{m=0}^{M-1} g\left(\Delta\left(\frac{m}{N-1}\right),\beta\right). \tag{5}
$$

A convenient way to rewrite equation (5) is to define the function $P_g(M)$ such that

$$
P_g(0,\beta) = 0, \qquad P_g(M,\beta) = \sum_{m=0}^{M-1} g\left(\Delta\left(\frac{m}{N-1}\right),\beta\right). \tag{6}
$$

Then

$$
\ln \frac{\rho_\beta(M)}{\rho_\beta(0)} = \ln \binom{N}{M} + P_g(M,\beta). \tag{7}
$$

The invariant distribution ρ depends upon the parameter β. As β becomes large, the choice probabilities converge to equiprobable draws from the best-response correspondence. If the strategy revision process were constructed with this stochastic choice rule instead of a finite-β rule, it may fail to be irreducible. In a game with multiple Nash equilibria there will be more than one ergodic distribution, each with support on a distinct communication class of the process. For each $\beta < \infty$ there is a unique invariant distribution ρ_β, and as β becomes large ρ_β converges to a particular invariant distribution for the best-response process.

Definition 1. *A state M is said to be* stochastically stable *iff* $\lim_{\beta\to\infty} \rho_\beta(M) > 0$.

If the state $M = N$ is stochastically stable, we will say that α *is stochastically stable,* and if $M = 0$ is stable, we will say that ζ *is stochastically stable.*

In particular, for a given β the log-odds ratios of the invariant distribution are the sum of a combinatorial expression and the value of the potential function. But for large β and functions g satisfying assumptions A.1–3, the P_g term will dominate: As β grows to $+\infty$, $\rho_\beta(K)/\rho_\beta(L)$ will diverge to $+\infty$ for any states K and L such that K is a global maximizer of P_g and L is not. In other words, the *stochastically stable* states are precisely the global maximizers of $P_g(\,\cdot\,, \beta)$. More details about the construction and asymptotic properties of stochastic strategy revision processes can be found in (Blume 1998).

3.2 Interpretation of $\mathbf{P_g}$

A potential for a symmetric two-person binary choice game is a function $Q : \{\alpha, \zeta\} \times \{\alpha, \zeta\} \to \mathbf{R}$ such that for all $y = \alpha, \zeta$,

$$Q(\zeta, y) - Q(\alpha, y) = u_{\zeta y} - u_{\alpha y}.$$

See (Monderer and Shapley 1996) for a discussion of potentials for general N-person games. A potential function for the population game with the same payoff matrix is a function $\widehat{P}$: $\{\alpha, \zeta\}^N \to \mathbf{R}$ such that for any player i and strategy choices s_{-i} for the other players,

$$\widehat{P}(s_{-i}, \zeta) - \widehat{P}(s_{-i}, \alpha) = \frac{1}{N-1} \sum_{j \neq i} U_{\zeta s_j} - U_{\alpha s_j}\,.$$

If a two-player symmetric game has a potential, than so does the corresponding population game.

Lemma 1. *If Q is a potential for the two-player game G, then*

$$\widehat{P}(s) = \frac{1}{2(N-1)} \sum_i \sum_{j \neq i} Q(s_i, s_j)$$

is a potential for the population game.

Proof. A direct calculation verifies the definition.[5] □

For two-player symmetric games we can easily construct a potential by hand. Let $Q(\zeta,\zeta)=0$. Then $Q(\alpha,\zeta)=Q(\zeta,\alpha)=U_{\alpha\zeta}-U_{\zeta\zeta}$, and $Q(\alpha,\alpha)=U_{\alpha\alpha}-U_{\zeta\alpha}+Q(\zeta,\alpha)$.

Notice that the value of $\widehat{P}(s)$ is determined only by the number of players playing α. Let A denote the set of players playing α and B the set of players playing ζ. Then, letting M denote the number of α players and recalling that $Q(\zeta,\zeta)=0$,

$$\begin{aligned}
\widehat{P}(s) &= \frac{1}{2(N-1)}\sum_i\sum_{j\neq i}Q(s_i,s_j)\\
&= \frac{1}{2(N-1)}\sum_i\sum_j Q(s_i,s_j)-\sum_i Q(s_i,s_i)\\
&= \frac{1}{2(N-1)}\Big(\sum_{i\in A}\sum_{j\in A}Q(\alpha,\alpha)+\sum_{i\in A}\sum_{j\in B}Q(\alpha,\zeta)+\sum_{i\in B}\sum_{j\in A}Q(\zeta,\alpha)+\\
&\qquad \sum_{i\in B}\sum_{j\in B}Q(\zeta,\zeta)-\sum_{i\in A}Q(\alpha,\alpha)-\sum_{i\in B}Q(\zeta,\zeta)\Big)\\
&= \frac{1}{2(N-1)}\big(M^2Q(\alpha,\alpha)+2M(N-M)Q(\alpha,\zeta)+(N-M)^2Q(\zeta,\zeta)-\\
&\qquad MQ(\alpha,\alpha)-(N-M)Q(\zeta,\zeta)\big)\\
&= \frac{M}{2(N-1)}\big((M-1)Q(\alpha,\alpha)+2(N-M)Q(\alpha,\zeta)\big)\\
&\equiv P(M)
\end{aligned}$$

Consequently $P(M+1)-P(M)$ is the gain in utility for a ζ player to switch to α when M of his opponents are playing α. Since $\widehat{P}(s)$ and $P(M)$ can each be determined from the other, I will also refer to $P(M)$ as a potential of the population game. Observe from the calculation that $P(0)=0$. It follows from the definition of the potential that

$$P(M+1)-P(M)=\Delta\Big(\frac{M}{N-1}\Big).$$

Suppose that $g(x,\beta)=\beta x$. In other words, the model of noisy choice is the log-linear model. The function $P_g(M,\beta)=\beta P(M)$. In other words, $P_g(M,\beta)/\beta$ is an exact, or cardinal, potential for the population game in the sense of Monderer and Shapley (1996).

According to Monderer and Shapley (1996), an ordinal potential for a game is is a function $\widehat{P}(s)$ such that $\widehat{P}(s_{-i}, \zeta) - \widehat{P}(s_{-i}, \alpha)$ is equal in sign to the utility gain to player i from switching to strategy ζ from strategy α when other players are playing according to strategy profile s_{-i}. For a population game, an ordinal potential satisfies

$$\operatorname{sgn}\big(\widehat{P}(s_{-i}, \zeta) - \widehat{P}(s_{-i}, \alpha)\big) = \operatorname{sgn}\left(\frac{1}{N-1}\sum_{j \neq i} U_{\zeta s_j} - U_{\alpha s_j}\right).$$

Now suppose that $g(x) = \beta h(x)$ and satisfies A.2–3. Then $P_g(M)$ will be an ordinal potential, because

$$P_g(M+1, \beta) - P_g(M, \beta) = \beta h\left(\Delta\Big(\frac{M}{N-1}\Big)\right)$$

and h is sign-preserving. The function P is, up to addition of a constant, the only potential for the population game. However there are many distinct ordinal potentials of the population game, including some which cannot be described as a $P_g(M, \beta)$ for any g satisfying A.1–3. So the ordinal potentials out of which the population dynamics can be built is a strict subset of the set of all ordinal potentials for the population game.

4 Stochastic Stability in Random Matching Games

In random matching environments, players (noisily) best-respond to the play of the entire population. Since players are expected utility maximizers, $\Delta(p)$ is linear in p, the fraction of the population playing strategy α. Since

$$\ln \sigma_\alpha(p)/\sigma_\zeta(p) = g\big(\Delta(p), \beta\big),$$

it follows that

$$\ln \sigma_\zeta(p)/\sigma_\alpha(p) = -g\big(\Delta(p), \beta\big).$$

On the other hand, if only payoff differences, and not the "names" of strategies matter to choice, then the log-odds of choosing ζ over α ought to stand in the same relationship to the payoff difference between ζ and α, which is $-\Delta(p)$. In other words, if only payoff differences matter, then

$$\ln \sigma_\zeta(p)/\sigma_\alpha(p) = g\big(-\Delta(p), \beta\big).$$

Consequently it would follow that $g(\Delta, \beta) = -g(-\Delta, \beta)$ for all Δ in the range of $\Delta(p)$. Since stochastic stability is a large β property, we really need this symmetry not for g but just for the term which is linear in β.

Definition 2. *A noise process $g(x, \beta) = \beta h(x) + r(x, \beta)$ is* skew-symmetric *if for almost all x, $h(x) = -h(-x)$.*

This symmetry property characterizes the largest class of noise processes containing log-linear choice and the mistakes model, and for which, in coordination games, the only stochastically stable state is the risk-dominant strategy. The main results of this section are twofold. First, if the noise process is skew-symmetric, then universal adoption of the risk-dominant strategy is the only stable state in two-by-two coordination games. Second, if the noise process is not skew-symmetric, one can construct a game for which the risk-dominated strategy is also stochastically stable.

Suppose that for some numbers $A, B > 0$, $\Delta(0) = -A$ and $\Delta(1) = B$. This describes the payoff differences of a coordination game, and strategy α is risk dominant if and only if $A < B$. We suppose this is the case unless explicitly stated otherwise.

Theorem 1. *If g is skew-symmetric, then for all large N $\lim_{\beta \to \infty} \rho_\beta(N) = 1$.*

The proof of this theorem and the next make use of the following obvious lemma. Define $H(y) = \int_{-y}^{y} h(x)\, dx$.

Lemma 2. *For almost all $y \in \mathbf{R}$, $H'(y)$ exists and equals $h(y) - h(-y)$. The noise process $g(x, \beta) = \beta h(x) + r(x, \beta)$ is skew-symmetric iff $H(x) \equiv 0$.*

Proof of Theorem 1.

$$
\begin{aligned}
\frac{1}{N\beta}\ln\frac{\rho_\beta(M)}{\rho_\beta(0)} &= \frac{1}{N\beta}\ln\binom{N}{M} + \frac{1}{N\beta}P_g(M,\beta)\\
&= \frac{1}{N\beta}\ln\binom{N}{M} + \frac{1}{N}\sum_{m=0}^{M-1} h\left(\Delta\left(\frac{m}{N-1}\right)\right) + \frac{1}{N\beta}\sum_{m=0}^{M-1} r\left(\Delta\left(\frac{m}{N-1}\right),\beta\right)\\
&\to \frac{1}{N}\sum_{m=0}^{M-1} h\left(\Delta\left(\frac{m}{N-1}\right)\right) \qquad \text{as } \beta\to\infty.
\end{aligned}
$$

The limit term has local maxima at $M = 0$, where its value is 0, and at $M = N$. We need to prove that at $M = N$, the summation is positive. It follows from equations (5) and (7) that for large N, this term is a Riemann-sum approximation to the integral $\int_0^1 h\big(\Delta(p)\big)\,dp$. Thus it suffices to show that the value of this integral is positive. To see this, carry out a change of variables.

$$
\begin{aligned}
\int_0^1 h\big(\Delta(p)\big)\,dp &= (B+A)^{-1}\int_{-A}^{B} h(x)\,dx\\
&= (B+A)^{-1}H(A) \;+\; (B+A)^{-1}\int_{A}^{B} h(x)\,dx
\end{aligned}
$$

According to Lemma 2, the first term in the sum is 0. The second term is positive because $h(x)$ is positive for positive x. □

All skew-symmetric noise processes give the same stochastic stability result in two-by-two coordination games. The converse is also true.

Theorem 2. *If g is not skew-symmetric, there is a payoff difference function $\Delta(p)$ for some coordination game such that, with noise process g, for all large N $\lim_{\beta\to\infty}\rho_\beta(0) = 1$.*

Proof. Suppose that g is not skew-symmetric. Then there is an x such that $H(x) \neq 0$. If $H(x) < 0$, then choose $A = x$ and $B = x + \epsilon$ for small enough ϵ that $H(x) + \int_x^{x+\epsilon} h(y)\,dy < 0$. For this game α is risk dominant. Making use of the Riemann sum approximation of Theorem 1's proof, we see that for large N, $(N\beta)^{-1}\ln\rho_\beta(N)/\rho_\beta(0) < 0$, and so only $M = 0$ is stochastically stable. If $H(x) > 0$. Choose $A = x$ and $B = x - \epsilon$ for small ϵ as above. In this case ζ is risk dominant, but $M = N$, all play α, is the only stochastically stable state. □

If the noise process is systematically biased, then the degree of bias determines the size of the set of games for which $M = 0$ is stochastically stable. Suppose for instance that $\tilde{h}(x)$ is skew-symmetric, and that $k(x) < 0$ for $x < 0$. Let $h(x) = \tilde{h}(x)$ for $x \geq 0$ and $\tilde{h}(x) + k(x)$ for $x < 0$. Thus errors downward are more likely than errors upward. Then

$$\begin{aligned}
\int_{-A}^{A} h(x)\,dx &= \int_{-A}^{0} h(x)\,dx + \int_{0}^{A} h(x)\,dx \\
&= \int_{-A}^{0} \tilde{h}(x)\,dx + \int_{-A}^{0} k(x)\,dx + \int_{0}^{A} \tilde{h}(x)\,dx \\
&= \int_{-A}^{0} k(x)\,dx \\
&\equiv -K < 0.
\end{aligned}$$

Then $M = N$ is not stochastically stable for large N if $\int_A^B h(x)\,dx < K$, and it is stochastically stable for large N if $\int_A^B \tilde{h}(x)\,dx > K$. Since this integral is increasing in B, for given A there is a $\hat{B}$ such that if $B < \hat{B}$, $M = 0$ alone is stochastically stable for large N. Now consider a noise process g' even more biassed in the sense that $h'(x) = \tilde{h}(x)$ for $x \geq 0$, $\tilde{h}(x) + k'(x)$ for $x < 0$, and such that for all $x < 0$, $k'(x) < k(x)$. Then $\int_{-A}^{0} k'(x) = -K' < -K$, and so the corresponding $\hat{B}'$ will exceed $\hat{B}$. If $M = N$ fails to be stochastically stable under g, then it will also fail under g', while the converse need not hold. For more general asymmetries, however, there is little one can say about the set of games for which 0 is stochastically stable.

The following variation on the errors model illustrates how the remainder term $r(x, \beta)$ can arise.

Example 1: A mistake when strategy α is the best response occurs with probability ϵ, while a mistake when strategy ζ is the best response occurs with probability $k\epsilon$. The birth and death rates become

$$\begin{aligned}
M \to M+1 &\quad \text{at rate} \quad (N - M)(1 - \epsilon) \\
M + 1 \to M &\quad \text{at rate} \quad (M + 1)\epsilon
\end{aligned}$$

when α is a best response, and

$$\begin{aligned} M \to M+1 &\quad \text{at rate} \quad (N-M)k\epsilon \\ M+1 \to M &\quad \text{at rate} \quad (M+1)(1-k\epsilon) \end{aligned}$$

when ζ is a best response.

Take $h(\Delta) = \operatorname{sgn}(\Delta)$, and let

$$r(\Delta, \beta) = \begin{cases} 0 & \text{if } \Delta > 0, \\ \dfrac{1 - k\epsilon(\beta)}{\beta k\bigl(1 - \epsilon(\beta)\bigr)} & \text{if } \Delta < 0. \end{cases}$$

where $\epsilon(\beta) = (1 + \exp \beta)^{-1}$. The noise process $g(\Delta) = \beta h(\Delta) + r(\Delta, \beta)$ has the property that the odds ratio of errors when α is the best response to errors when ζ is the best response is bounded. Now $0 < \beta r(\Delta, \beta) < k^{-1}$, so the stochastic stability results are driven by h alone. In this case as $\beta \uparrow \infty$, mass piles up on the risk-dominant state: $\rho_\beta(N) \to 1$. □

This example shows how the stochastic stability results are robust to small changes in the odds of mistakes. This is the obvious result that only if the odds of a mistake in favor of the risk-dominated strategy become infinitely more likely than the odds of a mistake in favor of the risk-dominant strategy as β grows will the risk-dominant selection result be lost.

The next example shows what happens when the arrival rate of strategy revision opportunities to a player depends upon the strategy she currently employs.

Example 2: Suppose that strategy revision opportunities arrive at rate 1 to all players whose current choice is ζ, and at rate γ to all players currently choosing α. The birth and death rates for the strategy revision process are given by

$$M \to M+1 \quad \text{at rate} \quad (N-M)\gamma^{N-M} \frac{\exp \beta h\Big(\Delta\Big(\dfrac{M}{N-1}\Big)\Big)}{1+\exp \beta h\Big(\Delta\Big(\dfrac{M}{N-1}\Big)\Big)}$$

$$M+1 \to M \quad \text{at rate} \quad (M+1)\frac{1}{1+\exp \beta h\Big(\Delta\Big(\dfrac{M}{N-1}\Big)\Big)}$$

where h is skew-symmetric, non-decreasing and increasing at 0. Let $r(x, \beta) = \gamma^{N-M}$ where $M = \lfloor (N-1)\Delta^{-1}(x) \rfloor$, the largest integer not greater than $(N-1)\Delta^{-1}(x)$. Let $g(x, \beta) = \beta h(x) + r(x, \beta)$. The conditions of Theorem 1 are met. Thus for large N, the state N wherein all play the risk-dominant strategy, will be the only stochastically stable state. In fact a tedious calculation shows that

$$\ln \frac{\rho_\beta(M)}{\rho_\beta(0)} = \ln \binom{N}{M} + \frac{(M+1)(2N-M)}{2} \ln \gamma + \beta P_g(M).$$

For large β this expression is dominated by the term $\beta P_g(M)$, and so the stochastic stability results are independent of γ. $\square$

In the model, the function g describes individual choice behavior. The aggregate birth-death chain is constructed from individual choice behavior, characterized by g and and the parameter β, and a process for assigning strategy revision opportunities. The conventional models have constant revision opportunity arrival rates. Here we have endogenized them by making the rate for a given player depend upon his current action. It is possible to allow for variation based on unobserved individual characteristics as well, but somewhat different techniques are required.

5 Population Games

For pairwise matching processes in two-by-two games, the stochastic stability results have been shown to be invariant to the noise process so long as the noise satisfies a basic symmetry requirement. This section extends the analysis of sections 3 and 4 to a larger class of *population games*. Population games model many important phenomena, including local public goods, clubs, and coordination problems of the kind which arise in macroeconomics and in the industrial organization of markets with many firms.

In section 3 we developed the payoff difference function $\Delta(p)$ by assuming that players valued each alternative according to its expected return against an opponent drawn at random from the population. The subsequent analysis depended only upon specific properties of the function $\Delta(p)$ and not upon the underlying random matching story. Other interaction stories will presumably generate different $\Delta(p)$ functions. Consequently, it is interesting to ask what general properties of the $\Delta(p)$ function will preserve the conclusions of Theorems 1 and 2.

For instance, suppose that action α is joining a network, and action ζ is staying out. The value of staying out is set at 0. The benefits of joining the network are a function of the proportion of the population p already signed up, and is described by the function $\Delta(p)$. Now the evolutionary process models the dynamics of group membership. A typical hypothesis about the returns to joining a network is that there are initially increasing returns to network size, but ultimately constant or decreasing returns sets in. Thus $\Delta(p)$ would be initially increasing and s-shaped. (Blume (1994) uses a construction like this to study the dynamics of investment cycles due to coordination failure.)

Linearity of $\Delta(p)$ is crucial to the proofs of Theorems 1 and 2, and so it should not be surprising that without linearity invariance typically fails even among skew-symmetric noise processes. This section illustrates invariance failure and then establishes a condition on $\Delta(p)$ which implies outcome invariance for skew-symmetric noise processes.

The following example demonstrates invariance failure.

Example 3: Let $\Delta(p) = -A + (B + A)p^2$ where A and $B + A$ are positive. The payoff difference function $\Delta(p)$ is increasing in p, creating a strategic complementarity. Suppose first that $g(x) = \text{sgn}(x)$, the mistakes model. The root of $\Delta(p)$ is $p^* = \sqrt{A/(B + A)}$, and α is preferred to ζ if and only if $p > p^*$. In the mistakes model α will be the unique stochastically stable outcome if and only if $p^* < 1/2$, that is, if and only if $3A < B$.

For the log-linear model, the stochastic stability criterion is different. Suppose N is large, so that the Riemann sum approximation is good. Then α will be the only stochastically stable outcome under the log-linear noise process if and only if $\int_0^1 g\big(\Delta(p), \beta\big)\, dp > 0$. Calculation shows that the integral will be positive if and only if $2A < B$. Thus when $B > 3A$, α is the unique stochastically stable outcome under both processes. When $B < 2A$, ζ is the only stochastically stable outcome under both processes. And when $2A < B < 3A$, α is stochastically stable under the log-linear process but not under the mistakes process. $\square$

Although invariance fails for some parameter values in this example, invariance does hold for a class of payoff difference functions that includes more than just linear functions. Let $p(\Delta)$ denote the inverse of $\Delta(p)$ when it exists. Proofs can be found at the end of the section.

Theorem 3. *Suppose that* $\Delta(p)$ *is strictly increasing and bounded, and that* $\limsup_{p<1} \Delta(p) > \max\{0, -\Delta(0)\}$. *Then for large enough* N*:*

1. *There is a skew-symmetric noise process* g *such that* $M = N$ *is the only stochastically stable state.*

2. *If* $p(0) > 1/2$, *then there is a skew-symmetric noise process for which* $M = N$ *is not stochastically stable.*

Part 2 of the Theorem states that $p(0) \leq 1/2$ is necessary for skew-symmetry to imply that $M = N$ is stochastically stable. When $\Delta(p)$ is linear, $p(0) \leq 1/2$ is also sufficient for $M = N$ to be a stochastically stable state for all skew-symmetric noise processes (and uniquely so if $p(0) < 1/2$). But in general the $p(0) \leq 1/2$ condition is only necessary, as the following example shows:

Example 4: Let

$$\Delta(p) = \begin{cases} 3p - 1 - \epsilon & \text{if } 0 \leq p < 1/3, \\ (21p - 25)\epsilon/18 & \text{if } 1/3 \leq p < 11/12, \\ (24 - 12\epsilon)p - 22 + 12\epsilon & \text{if } 11/12 \leq p \leq 1. \end{cases}$$

where $\epsilon > 0$ is very small, and take the log linear rule, $g(x, \beta) = \beta x$. Then $\int_0^1 g\big(\Delta(p), \beta\big) dp \approx -1/12$, and so for large enough N, $M = 0$ will be the only stochastically stable state. □

The next Theorem gives some sufficient conditions for invariance of the stochastically stable set to skew-symmetric noise processes.

Theorem 4. *Suppose that $\Delta(p)$ is strictly increasing and C^1, and that $\Delta(1) > -\Delta(0)$. Under any of the following conditions, for every skew-symmetric noise process and for large enough N, only the state $M = N$ will be stochastically stable.*

1. *$\Delta(p)$ is concave,*
2. *$\Delta(p)$ is convex, $p(0) < 1/2$ and h is differentiable on $\big(0, \Delta(1)\big)$.*
3. *$\Delta(p)$ is skew-symmetric around some p^*; that is, there is a p^* such that for all ρ such that $0 \leq p^* - \rho \leq p^* + \rho \leq 1$, $\Delta(p^* - \rho) = -\Delta(p^* + \rho)$.*

In Example 3, $\Delta(p)$ is convex. The condition $3A < B$ guarantees that $p^* = p(0) < 1/2$, so all skew-symmetric rules have the same unique stochastically stable state $M = N$. Alternatively, if $3A < B$, then the conclusion holds for the mistakes model, and so it holds for all models.

Proof of Theorem 3. (1) Let $r(\Delta, \beta) \equiv 0$. Define $h(\Delta)$ skew-symmetric and such that

$$h(\Delta) = \begin{cases} \epsilon & \text{if } 0 \leq \Delta \leq -\Delta(0), \\ 1 & \text{if } -\Delta(0) < \Delta. \end{cases}$$

Let p^* denote the root of $\Delta(p) = 0$ if it exists, and 0 otherwise. The assumptions imply that there is a δ between $-\Delta(0)$ and $\limsup_{p<1} \Delta(p)$, and a p' strictly between p^* and 1 such that

$\Delta(p') \geq \delta$. Then

$$\begin{aligned}\int_0^1 h\big(\Delta(p)\big)\,dp &\geq -\epsilon p^* + (p' - p^*)\epsilon + (1 - p')\delta \\ &= (p' - 2p^*)\epsilon + (1 - p')\delta\end{aligned}$$

Since the last term in the equality is strictly positive, the whole expression can be made positive by choosing ϵ sufficiently small. The Riemann approximation argument used in proving Theorems 1 and 2 prove the claim.

(2) A calculation shows that $M = 0$ will be the unique stochastically stable state for the mistakes noise process. $\square$

Proof of Theorem 4. As in the proofs of Theorems 1 and 2, the remainder term r can be neglected, and we need to show that $\int_0^1 h\big(\Delta(p)\big)dp > 0$. Let $P(z) = p(z) + p(-z) - 2p(0)$. Computing,

$$\begin{aligned}\int_0^1 h\big(\Delta(p)\big)dp &= \int_{\Delta(0)}^{\Delta(1)} h(z)p'(z)dz \\ &= \int_{\Delta(0)}^{0} h(z)p'(z)dz + \int_0^{-\Delta(0)} h(z)p'(z)dz + \int_{-\Delta(0)}^{\Delta(1)} h(z)p'(z)dz \\ &= -\int_{\Delta(0)}^{0} h(-z)p'(z)dz + \int_0^{-\Delta(0)} h(z)p'(z)dz + \int_{-\Delta(0)}^{\Delta(1)} h(z)p'(z)dz \\ &= \int_{-\Delta(0)}^{0} h(z)p'(-z)dz + \int_0^{-\Delta(0)} h(z)p'(z)dz + \int_{-\Delta(0)}^{\Delta(1)} h(z)p'(z)dz \\ &= \int_0^{-\Delta(0)} h(z)\big(p'(z) - p'(-z)\big)dz + \int_{-\Delta(0)}^{\Delta(1)} h(z)p'(z)dz \qquad (8) \\ &= h\big(-\Delta(0)\big)P\big(-\Delta(0)\big) - \int_0^{-\Delta(0)} h'(z)P(z)dz + \int_{-\Delta(0)}^{\Delta(1)} h(z)p'(z)dz \\ &\geq h\big(-\Delta(0)\big)\big(1 - 2p(0)\big) - \int_0^{-\Delta(0)} h'(z)P(z)dz \qquad (9)\end{aligned}$$

The first line is a change of variables. The third line comes from skew symmetry. The forth line is another change of variables. The sixth line is an integration by parts. Equation (9) holds because, since $h(z)$ is non-decreasing and $p'(z)$ is non-negative, the last integral on the preceding line is bounded below by $h\big(-\Delta(0)\big)\Big(p\big(\Delta(1)\big) - p\big(\big(-\Delta(0)\big)\Big)$.

(1) Consider now the last equality. If $\Delta(p)$ is concave, $p'(z) - p'(-z) \geq 0$ for $z \geq 0$. The function $h(z)$ is non-negative for $z \geq 0$, and the second integral in (8) is strictly positive.

(2)If $\Delta(p)$ is convex, then $P(z) \leq 0$ and so from (9) it follows that the integral is no more than $h\big(-\Delta(0)\big)\big(1 - 2p(0)\big)$. If $p(0) < 1/2$ the integral will be positive.

(3) $h\big(\Delta(p^* - \rho)\big) = -h\big(\Delta(p^* + \rho)\big)$. Consequently the integral will be positive so long as $p^* < 1/2$, which with skew-symmetry of $\Delta(p)$, is implied by the hypothesis that $\Delta(1) > -\Delta(0)$. □

Observe from (9) that if $\Delta(p)$ is strictly convex and $h(z)$ is increasing, then $p(0)$ can be slightly greater than 1/2 and $M = N$ will still be the only stochastically stable state.

6 Random Utility Models

Random utility models are one source for the kind of stochastic choice behavior postulated in evolutionary models. In these models the utilities of the choices α and ζ are random variables with fixed means given by the payoff matrix. The actual utility realizations are observations of each alternative's random utility, and the chosen alternative is that with the higher actual value. These observations are independent across objects of choice, choice opportunities, and, in games, players. Let ϵ_α and ϵ_ζ be i.i.d. mean 0 random variables, and suppose that the cumulative distribution of $\epsilon_\zeta - \epsilon_\alpha$ is F. Then the probability that α will be chosen when the mean payoff difference is $\Delta(p)$ is $F\big(\Delta(p)\big)$ — the probability that the utility to α, $U_\alpha(p) + \epsilon_\alpha$ exceeds the utility to ζ, $U_\zeta(p) + \epsilon_\zeta$ (where $U_\alpha(p)$ and $U_\zeta(p)$ are the mean payoffs to α and ζ, respectively, when the population fraction choosing α is p). Then

$$\ln \sigma_\alpha(p) - \ln \sigma_\zeta(p) = \ln F\big(\Delta(p)\big) - \ln\Big(1 - F\big(\Delta(p)\big)\Big).$$

The applicability of the last section's analysis depends upon how the noise is parametrized. The random utility interpretation for the log-linear model is well known (see Brock and Durlauf, 1999).

The mistakes model, on the other hand, is not a random utility model. For any given game, a random utility model can be written down that generates the mistakes model for that particular game. But there is no random utility model which generates the mistakes model in every game. Consequently, random utility models are not a generic way to generate the mistakes model of noisy choice. To see the first claim, suppose that observations on the utility of the two alternatives are independent. With probability $1 - 2\epsilon$ the player correctly observes the utility of an alternative. With probability 2ϵ they observe the utility plus a random variable $\tilde{x}$, which is takes on the values x_H and x_L with probability 1/2 each. The values are such that $x_1 > \Delta(1)$ and $x_2 < \Delta(0)$. Then with probability $\epsilon(1 - \epsilon)$ the observed utility difference will be incorrect, and it will be correct with the complementary probability. To see the second claim, observe that in any random utility model which gives rise to the mistakes model in any game $\Delta(p)$ with error probability ϵ,

$$\ln \frac{F\big(\Delta(p)\big)}{1 - F\big(\Delta(p)\big)} = \operatorname{sgn}\big(\Delta(p)\big) \ln \frac{1-\epsilon}{\epsilon}$$

and so

$$F\big(\Delta(p)\big) = \frac{(1-\epsilon)^{\operatorname{sgn}\left(\Delta(p)\right)}}{\epsilon^{\operatorname{sgn}\left(\Delta(p)\right)} + (1-\epsilon)^{\operatorname{sgn}\left(\Delta(p)\right)}}$$

Since $\operatorname{sgn}\big(\Delta(p)\big)$ can be only -1, 0 or 1, $F(x)$ can take on only three values, none of which is 0 or 1. Consequently F is not a cumulative distribution function for any random variable, which contradicts the random utility hypothesis.

Other parametrizations of noise reduction are possible, and they require a different analysis. One natural parametrization is linear rescaling. Suppose that $\epsilon_\zeta - \epsilon_\alpha$ is distributed according to cdf F, whose support is all of $\mathbf{R}$. Suppose now that for parameter value β choice α is chosen if and only if $U_\alpha(p) + \epsilon_\alpha/\beta > U_\zeta(p) + \epsilon_\zeta/\beta$. As β becomes large, the influence of the random term disappears and the choice rule converges to best-response. For arbitrary β,

$$\ln \sigma_\alpha(p) - \ln \sigma_\zeta(p) = \ln F\big(\beta\Delta(p)\big) - \ln\Big(1 - F\big(\beta\Delta(p)\big)\Big). \tag{10}$$

Stochastic stability results will depend upon the behavior of the tails of F.

Theorem 5. *Suppose that the log-odds of choice are given by equation* (10)*. Suppose that* $\Delta(p)$ *is strictly increasing and that* $\sup_p\{\Delta(p) \leq 0\} < 1/2$*. For any cdf* F *satisfying the conditions:*

1. *For all* $x > 0$*,* $1 - F(x) + F(-x) > 0$*, and*
2. $\liminf_{x\to\infty} \dfrac{F(-x)}{\big(1 - F(x)\big) + F(-x)} > 0$ *and* $\limsup_{x\to\infty} \dfrac{F(-x)}{\big(1 - F(x)\big) + F(-x)} < 1.$

For all large N*, the state* $M = N$ *will be the unique stochastically stable state.*

The condition of the theorem bounds away from 0 the lim inf of the conditional probability of an observation occurring in a particular tail, upper or lower, given that it is in the tails. Notice that the stated conditions for the lower tail implies the same conditions for the upper tail. It is possible to construct mean 0 random variables failing this condition. The proof makes it clear that the condition is not necessary. The Riemann sum argument of the previous section no longer works because the integrand changes with β. The proof instead relies on a direct calculation.

Proof. Suppose that α is risk-dominant, and let k^* denote the largest non-negative M such that the utility difference $\Delta\big(M/(N-1)\big)$ is less than or equal to 0, if it exists, and 0 otherwise. The integer k^* must be no greater than p^*N where p^* is the root of $\Delta(p)$, and of course $p < 1/2$. We suppose that N is large enough that $k^* \leq (N/2) - 1$.

A calculation shows that

$$\ln \frac{\rho(M)}{\rho(0)} = \ln \binom{N}{M} + \sum_{m=1}^{M} \ln F\bigg(\beta\Delta\Big(\frac{m-1}{N-1}\Big)\bigg) - \ln\bigg\{1 - F\bigg(\beta\Delta\Big(\frac{m}{N-1}\Big)\bigg)\bigg\}$$

for $M \geq 1$. Since $\Delta(p)$ is strictly increasing, it clear that the sum is either U-shaped or strictly increasing as a function of M. It suffices to show that $\lim_{\beta\to\infty} \ln\big(\rho(N)/\rho(0)\big) = \infty$. Write this log-odds ratio as

$$\sum_{k=0}^{k^*} \ln \frac{F\Big(\beta\Delta\big(\frac{k^*-k}{N-1}\big)\Big)}{\Big\{1 - F\Big(\beta\Delta\big(\frac{k^*+k}{N-1}\big)\Big)\Big\}} + \sum_{k=0}^{k^*} \ln \frac{F\Big(\beta\Delta\big(\frac{k^*+k}{N-1}\big)\Big)}{\Big\{1 - F\Big(\beta\Delta\big(\frac{k^*-k}{N-1}\big)\Big)\Big\}} +$$

$$\sum_{k=2k^*+1}^{N-1} \ln \frac{F\Big(\beta\Delta\big(\frac{k}{N-1}\big)\Big)}{\Big\{1 - F\Big(\beta\Delta\big(\frac{k}{N-1}\big)\Big)\Big\}}$$

(The reversal of signs in the argument to Δ between the numerator and denominator of the first two terms is intentional.) The third term grows to $+\infty$ as β grows, and the middle term converges to 0. Only the first term needs to be bounded, and this is accomplished by one of the lim sup and one of the lim inf conditions together with the hypothesis that the probability of being in the tails is always positive. □

The condition of Theorem 5 requires that one tail of F not be too skewed relative to the other. The condition is satisfied if $\epsilon_\zeta - \epsilon_\alpha$ is symmetric. For this to be true it suffices that ϵ_α and ϵ_ζ be identically distributed but not necessarily symmetrically distributed, or that they be symmetrically distributed but not necessarily identically distributed.

The following counterexample demonstrates why control of the tails is necessary in Theorem 5.

Example 5: Suppose that ϵ_α is surely 0 and ϵ_ζ is distributed with cdf

$$F(x) = \begin{cases} |x|^{-2}/2 & \text{if } x < -1, \\ 1/2 & \text{if } -1 \leq x < 1, \\ 1 - (x^{-1}/2) & \text{if } 1 \leq x. \end{cases}$$

For the random variable ϵ_ζ the probability of being in the lower tail conditional upon being in any tail converges to 0. Calculations show that $\rho(N)/\rho(0)$ converges to 0. □

This example works, as does the Bergin and Lipman (1996) construction, because the odds ratios of different kinds of errors are not bounded.

7 Conclusion

There are many stories behind the "unmodeled noise" which drives stochastic stability results, including random utility, experimentation, and unobserved characteristics. These various processes are applied to a variety of games. Although we explicitly do not want to concern ourselves with the fine details of how the noise is generated, we do need to ask under what kinds of general noise process the usual stochastic stability results are robust. Rather than focus directly on the trembles which are themselves derived from some noise process, this paper directly addresses the source of the noise.

This paper has argued several points. Most important, the conclusions of the so-called evolutionary stochastic stability results, while not robust to all manner of mutation or noise, are robust to a large class of such processes. Moreover, this class has an intuitively appealing description. It is the class of noise or mutation processes for which labels do not matter. The probability of a deviation from the best response may depend upon the payoff difference between the two strategies, who is choosing and what she is currently playing. But it may not depend upon the name of the best response. That is, the probability of player 1 choosing ζ when α has a payoff advantage of Δ and she is currently choosing α is the same as that of her choosing ζ when she is currently playing α and ζ has a payoff advantage of Δ.[6]

Second, the robust stochastic stability results of pairwise matching models extend to a larger class of games, which are here called *population games.* The robustness criterion for noise processes, skew-symmetry, was first derived for the class of games in which payoff difference functions are linear in the distribution of play in the population. This class of games includes the random matching paradigm of much biologically-motivated evolutionary game theory, but little of economic interest. Section 5 introduces a class of coordination games which model such phenomena as local public goods, clubs, and coordination problems of the kind which arise in macroeconomics and in the industrial organization of markets with many firms. The robust selection results for skew-symmetric noise processes carry over to this richer setting, provided the non-linearities of the payoff difference functions do not themselves introduce large asymmetries into the strategy selection process.

Third, this paper argues for an alternative technology to the complex computations required by the discrete-time versions of the Friedlin-Wentzel/Kifer techniques which have become popular in the analysis of population games. This alternative technology has two components. The first component is the exploitation of simple Markov processes. By supposing that only one player has a strategy revision opportunity at any given time, the Markov process on states becomes a birth-death process. Single-type birth-death processes are among the easiest Markov processes to study. Not only is it trivial to compute stationary distributions, but expected first-passage times and recurrance intervals are easily computed as well. Continuous time jump processes are a natural way to describe the onset of discrete cognitive events for individual agents. Furthermore, the continuous-time birth-death processes are a natural limit of the more usual discrete time processes as the length of a period grows small. But the choice of continuous time is only a convenience; these models can just as easily be built and studied in discrete time.

The other component is to exploit the fact that all two-by-two symmetric games (and most other games for which solid stochastic stability results exist) are in fact potential games. The strategy revision process for the log-linear choice model turns out to be stochastic hill climbing on the graph of the potential; hence the selection result that when noise is small, the process accumulates on the global maximum. This phenomenon was first observed in Blume (1993) and provides a justification for Monderer and Shapley's (1993) conjecture that selecting the global maxima of the potential in a potential game might prove to be an interesting equilibrium refinement. Other noise processes correspond to the choice of other ordinal potential functions. The potential function is unique up to the addition of a constant, but there can be many quite distinct ordinal potentials. The hypotheses of the robust stochastic stability results presented here can be viewed as the construction of a class of ordinal potentials that all have the same global maximum as the exact potential.

The potential function approach gives information about other quantities of interest. For instance, moments of the distribution of the first passage time between any two states can be computed directly. For a given noise process — a given ordinal potential — the log of the expected first-passage time from one equilibrium to the other is, for large β, on the order of the difference between the value of the ordinal potential at the first equilibrium (a local maximum) and the minimum of the ordinal potential on the interval between the two equilibria. Of course ordinal potential games are a small class of games, so I am not offering these techniques as general solutions to the problems of evolutionary dynamics. But they do extend to ordinal potential games with more than two strategies and also to some asymmetric two-strategy games. Moreover, as Rosenthal (1973) has shown, many economic models naturally set up as potential games. Thus the importance of the class is not connected to its size.

Notes

[1] "Current beliefs" are constructed from the aggregate behavior of the population, and there are different ways to do that. Blume (1993) and Kandori, Mailath, and Rob (1993) take beliefs to be the current distribution of play. A richer cognitive model can be found in (Young 1993). Of course, still other dynamics not involving best response at all, such as imitation dynamics, have also been studied.

[2]Other alternatives in the literature include having players best respond to some beliefs that are derived from the history of play, as in (Young 1993), and having players consider the future rather than just responding myopically to current beliefs as in (Blume 1995). Yet other alternatives are possible.

[3]In populations of myopic players interacting through pairwise random matches, the return to playing a given strategy in a future match given a player's beliefs p is just the expected value of the payoffs from that strategy with respect to probability distribution p.

[4]Some models which are not posed in this manner have equivalent formulations which are.

[5]In fact this lemma will remain true in any local interaction game where the neighborhood relation is symmetric in the sense that i is a neighbor of j if and only if j is a neighbor of i. See (Blume 1993) for an application to lattice games and (Blume 1998) for a general stochastic stability result.

[6]Even this is an overstatement. Labels may matter in the $r(x, \beta)$ terms that are "lower order" with respect to β than the $\beta h(x)$ term. Also, as we have seen, the arrival rate of revision opportunities may be label-dependent, and more generally, state dependent in a way that does not involve β.

References

Bergin, J., and B. Lipman (1996). "Evolution with State-Dependent Mutations," *Econometrica*, 64, 943–56.

Blume, L. E. (1993). "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 4, 387–424.

——— (1994). "An Evolutionary Analysis of Keynsian Coordination Failure," Cornell University.

——— (1995). "Evolutionary Equilibrium with Forward-Looking Players," Cornell University.

——— (1998). "Population Games," in *The Economy as an Complex Adaptive System II*, ed. by W. B. Arthur, D. Lane, and S. Durlauf. Addison-Wesley.

Brock, W., and S. Durlauf (1995). "Discrete Choice with Social Interactions I: Theory," SSRC Working Paper No. 9521, University of Wisconsin.

——— (1999). "Interactions Based Models," SSRC Working Paper No. 9910, University of Wisconsin.

Canning, D. (1992). "Average Behavior in Learning Models," *Journal of Economic Theory*, 57, 442–472.

Durrett, R. (1991). *Probability: Theory and Examples*. Wadsworth & Brooks, Pacific Grove, CA.

Foster, D., and H. P. Young (1990). "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology*, 38, 219–232.

Fudenberg, D., and C. Harris (1992). "Evolutionary Dynamics with Aggregate Shocks," *Journal of Economic Theory*, 57(2), 420–41.

Kandori, M., G. Mailath, and R. Rob (1993). "Learning, Mutation and Long Run Equilibrium in Games," *Econometrica*, 61, 29–56.

Karlin, S. (1966). *A First Course in Stochastic Processes.* Academic Press, New York.

Kingman, J. F. C. (1993). *Poisson Processes.* Oxford University Press, New York.

Monderer, D., and L. Shapley (1996). "Potential Games," *Games and Economic Behavior*, 1, 124–43.

Rosenthal, R. (1973). "A Class of Games Possessing Pure Strategy Nash Equilibria," *Int. J. Game Theory*, 2, 65–67.

Samuelson, L. (1994). "Stochastic Stability in Games with Alternative Best Replies," *Journal of Economic Theory*, 64, 35–65.

Young, H. P. (1993). "The Evolution of Conventions," *Econometrica*, 61, 57–84.