HOW NOISE MATTERS

Lawrence E. Blume

June 1994

This research is 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 for helpful discussions and to Murali Agastya for prodding him into writing this paper and for comments.

Abstract

Recent advances in evolutionary game theory have employed stochastic processes of noise in decisionmaking to select in favor of certain equilibria in coordination games. 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, 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 reduction 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 selection results of the evolutionary literature.

JEL Classification: C78

Correspondent:

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

Printed: July 27, 1994

"A little less noise there, a little less noise!"
Father Darling in J. M. Barrie's *Peter Pan*

1. Introduction

In recent years a dynamic analysis of evolution in strategic environments has emerged which demonstrates the importance of small amounts of random noise in driving population behavior towards particular equilibria and away from others. Since the path-breaking work of Foster and Young (1990), and of Kandori, Mailath and Robb (1993) and Young (1993) (hereinafter referred to as KMRY), selection in favor of risk-dominant equilibria has been uncovered repeatedly in a variety of strategic settings.

In the development of evolutionary models we take seriously the structure of the payoffs and the details of the interaction technology. We do not agonize over the specification of the noise. In fact, the power of the evolutionary approach comes from the fact that the noise process does not need to be explicitly modelled. Or does it? Are the equilibrium selection results in two-by-two matching games invariant to the choice of a noise process? And what happens in other strategic environments? This paper answers these questions.

The news is mixed. The good news is that so 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 when randomly chosen members of the population are matched to play a symmetric two-by-two game. Moreover, this symmetry condition turns out to be sufficient for invariance results even in strategic environments 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.

These results complement the related work by Bergin and Lipman (1994). They demonstrate that, without any structure on mutations in evolutionary models of the KMRY variety (their setting is in fact even more general) any equilibrium can be selected for by a suitable specification of noise and noise reduction. How general, then, are the known selection results? I answer that question here.

2. Matching Processes

The most popular setting for evolutionary analysis is a population of players repeatedly playing against each other. The basic setup, devoid of any motivation, goes like this. At each tick of the clock, players are chosen at random to reconsider their strategic choice. The probability that any given player will have a strategy revision opportunity in any given period is p, independently of the availability of strategy revision opportunities to other players. When a player has a strategy revision opportunity, she will with high probability choose to play a best response to the aggregate play of the population. With low probability she will do something else. What is this probability? This is the unmodelled noise. One popular process is the "mistakes" process of KMRY wherein, regardless of payoff, she best-responds with fixed probability $1-\epsilon$ and "errs" with probability ϵ. A proper noise process is the "log-linear" process examined in Blume (1993). There, the log-odds ratio of α versus ζ is proportional to the difference in payoffs between α and ζ. Consequently if the payoff to α exceeds that to ζ, then as β grows large the odds ratio of choosing α to choosing ζ grows to $+\infty$.

These processes of randomly choosing players who then randomly choose a strategy creates a Markov process on the unit interval which measures the fraction of the population choosing, say, strategy α. This Markov process will have a unique invariant distribution whose features depend both on the payoffs and on the unmodelled noise. Now we, the modellers, perform the experiment of shrinking the noise, for instance by letting ϵ go to 0 in the mistakes model or letting the constant of proportionality go to infinity in the log-linear model. When these experiments are performed, the invariant distributions will tend to some limit, typically a point mass on some population fraction. When the two-by-two game is a pure coordination game with a unique risk-dominant strategy, for both noise processes the invariant distributions converge to point mass on all of the population choosing the risk-dominant strategy. This is the equilibrium selection result of the new evolutionary analysis. In this section I investigate the robustness of this result with respect to the specification of the noise process.

The evolutionary processes of strategy revision has two parts: The description of individual choice behavior and the construction of the popupulation process. I begin with the former. Given a payoff matrix for a symmetric two-strategy game, derive the payoffs $U_\alpha(p)$ and $U_\zeta(p)$ for strategies α and ζ as a function of p, the likelihood that the player assigns to playing against strategy α. In the matching processes described by KMRY, Blume (1993) and others, p is the fraction of the players opponents choosing α. The difference in payoff between strategy α and ζ, $\Delta(p) = U_\alpha(p) - U_\zeta(p)$, is a linear function of p. In a pure coordination game $\Delta(p)$ will be increasing. Strategy α will be risk-dominant if and only if $\Delta(p) = 0$ for some $p < 1/2$ or $\Delta(0) > 0$. (I shall completely ignore this last case because it is obvious that almost any noise process will select away from strictly dominated strategies.)

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 α is p. I will model general noise processes with shrinkage parameter β by supposing that

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

I will assume that the function g is non-decreasing and that $g(0) = 0$. These assumptions capture the following phenomena: That the greater the payoff advantage to α, the more likely it is to be chosen, and that the noise process is *unbiased,* — when the two strategies pay off the same, 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 noise reduction.) Nonetheless the most popular parametrizations of noise and noise reduction fit into this scheme. For instance, the mistakes model can be represented by choosing $g(\Delta)$ to equal $\log(1-\epsilon)/\epsilon$, 0, or $-\log(1-\epsilon)/\epsilon$ as Δ is positive, 0, or negative. The log-linear model has $g(\Delta) = \beta\Delta$. Most general would be a specification where the log-odds ratio was of the form $g(\Delta, \beta)$ where $g(\Delta, \beta)$ converges to $-\infty$, 0 or $+\infty$ as $\operatorname{sgn}\Delta$ is negative, positive or 0. I shall have something to say about more general specifications when discussing the results of the next section.

The explication of the population process entails a description of how the choices of individual players are coupled together. It is customary to consider discrete-time stochastic processes of strategy revision where 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. I find it more convenient to work in continuous time. Each player has a rate 1 Poisson alarm clock: At randomly chosen intervals each player gets a strategy revision opportunity. The random interval lengths are drawn from an exponential distribution with rate parameter 1, and the draws are independent across players and across times. 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 iid exponentially distributed with parameter 1. If there are N players, strategy revision oppotunities for the population as a whole arrive at intervals independently exponentially distributed with rate N. Strategy revision opportunities for any group of M players arrive at rate M. The probability of any two players revising at the same moment of time is a probability 0 event. Another way to view this model is to start with the rate N Poisson alarm clock for the collection of all strategy revision opportunities. When one arrives, a player is chosen at random, equiprobably, to receive it.

This model is easy to work with. The dense forests of B-trees required by KMRY can be avoided altogether in favor of the greener pastures of birth-death process analytics. Moreover, the story just described is a continuous-time limit of the KMRY discrete-time processes. Imagine a succession of KMRY processes where the period length is τ and, for small τ, the probability of a given player's receiving a strategy revision opportunity is $\tau + o(\tau)$. The limit process found by taking τ to 0 gives the continuous-time process described above with revision opportunities arriving at rate 1.

The process $\{M_t\}_{t\geq 0}$ which records the number of players choosing strategy α is called a *strategy revision process.* One consequence of the forgoing 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. The process $\{M_t\}_{t\geq 0}$ is a birth-death process.

A birth-death process on the non-negative integers is characterized by a sequence of birth rates $\{\lambda_{m,m+1}\}$ and death rates $\{\mu_{m+1,m}\}$. Given that the process is in state m at time t, the probability of a transition to $m+1$ at time $t+h$ is $\lambda_{m,m+1} + o(h)$, and the probability of a transition to $m-1$ at time h is $\mu_{m,m-1} + o(h)$. It will always be the case that for all $0 \le m < N$ $\lambda_{m,m+1} > 0$, and for all $0 < m \le N$ $\mu_{m,m-1} > 0$. All other birth and death 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 relations

$$\frac{\rho(M)}{\rho(M-1)} = \frac{\lambda_{M-1,M}}{\mu_{M,M-1}} \quad \text{and} \quad \frac{\rho(M)}{\rho(0)} = \prod_{m=1}^{M} \frac{\lambda_{m-1,m}}{\mu_{m,m-1}}$$

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

To compute these rates for the process $\{M_t\}_{t\ge 0}$, suppose the process is in state m. This means that m players are choosing strategy α and the remaining players are choosing strategy ζ. For the process to increase, a strategy revision opportunity for one of the $N-m$ players choosing ζ arrives at rate $N-m$. The probability that a randomly chosen ζ player will choose α is $\sigma_\alpha(m/N)$. Hence the rate at which the process will increase is $(N-m)\sigma_a(m/N)$. A similar calculation gives the death rates. In summary,

$$\begin{aligned} M \to M+1 \quad &\text{at rate} \quad (N-M)\frac{\exp \beta g\Big(\Delta\Big(\dfrac{M}{N-1}\Big)\Big)}{1 + \exp \beta g\Big(\Delta\Big(\dfrac{M}{N-1}\Big)\Big)} \\ M+1 \to M \quad &\text{at rate} \quad (M+1)\frac{1}{1 + \exp \beta g\Big(\Delta\Big(\dfrac{M}{N-1}\Big)\Big)} \end{aligned} \tag{2.1}$$

A word of explanation is required: Why $M/(N-1)$? Suppose there are M players choosing α. A β player sees that of his $N-1$ opponents, M are choosing α and so the frequency of αs he faces is $M/(N-1)$. If $M+1$ players choose α the frequency with which a α player encounters α is $M/(N-1)$ because she can count herself on neither the top nor the bottom.

Consequently we have the following relations for the invariant distribution:

$$\log \frac{\rho(M)}{\rho(M-1)} = \log \frac{N-M-1}{M} + \beta g\Big(\Delta\Big(\frac{M-1}{N-1}\Big)\Big)$$

and

$$\log \frac{\rho(M)}{\rho(0)} = \log \binom{N}{M} + \beta \sum_{m=0}^{M-1} g\Big(\Delta\Big(\frac{m}{N-1}\Big)\Big). \tag{2.2}$$

There is another way of describing equilibrium which will also be useful for constructing the dynamics. Every symmetric two-strategy game has a *potential*

$$P(M) = \sum_{m=0}^{M-1} \Delta\Big(\frac{m}{N-1}\Big).$$

The significance of the potential is that the change in utility to a player from adopting a new strategy is a first difference of $P(M)$. Consequently, the equilibrium configurations of the player population are precisely the local maxima of the potential $P(M)$.

If $g : \mathbf{R} \to \mathbf{R}$ is sign-preserving, then

$$P_g(M) = \sum_{m=0}^{M-1} g\bigg(\Delta\Big(\frac{m}{N-1}\Big)\bigg) \tag{2.3}$$

is an *ordinal potential* for the game. Now the signs of the change in utility to a player from adopting a new strategy is the sign of a first difference of $P_g(M)$. Again the equilibrium configurations of the player population are precisely the local maxima of $P_g(M)$. For more on the role of potential functions and ordinal potential functions in equilibrium analysis, see Monderer and Shapley (1993). Their significance for this paper is that for a given noise process g, $P_g(M)$ is an ordinal potential function. If β is large enough then the expected movement of the process increases $P_g(M)$. Consequently when β is large, the process puts most of its mass on the global maxima of $P_g(M)$. Equation (2.2) can be rewritten as

$$\log \frac{\rho(M)}{\rho(0)} = \log \binom{N}{M} + \beta P_g(M). \tag{2.4}$$

3. Linear Payoff Differences

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

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

it follows that

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

On the other hand, if only payoff differences matter to choice, and not the "names" of strategies, 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 $-\Delta(p)$. If this symmetry across strategies holds, then

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

Consequently it would follow that $g(d) = -g(-d)$ for all d in the range of $\Delta(p)$. This symmetry property characterizes the largest class of noise processes containing log-linear choice and the mistakes model, and under which equilibrium selection is invariant.

Definition 3.1: A noise process g is *skew-symmetric* if $g(x) = -g(-x)$.

First I show that if the noise process is skew-symmetric, shrinking the noise always selects for the risk-dominant equilibrium in two-by-two coordination games. Next I show that if the noise process is not skew-symmetric, it is possible to construct a game for which shrinking the noise selects the risk-dominated equilibrium.

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 3.1: If g is skew-symmetric, then for all large N $\lim_{\beta\uparrow\infty} \rho(N) = 1$.

Proof: From equation (2.4) and the assumption that g is non-decreasing, it follows that $P_g(M)$ is U-shaped, with local maxima at the two equilibria $M = 0$ and $M = N$. Immediately (2.2) implies that for every $0 < M < N$ at least one of the odds ratios $\rho(M)/\rho(0)$ and $\rho(M)/\rho(0)$ is going to 0 as β grows. The Theorem claims that for N large enough, N will be the unique global maximum, and consequently $\rho(M)/\rho(N)$ will go to 0 as β grows for all $M < N$.

Since $\beta^{-1} \log \rho(N)/\rho(0)$ is a Riemann-sum approximation to the integral $\int_0^1 g\big(\Delta(p)\big)\, dp$, it suffices to show that the value of this integral is positive. To see this, carry out a change of variables.

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

The first term in the sum is 0 because of the symmetry assumption and the second term is positive because $g(x)$ is positive for positive x. $\square$

All skew-symmetric noise processes give the same equilibrium selection result in coordination games. The converse is also true.

Theorem 3.2: If g is not skew-symmetric, there is a coordination game Δ for which g selects the risk-dominated equilibrium.

Proof: Let x^* denote the greatest lower bound of the set of all $x > 0$ such that $g(x) \neq -g(-x)$. There is an $\epsilon > 0$ such that $g(x) \neq -g(-x)$ on all of $(x^*, x^* + \epsilon)$. Suppose that on this interval, $g(x) < -g(-x)$. Choose Δ such that $\Delta(0) = -A$, $\Delta(1) = B$,

$x^* + \epsilon > B > A > 0$. Then

$$\int_0^1 g\big(\Delta(p)\big)\,dp = (B+A)\int_{-A}^{B} g(x)\,dx$$
$$= (B+A)\int_{-A}^{A} g(x)\,dx + (B+A)\int_{A}^{B} g(x)\,dx$$

The first term is negative due to the asymmetry. The second term can be made arbitrarily small by choosing B sufficiently close to A. Thus the value of the integral is negative, and so the Riemann sum will be less than 0 for large enough N. Consequently the risk-dominated equilibrium is selected for.

If $g(x) > -g(-x)$ on the interval $(x^*, x^* + \epsilon)$, take $0 < B < A$. In this case strategy ζ will be risk-dominant. The value of the integral becomes

$$(B+A)\int_{-A}^{-B} g(x)\,dx + (B+A)\int_{-B}^{B} g(x)\,dx.$$

The first second term is positive, and the first term can be made arbitrarily small by choosing A near to B. The value of the integral is positive, and so g selects for strategy α. □

The proofs of Theorem 3.1 and 3.2 extend to the case where the noise process is of the form $g(\Delta) + r(\Delta, \beta)$ where $\beta\|r(\Delta,\beta)\|$ is bounded. For a fixed noise process g, $g(\Delta) \neq 0$ implies that the log-odds of choosing α over ζ go to $\pm\infty$ as β becomes large. The extension to the case of a β-dependent remainder shows that noise in choice which does not make the log-odds ratios unbounded has no effect on selection results.

Consider the following variation on the errors model. 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

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

when α is a best response, and

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

when ζ is a best response.

Take $g(\Delta) = \text{sgn}(\Delta)$, and let

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

where $\epsilon(\beta) = (1 + \exp\beta)^{-1}$. The noise process $g(\Delta) + r(\Delta, \beta)$ 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 selection results are driven by g alone. In this case as $\beta \uparrow \infty$ ($\epsilon \downarrow 0$), strategy α is selected.

Any reformulation of the dynamics that adds terms independent of β or bounded functions of β to equation (2.4) will have no effect on the selection results. One example of this phenomenon is when strategy revision opportunities arrive at different rates to different players, depending upon the strategy they now employ. For instance, 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.

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

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

Consequently the log-odds ratio for the invariant distribution are

$$\log \frac{\rho(M)}{\rho(0)} = \log \binom{N}{M} + \frac{M(2N+3)}{2} \log \gamma + \beta P_g(M).$$

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

4. 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 utilty realizations are independent draws across objects of choice, choice opportunities, and, in games, players. Let ϵ_α and ϵ_ζ be iid random 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

$$\log\sigma_\alpha(p)-\log\sigma_\zeta(p)=\log F\bigl(\Delta(p)\bigr)-\log\Bigl(1-F\bigl(\Delta(p)\bigr)\Bigr).$$

The applicability of the last section's analysis depends upon how the noise is parametrized. Both the mistakes model and the log-linear choice model are random utility models. A generalization of the mistakes model which fits the previous analysis goes like this. With probability $1-\epsilon$ the player correctly observes the utility difference $\Delta(p)$. With probability ϵ she observes utility difference $\Delta(p)+\tilde{x}$ where $\tilde{x}$ is a random variable with support strictly containing the interval $\Delta([0,1])$. It is straightforward to see that for small ϵ, the log of the ratio $\lambda_{M,M+1}/\mu_{M+1,M}$ is approximately $\beta\,\mathrm{sgn}\,\Delta\bigl(M/(N-1)\bigr)$. This approximation is uniform in M, so the Riemann sum argument applies without any further constraint on the distribution of $\tilde{x}$. This is just the simple mistakes model again, and gives the same results.

Other parametrizations of noise reduction are possible which require a different analysis. 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 $\Delta(p)+(\epsilon_\alpha-\epsilon_\zeta)/\beta>0$. Then for parameter β,

$$\log\sigma_\alpha(p)-\log\sigma_\zeta(p)=\log F\bigl(\beta\Delta(p)\bigr)-\log\Bigl(1-F\bigl(\beta\Delta(p)\bigr)\Bigr). \qquad (4.1)$$

Equilibrium selection results will depend upon the behavior of the tails of F.

Theorem 4.1: Suppose that the log-odds of choice are as in equation (4.1). Then in a game with a unique risk dominant equilibrium, that equilibrium will be selected for if N is sufficiently large, for all $x>0$, $1-F(x)+F(-x)>0$, and

$$\liminf_{x\to\infty}\frac{F(-x)}{1-F(x)+F(-x)}>0 \quad\text{and}\quad \limsup_{x\to\infty}\frac{F(-x)}{1-F(x)+F(-x)}<1.$$

The condition of the theorem bounds away from 0 the liminf 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 fdailing 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 as β grows, a fixed number of rectangles covering an ever larger domain is being approximated by an integral, so this approximation gets worse as β grows. The proof instead relies on a direct calculation.

Proof: Suppose that α is risk-dominant, and let k^* denote the largest M such that the utility difference $\Delta\big(M/(N-1)\big)$ is less than or equal to 0. Without loss of generality we can assume $\Delta\big(k^*/(N-1)\big) = 0$, since a standard coupling argument shows that if the result can be established in this case, it follows whenever $\Delta\big(k^*/(N-1)\big) < 0$. 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

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

By convention the sum equals 0 when $M = 0$. Since $\Delta(p)$ is upward sloping, it clear that for large β this function is U-shaped. It suffices to show that $\lim_{\beta\to\infty} \log\big(\rho(N)/\rho(0)\big) = \infty$. Write this log-odds ratio as

$$\sum_{k=0}^{k^*} \log \frac{F\Big(\beta\Delta\big(\frac{k^*-k}{N-1}\big)\Big)}{\Big\{1 - F\Big(\beta\Delta\big(\frac{k^*+k}{N-1}\big)\Big)\Big\}} + \log \frac{F\Big(\beta\Delta\big(\frac{k^*+k}{N-1}\big)\Big)}{\Big\{1 - F\Big(\beta\Delta\big(\frac{k^*-k}{N-1}\big)\Big)\Big\}} + \sum_{k=2k^*+1}^{N-1} \log \frac{F\Big(\beta\Delta\big(\frac{k}{N-1}\big)\Big)}{\Big\{1 - F\Big(\beta\Delta\big(\frac{k}{N-1}\big)\Big)\Big\}}$$

The 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 limsup and one of the liminf conditions together with the hypothesis that the probability of being in the tails is always positive. The same kind of argument when ζ is risk dominant also works using the remaining liminf and limsup conditions. $\square$

The condition of Theorem 4.1 requires that the tails of F not be too skewed. 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.

Theorems 3.1 and 4.1 identify a large class of random utility models which give rise to risk-dominant selection. Here is an example of a parametric random utility model for which risk-dominant selection fails.

Let $-A = \Delta(0)$ and $B = \Delta(1)$. Assume $B > A$ so that α is risk-dominant. Suppose that the return to strategy α is observed without error, and that the return to ζ is observed with error ϵ_ζ. Let p^* denote the root of $\Delta(p)$, and choose $n > (1 - p^*)/p^* = z$. Suppose that, for parameter value β, the distribution of ϵ_ζ places mass β^{-n} on $-2B$, mass β^{-1} on $2B$ and the remaining mass on $C(\beta) = -2B\beta^{-1}(1-\beta^{1-n})$. For all β the random variable ϵ_ζ has mean 0, and as β gets large the distribution converges weakly to point mass at 0.

Choose N large enough that $k^* = \sup\{k : k \leq p^* N\}$ is less than $n/2 - 1$, and observe that for large enough β, $0 > C(\beta) > \Delta\big(k^*/(N-1)\big)$. Thus for large enough β choice is only perturbed by observations of the extreme values of ϵ_ζ. Then

$$\begin{aligned}
\log \frac{\rho(N)}{\rho(0)} &= \sum_{k=0}^{N-1} \log \frac{F\Big(\beta\Delta\big(\frac{k}{N-1}\big)\Big)}{\Big\{1 - F\Big(\beta\Delta\big(\frac{k}{N-1}\big)\Big)\Big\}} \\
&= k^* \log \frac{\beta^{-n}}{1-\beta^{-n}} + (N - k^* - 1)\log \frac{1-\beta^{-1}}{\beta^{-1}} \\
&< Np^* \log \frac{\beta^{-n}}{1-\beta^{-n}} + N(1-p^*)\log \frac{1-\beta^{-1}}{\beta^{-1}} \\
&< Np^* \left\{ \log \frac{\beta^{-n}}{1-\beta^{-n}} - \log \frac{\beta^{-z}}{(1-\beta^{-1})^z} \right\} \\
&= Np^* \big\{ \log \beta^{z-n} + \log(1-\beta^{-1})^{z-n} \big\}
\end{aligned}$$

The last expression converges to $-\infty$ as β grows large. Similar arguments show that $\rho(M)/\rho(0)$ converges to 0 for all $M > 0$, so the risk dominated strategy is selected.

5. Population Games

For pairwise matching processes in two-by-two games, the evolutionary equilibrium selection results have been shown to be invariant to the noise process so long as some the noise satisfies some symmetry requirements. This section extends the analysis of section 3 to a larger class of games, which I call population games.

Imagine a population of identical players, each interacting with the entire population. Each player can choose one of two strategies, and her payoffs depend upon the fraction of the population choosing each strategy. The payoff differences can be summarized by a function $\Delta(p)$. The payoff function for the pairwise matching process was constructed by assuming that players perceive that the population play is unchanging, and that they have an equal chance of being matched against anyone in the population. Thus the value of a strategy is just its expected return against the distribution of play in the population.

Other formulations fit into the same mode. 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 $\Delta(p)$, which since the return to β is 0 also represents the value of α. Now the evolutionary process models the dynamics of group membership.

In the more general context of population games the function $\Delta(p)$ need not be linear. Linearity was crucial to the proofs of the theorems in the last two sections, and so it should not be surprising that without linearity invariance typically fails even amoung skew-symmetric noise processes. This section illustrates this, and then establishes a condition on $\Delta(p)$ which implies outcome invariance for skew-symmetric noise processes. For the remainder of the paper only the parametrization of noise shrinkage introduced in section 2 will be studied.

An example of invariance failure is provided by the payoff difference function $\Delta(p) = -A + (B + A)p^2$ where A and $B + A$ are positive. The payoff difference function $\Delta(p)$ is increasing in p, and so a strategic complementarity is present. 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 $p > p^*$. The mistakes model will select α if and only if $p^* < 1/2$, that is, if and only if $3A < B$.

For the log-linear model, the selection criterion is different. Suppose N is large, so that the Riemann sum approximation is good. Then α will be selected by the log-linear noise process if and only if $\int_0^1 g\big(\Delta(p)\big)dp > 0$. Calculation shows that the integral will be positive if and only if $2A < B$. Thus when $B > 3A$, both processes select α. When $B < 2A$ both processes select ζ. And when $2A < B < 3A$ the log-linear noise process selects α and the mistakes selects ζ.

Although invariance fails in for some parameter values in this example, invariance holds for a class of payoff difference functions that includes more than linear functions. Suppose $\Delta : [0, 1] \to \mathbf{R}$ is increasing and C^1. Let $p(\Delta)$ denote the inverse function of $\Delta(p)$. I will assume that the noise process g is strictly increasing. This rules out the mistakes model, but of course the mistakes model can be approximated arbitrarily well by noise processes in this class.

Theorem 5.1: If $-\Delta(0) < \Delta(1)$ then there is a skew-symmetric noise process that selects α. If in addition $p(\Delta) + p(-\Delta) \le 1$, then all skew-symmetric noise processes select α. If $-\Delta(0) = \Delta(1)$ and $p(\Delta) + p(-\Delta) < 1$, then all skew-symmetric noise processes select α. If $-\Delta(0) > \Delta(1)$, then not all skew-symmetric noise processes select α, and there exists a skew-symmetric noise process that selects α if and only if $p(\Delta) + p(-\Delta) < 1$.

The payoff difference function $\Delta(p)$ is said to be *invariant to noise* if it satisfies the invariance condition of theorem 5.1. The invariance condition is not easy to interpret. Examination of the proof shows that, if it is met, then the effects of asymmetry on the interval $\big(\Delta(0), -\Delta(0)\big)$ in the worst possible case are just offset by the consequences of noisy choice on the interval $\big(-\Delta(0), \Delta(1)\big)$

Proof: It is necessary to prove only the first two parts of the Theorem, because the third case, where $-\Delta(0) > \Delta(1)$, follows by applying the first part to the selection of ζ.

So long as $-\Delta(0) = A < \Delta(1) = B$ there is a skew-symmetric model that will select α for large enough N. Take $g'(x) = \epsilon$ for $0 \le x \le -D(0)$ and $g'(x) = 1$ for $x > -D(0)$. Extend $g'(x)$ to $x < 0$ by symmetry. For ϵ small enough g will select α. The function $g\big(\Delta(p)\big)$ is increasing in p. Thus the Riemann sum $P_g(N)$ always underestimates the integral. Thus it has to be shown that precisely under the stated conditions, $\int_0^1 g\big(\Delta(p)\big)dp > 0$ for all skew-symmetric g. Changing variables,

$$\int_0^1 g\big(\Delta(p)\big)\,dp = \int_{-A}^0 g(\Delta)p'(\Delta)\,d\Delta + \int_0^A g(\Delta)p'(\Delta)\,d\Delta + \int_A^B g(\Delta)p'(\Delta)\,d\Delta.$$

If the conditions of the theorem are not satisfied, there is a skew-symmetric g for which the value of this expression is negative. To see this, suppose that $g'(x) = \epsilon$ for $x > A$. Then the value of the last term is $g(A)\big(1 - p(A)\big) + \delta$ where $\delta > 0$ goes to 0 with ϵ. Symmetry implies that the value of the first term is

$$\int_{-A}^0 g(-\Delta)p'(\Delta)\,d\Delta = -\int_A^0 g(\Delta)p'(-\Delta)\,d\Delta \tag{5.1}$$

where the equality follows from a change of variables. Thus the sum of the first two terms is

$$\int_0^A g(\Delta)\big(-p'(-\Delta) + p'(\Delta)\big)\,d\Delta.$$

Integrating by parts and using the facts that $p(-A)$ and $g(0)$ are both 0 gives for this sum

$$g(A)p(A) - \int_0^A g'(\Delta)\big(p(-\Delta) + p(\Delta)\big)\,d\Delta.$$

Thus the value of the left hand side of (5.1) is

$$g(A) + \delta - \int_0^A g'(\Delta)\big(p(-\Delta) + p(\Delta)\big)\,d\Delta.$$

If the conditions of the theorem are not satisfied, then the value of this expression can be made negative. Choose an interval (d, e) on which $p(-\Delta) + p(\Delta)$ exceeds $1 + \gamma$ where γ is some small positive number.

Make $g'(d)$ equal to ϵ outside of this interval, and equal to 1 on the interval $(d+\epsilon, e-\epsilon)$. (This is the same ϵ as earlier.) Now let ϵ go to 0. Then δ converges to 0, and the integral converges to a term which can be bounded from below by $(1+\gamma)g(A)$. Thus for small enough ϵ the expression is negative. Conversely, since making g' small on $[A, B]$ is a worst case, it is clear that if the condition of the Theorem is satisfied, then the value of the integral must always be at least $\int_A^B g(\Delta)p'(\Delta)\,dp$, which exceeds 0.

Finally, a careful check of the calculation shows that when $A = B$, $p(-\Delta) + p(\Delta) < 1$ is required for the value of the integral to be positive. $\square$

Theorem 5.1's criterion is a bound on the asymmetry of the payoff difference function $\Delta(p)$. If $p(\Delta)$ is large, then $p(-\Delta)$ is required to be small. Notice that the criterion is "one way". If $p(-\Delta)$ is small, there is no corresponding criterion that $p(\Delta)$ be large. This is because asymmetries in this direction do not hurt the case for selection of α.

Checking theorem 5.1's criterion in the quadratic example from the beginning of the section, it appears that $p(-\Delta) + p(\Delta) \leq 1$ precisely when $3A \leq B$. We have already see that when $3A > B > 2A$ the log-linear and error models select for different strategies. So long as $A < B$, $\Delta(1) > A$ and there will exist a skew-symmetric noise process g that selects strategy α. Choose g so that $g(x) = 1$ for $x > A$ and $g(x) < \epsilon$ for $0 < x \leq A$. If ϵ is small enough, g will select for strategy α because the value of the potential function $P_g(N)$ is approximately the integral of $g(x)$ over the payoff range where $g(x)$ equals 1. If $A > B > 0$, then the condition of the theorem will be met for the difference $-\Delta(1-q)$ where q is the fraction of the population choosing ζ and $-\Delta$ is the payoff advantage of ζ over α. Consequently all skew-symmetric noise processes will select for ζ. Finally, if $B < 0$, the only equilibrium has everyone playing ζ.

When $\Delta(p)$ is linear, the requirement for the selection of α by all skew-symmetric noise processes is that the root p^* of $\Delta(p)$ is less than 1/2. This condition, which measures the relative sizes of basins of attraction of the two strategies is also connected with the invariance condition of Theorem 5.1. For the invariance condition to be satisfied, this condition is necessary, for otherwise the condition of theorem 5.1 must fail in a sufficiently small neighborhood of $\Delta = 0$. When $\Delta(p)$ is convex, this condition is also sufficient.

Corollary 5.1: Let p^* denote the root of $\Delta(p)$. The invariance condition $p(\Delta)+p(-\Delta) \leq 1$ fails if $p^* > 1/2$. If $\Delta(p)$ is convex and $\Delta(1) > -\Delta(0)$, then $p(\Delta) + p(-\Delta) \leq 1$ if and only if $p^* \leq 1/2$.

Proof: (Necessity) Suppose that $p^* > 1/2$. Then for all Δ sufficiently near to 0, $p(\Delta) + p(-\Delta) \approx p(0) + p(0) = 2p^*$ which is greater than 1.

(Sufficiency) If $\Delta(p)$ is convex, then $p(\Delta)$ is concave, and so $p(\Delta) + p(-\Delta)$ has a maximum where $p'(\Delta) = p'(-\Delta)$. The unique solution to this first order condition is $\Delta = 0$. IF $p^* < 1/2$, then $p(\Delta) + p(-\Delta)$ is less than 1 at $\Delta = 0$, and hence at all Δ in $\big(\Delta(0), -\Delta(0)\big)$. $\square$

This theorem is confirmed in the quadratic example. There the errors model fails right at the boundary of the parameter set on which selection is invariant to noise. The error model works explicitly by measuring p^*; Np^* is an estimate of the number of deviations from best response required to move the population process from everyone playing ζ into the basin of attraction of α, and $N(1-p^*)$ estimates the number of deviations needed to move in the reverse direction. The errors model selects for the strategy with the largest basin of attraction.

The special role of p^* can be seen in the following consequences of the Theorem 5.1.

Corollary 5.2: Suppose that $\Delta(p)$ is symmetric around some $1 > p^* > 0$; that is, for all ρ such that $p^* \pm \rho$ is in range, $D(p^* + \rho) = -\Delta(p^* - \rho)$. Then the condition of Theorem 5.1 is satisfied if and only if $p^* \leq 1/2$.

Proof: Let $\rho(\Delta)$ solve $\Delta(p^* + \rho) = \Delta$. Then

$$\begin{aligned} p(-\Delta) + p(\Delta) &= 2p^* + \rho(\Delta) + \rho(-\Delta) \\ &= 2p^* + \rho(\Delta) - \rho(\Delta) \\ &= 2p^* \end{aligned}$$

Again the criterion is satisfied if and only if $p^* \leq 1/2$ (with strict inequality if $-\Delta(0) = \Delta(1)$. □

When $\Delta(p)$ is convex, p^* will be above that of the linear model with the same end-points, and skew-symmetric noise will select for α if and only if $p^* < 1/2$ (with a small qualification for symmetric range). When $\Delta(p)$ is concave, p^* is less than what it would be in the linear case. In this case strategy α is always selected for.

Corollary 5.3: If $\Delta(p)$ is concave in p, then a sufficient condition for all skew-symmetric noise processes to select for α is that $\Delta(1) > -\Delta(0)$.

Proof: In this case $p(\Delta)$ is convex, and therefore so is $p(\Delta) + p(-\Delta)$. It takes its maxima on the range $[\Delta(0), -\Delta(0)]$ at the endpoints, where its value is $p\big(-\Delta(0)\big)$. This value will be less than 1 so long as $\Delta(1) > -\Delta(0)$. □

6. Heterogeneous Noise Processes

Throughout the analysis I have assumed that, whatever the noise process, it is the same for all players. In this section I investigate the dependence of selection results on this assumption by introducing players with different noise processes. The chief modification required of the foregoing analysis is that the process $\{M_t\}_{t \geq 0}$ which records the number of players choosing α is no longer Markov. Now it matters not only how many players are choosing α but which players are choosing α. The process which records the play of each individual will be a Markov process, and although the resulting aggregate process is no longer a Markov process, the analysis of the previous section can be used as a proof device to demonstrate that selection does not depend on homogeneity of the noise process across traders.

To prove this result some apparatus is required. Denote the set of players by I. A *configuration* of the population is a map $\phi : I \to \{\alpha, \zeta\}$ which describes the current choice of each player. We shall construct Markov processes on the space C of configurations. For a given configuration ϕ, ϕ_i denotes the configuration which is identical to ϕ for all players other than i: $\phi(j) = \phi_i(j)$ for $j \neq i$, and $\phi(i) \neq \phi_i(i)$. Also let ϕ^α denote the configuration in which each player is playing α.

All players see the same payoff difference function $\Delta(p)$. To each player $i \in I$ is assigned a noise process $g_i(\Delta)$ and a shrinkage parameter β_i. Each $g_i(\Delta)$ is a skew-symmetric noise process.

The state space C for the strategy revision process is finite. To describe a continuous time Markov process $\{\phi_t\}_{t\geq 0}$ on C it suffices to describe the transition rates between one configuration and the next. As before, each player has a strategy revision opportunity whenever her rate-1 Poisson alarm clock rings. Let $1_i(\phi)$ denote the indicator function which is $+1$ when $\phi(i) = \alpha$ and 0 otherwise. Also violate the notation by defining $\Delta(\phi)$ to be $\Delta(M/N - 1)$ where M is the number of players in ϕ playing α, $\#\phi^{-1}(\alpha)$, and N is the number of players.

$$\begin{aligned}\phi \to \phi_i \quad \text{at rate} \quad & \Big(1 + \exp -\beta_i g_i\big(\Delta(\phi)\big)\Big)^{-1} 1_i(\phi) + \\ & \Big(1 + \exp \beta_i g_i\big(\Delta(\phi)\big)\Big)^{-1} \big(1 - 1_i(\phi)\big).\end{aligned} \tag{6.1}$$

All other transitions require at least two players to simultaneously revise their strategies, and these occur at rate 0. It is easy to see that when all the β_i and all the $g_i(\Delta)$ are equal, that all sites playing the same strategy switch at the same rate, and from this and the fact that payoffs are determined only by M_t, the birth-death process of the previous sections can be derived. In this sense the current framework generalizes the framework of the last section.

Players now behave in different ways. They share in common the assessment of payoff differences, $\Delta(p)$, and the fact that as their individual sensitivity parameters β_i grow large, their behavior converges to best-response to the play of the population. This is enough to prove the following Theorem:

Theorem 6.1: If $\Delta(p)$ satisfies the conditions of Theorem 5.1 and if the noise processes $g_i(\Delta)$ are skew-symmetric, then as $(\beta_i)_{i\in I}$ grow to $+\infty$, the stationary distribution for the Markov process with rates (6.1) converges to point mass on the configuration ϕ^α.

The remainder of this section is devoted to proving this theorem. The ideas of the proof are simple, but it takes some apparatus to work things out. The trick to analyzing this process is to compare it to a process of the kind whose behavior is already known. One process will be the one we want to study. The other process will be one with homogeneous players in which it is always harder to get a given player to switch to α and easier to get her to switch to ζ from comparable configurations. If the second process selects for α, than so should the first. To work this comparison of the two processes out involves building them on the same probability space in a way that makes the comparison easy to see. This procedure is called "coupling". A discussion of coupling techniques can be found in Durrett (1991). The construction of the processes is the tedious part of the analysis. The argument follows quite quickly once the coupling is set up.

Let $g(\Delta) = \min_i g_i(\Delta)$ for $\Delta \geq 0$, and let $g(\Delta) = -g(-\Delta)$ for $\Delta < 0$. Then for $\Delta < 0$ and all i, $g(\Delta) \geq g_i(\Delta)$ and $g(\Delta)$ is skew-symmetric. Also let $\beta = \min_i \beta_i$. Let $\{\hat{\phi}_t\}_{t \geq 0}$ denote the Markov process given by equation (6.1) with $g_i(\Delta) = g(\Delta)$ and $\beta_i = \beta$ for all i. This is a process whose behavior we already know: It is ergodic, and the limit distribution converges to point mass on the configuration ϕ^α as β grows large. It is harder to adopt α in the $\{\hat{\phi}\}$ process and easier to switch to ζ.

The most elementary way to compare these processes is to convert the problem to the study of discrete time Markov chains. The continuous time process works as follows: Start in some initial state ϕ_0. Wait a random amount of time and then switch to some new state ϕ_1. Wait another random amount of time and then flip to ϕ_2, and so forth. Let τ_k denote the time of the kth arrival of a strategy revision opportunity (and $\tau_0 = 0$). Then the $\tau_k - \tau_{k-1}$ are, by assumption, independently and exponentially distributed with mean $1/N$. The τ_k are stopping times, and the strong Markov property says that the process $\{\phi_{\tau_k}\}_{k=1}^\infty$ is also Markov. The configuration ϕ_{τ_k} is the state of the process immediately after the kth change of state. This process is a discrete-time Markov chain and is called the *embedded chain* for the process described by equations (6.1). Henceforth I shall write ϕ_k for ϕ_{τ_k}. it is simple to show that the invariant distributions of the embedded chain are precisely the invariant distributions of the original process.

Let $\phi^{i\alpha}$ and $\phi^{i\zeta}$ respectively denote the configuration ϕ with an α or ζ in the ith position. When a strategy revision opportunity arrives, with probability $1/N$ it belongs to player i. When it belongs to player i the configuration changes according to the following transition matrix:

$$\begin{aligned} P(\phi^{i\alpha} \mid \phi) &= \Big(1 + \exp -\beta_i g_i\big(\Delta(\phi)\big)\Big)^{-1} = p_i(\alpha, \phi) \\ P(\phi^{i\zeta} \mid \phi) &= \Big(1 + \exp \beta_i g_i\big(\Delta(\phi)\big)\Big)^{-1} = p_i(\zeta, \phi) \\ P(\phi' \mid \phi) &= 0 \quad \text{for all other } \phi' \end{aligned} \tag{6.2}$$

A similar construction applies to the process $\hat{\phi}$. Here the transitions are

$$\begin{aligned} \hat{P}(\hat{\phi}^{i\alpha} \mid \hat{\phi}) &= \Big(1 + \exp -\beta g\big(\Delta(\hat{\phi})\big)\Big)^{-1} = p(\alpha, \hat{\phi}) \\ \hat{P}(\phi^{i\zeta} \mid \phi) &= \Big(1 + \exp \beta g\big(\Delta(\hat{\phi})\big)\Big)^{-1} = p(\zeta, \phi) \\ \hat{P}(\phi' \mid \phi) &= 0 \quad \text{for all other } \phi' \end{aligned} \tag{6.3}$$

The comparison is carried out by constructing a process on $C \times C$ whose marginal process in the first component is the $\hat{\phi}$ process and whose marginal process in the second component is the ϕ process, and whose coordinates almost surely stand in a relation to each other that facilitates comparison between the behavior of the two processes.

Define $\phi \geq \hat{\phi}$ if $\hat{\phi}^{-1}(\alpha) \subset \phi^{-1}(\alpha)$; that is, every player who chooses α in the $\hat{\phi}$ configuration also chooses α in the ϕ configuration. Let $K \subset C \times C$ denote the set of all pairs $(\hat{\phi}, \phi)$ such that $\phi \geq \hat{\phi}$.

Having an ordering on C, we can define stochastic dominance of measures on C. Say that μ *stochastically dominates* ν ($\mu \geq \nu$) if $\int f \, d\mu \geq \int f \, d\nu$ for every increasing function $f : C \to \mathbf{R}$. The following Lemma is well-known in the stochastic dominance literature (Strassen 1965).

Lemma 6.1: Let X be a compact metric space with a partial ordering such that $K = \{(x, y) \in X \times X : x \geq y\}$ is closed in the product topology. The probability distribution μ_1 on X stochastically dominates probability distribution μ_2 if and only if there exists a probability distribution μ on $X \times X$ such that for all Borel sets A of X,

1. $\mu\{(x, y) : x \in A\} = \mu_1(A)$,
2. $\mu\{(x, y) : y \in A\} = \mu_2(A)$, and
3. $\mu(K) = 1$.

Now for the construction. First observe that if $\phi \geq \hat{\phi}$, then for all i $p_i(\alpha, \phi) \geq p(\alpha, \hat{\phi})$. This follows from the construction of p, the fact that each $g_i(\Delta)$ is non-decreasing and the fact that $\Delta(p)$ is non-decreasing. Define the following transitions:

$$Q(\hat{\eta}, \eta \,|\, \hat{\phi}, \phi) = \hat{P}(\hat{\eta} \,|\, \hat{\phi}) P(\eta \,|\, \phi)$$

if $(\hat{\phi}, \phi) \notin K$. For $(\hat{\phi}, \phi) \in K$ and $\phi(i) = b$,

$$\begin{aligned} Q(\hat{\phi}_i^{\alpha}, \phi_i^{\alpha} \,|\, \hat{\phi}, \phi) &= \hat{P}(\phi_i^{a} \,|\, \hat{\phi}) \\ Q(\hat{\phi}, \phi_i^{a} \,|\, \hat{\phi}, \phi) &= P(\phi_i^{a} \,|\, \hat{\phi}) - \hat{P}(\hat{\phi}_i^{a} \,|\, \phi) \\ Q(\hat{\phi}, \phi \,|\, \hat{\phi}, \phi) &= 1 - Q(\hat{\phi}_i^{a}, \phi_i^{a} \,|\, \hat{\phi}, \phi) - Q(\hat{\phi}, \phi_i^{a} \,|\, \hat{\phi}, \phi) \end{aligned}$$

For $(\hat{\phi}, \phi) \in K$ and $\hat{\phi}(i) = a$,

$$\begin{aligned} Q(\hat{\phi}_i^{\zeta}, \phi_i^{\zeta} \,|\, \hat{\phi}, \phi) &= P(\phi_i^{b} \,|\, \phi) \\ Q(\hat{\phi}_i^{b}, \phi \,|\, \hat{\phi}, \phi) &= \hat{P}(\phi_i^{b} \,|\, \hat{\phi}) - P(\hat{\phi}_i^{b} \,|\, \phi) \\ Q(\hat{\phi}, \phi \,|\, \hat{\phi}, \phi) &= 1 - Q(\hat{\phi}_i^{b}, \phi_i^{b} \,|\, \hat{\phi}, \phi) - Q(\hat{\phi}_i^{b}, \phi \,|\, \hat{\phi}, \phi) \end{aligned}$$

For $(\hat{\phi}, \phi) \in K$ and $b = \hat{\phi}(i) \neq \phi(i) = a$,

$$\begin{aligned} Q(\hat{\phi}_i^{a}, \phi \,|\, \hat{\phi}, \phi) &= \hat{P}(\phi_i^{a} \,|\, \hat{\phi}) \\ Q(\hat{\phi}, \phi_i^{b} \,|\, \hat{\phi}, \phi) &= P(\phi_i^{b} \,|\, \phi) \\ Q(\hat{\phi}, \phi \,|\, \hat{\phi}, \phi) &= 1 - Q(\hat{\phi}_i^{a}, \phi \,|\, \hat{\phi}, \phi) - Q(\hat{\phi}, \phi_i^{b} \,|\, \hat{\phi}, \phi) \end{aligned}$$

There is a simple story behind this complicated list of transitions. When the initial state of the process is not in K, the two coordinates move independently of each other according to their respective transition probabilities $\hat{P}$ and P. When the state is in K and both i coordinates are equal, both change together with the smallest of the transition probabilities prescribed by $\hat{P}$ and P. The remaining coordinate changes by itself with the residual probability allowed by its transition law. When the state is in K but the two states are not equal, they move alone according to their prescribed transition probabilities. The net effect of all this, which can be verified by simple calculations, is first that if $(\hat{\phi}_k, \phi_k) \in K$, then so are the states $(\hat{\phi}_{k+l}, \phi_{k+l})$ for all $l > 0$. That is, K is invariant under the process. Second, that starting from any initial configuration, K is reached with probability 1. Third, that for any joint distribution μ on $C \times C$ at time k with marginal distribution μ^1 on the first coordinate $\hat{\phi}$, the distribution of $\hat{\phi}$ at time $k+1$ is given by $\sum_{\hat{\eta}} \hat{P}(\hat{\phi} \,|\, \hat{\eta})\mu^1(\hat{\eta})$, and similarly for μ^2 and ϕ_{k+1}. A consequence of the third fact is that if μ is invariant for Q then the marginal distribution μ^1 is invariant for $\hat{P}$ and the marginal distribution μ^2 is invariant for P.

This construction is a standard "coupling" of the $\{\hat{\phi}_k\}$ and $\{\phi_k\}$ processes. We can use this coupling to quickly finish off the proof. The process Q is easily seen to be irreducible, and so it has a unique invariant distribution μ^*. Let μ_0 denote an given initial distribution of $(\hat{\phi}, \phi)$ whose support is contained in the diagonal; that is, $\mu_0\{\hat{\phi} = \phi\} = 1$. Then $\mu_0^1 = \mu_0^2$. Since the diagonal is a subset of K, $\mu_k\{K\} = 1$ for all k. Since K is closed, taking limits shows that $\mu^*(K) = 1$. From Strassen's lemma conclude that the invariant distribution μ^{1*} for the $\{\hat{\phi}\}$ process is stochastically dominated by the distribution μ^{2*} for the $\{\phi\}$ process. In other words, the limit probability that any given subset of players will play α in the $\{\phi\}$ process exceeds that in the $\{\hat{\phi}\}$ process. This is true for all $(\beta_i)_{i \in I}$. Now let all the β_i grow large, so that β grows to $+\infty$. The limit distribution for the $\{\hat{\phi}\}$ process converges to point mass on the configuration ϕ^α — we know this from the analysis of the previous section because the $\{\hat{\phi}\}$ process is a strategy revision process with homogeneous, skew-symmetric noise and $\Delta(p)$ satisfies the conditions of Theorem 5.1. Stochastic dominance is preserved by limits and the point mass on ϕ^α is the unique maximal distribution under the stochastic dominance relation, so the invariant distributions for the $\{\phi\}$ process must also be converging to point mass on the configuration ϕ^α. This quick argument after the long setup proves the theorem.

7. Conclusion

In this paper I have argued several points. Most important, the conclusions of the so-called evolutionary equilibrium selection 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 intuitive 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.

Second, the robust equilibrium selection results extend to a larger class of games, here called *population games*. The robustness criterion for noise processes 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 providing the non-linearities of the payoff difference functions don't 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/Kiefer techniques developed by KMRY. 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. I worked in continuous time because I find continuous time jump processes to be natural for economic processes and because the continuous-time birth-death processes are a natural limit of the processes studied by KMRY 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 selection 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; thus 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 the global maxima of the potential in a potential game might prove to be an interesting 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 robust selection results presented here can be viewed as conditions on the construction of a class of ordinal potentials that guarantee they all have the same argmax as the potential.

This potential function technology has other uses. For instance, it can be shown that 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. And the mutation counting, tree-costing techniques of KMRY are no treat with more than two strategies either. I hope that others too will find these techniques useful.

REFERENCES

Bergin, J. and B. Lipman (1994) "Evolution with state-dependent mutations," unpublished, Queens University.

Blume, L.. (1993) "The statistical mechanics of strategic interaction," *Games and Economic Behavior,* 4, 387–424.

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

Foster, D. and H. P. Young (1990) "Stochastic evolutionary game dynamics," *Theoretical Population Biology,* 38, 219–32.

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

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

Monderer, D. and L. Shapley (1993) "Potential games," unpublished, The Technion.

Strassen, V. (1965) "The existence of probability distributions with given marginals," *Ann. Math. Stat.,* 36, 423–39.

Young, H. P. (1993) "The evolution of choice behavior," *Econometrica,* 61, 57–84.