

Evolutionary game dynamics in finite populations with strong selection and weak mutation

Drew Fudenberg^a, Martin A. Nowak^{b,c}, Christine Taylor^{b,c}, Lorens A. Imhof^{d,*}

^a*Department of Economics, Harvard University, Cambridge, MA 02138, USA*

^b*Program for Evolutionary Dynamics, Harvard University, Cambridge, MA 02138, USA*

^c*Department of Mathematics, Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138, USA*

^d*Statistische Abteilung, Universität Bonn, Adenauerallee 24-42, D-53113 Bonn, Germany*

Received 9 September 2005

Available online 12 August 2006

Abstract

We study stochastic game dynamics in finite populations. To this end we extend the classical Moran process to incorporate frequency-dependent selection and mutation. For 2×2 games, we give a complete analysis of the long-run behavior when mutation rates are small. For 3×3 coordination games, we provide a simple rule to determine which strategy will be selected in large populations. The expected motion in our model resembles the standard replicator dynamics when the population is large, but is qualitatively different when the population is small. Our analysis shows that even in large finite populations the behavior of a replicator-like system can be different from that of the standard replicator dynamics. As an application, we consider selective language dynamics. We determine which language will be spoken in finite large populations. The results have an intuitive interpretation but would not be expected from an analysis of the replicator dynamics.

© 2006 Elsevier Inc. All rights reserved.

Keywords: Coordination game; Favored strategy; Frequency-dependent selection; Imitation process; Language dynamics; Moran process; Replicator dynamics; Risk-dominance; Selected strategy; Stochastic dynamics

1. Introduction

Evolutionary processes describing frequencies of phenotypes in biological populations are generally driven by selection and mutation. Mutations give rise to variety and selection favors some phenotypes over others. Game theory provides a means to study frequency-dependent selection, where the fitness of a phenotype depends on the composition of the population. In evolutionary biology, game dynamics are typically modelled by deterministic differential equations that describe the evolution of infinite populations. A widely used system that focuses on the effects of frequency-dependent selection is the replicator dynamics by Taylor and Jonker (1978) and Hofbauer et al.

(1979). See Fudenberg and Levine (1998), Hofbauer and Sigmund (1998, 2003) and Nowak and Sigmund (2004) for surveys of related results. Foster and Young (1990, 1997) and Fudenberg and Harris (1992) have introduced evolutionary models for infinite populations with stochastic shocks. In the model of Fudenberg and Harris, the payoffs are subject to aggregate stochastic shocks, but they argue that, in infinite populations, mutations are best modelled deterministically. They give a complete analysis for 2×2 games; and stochastic replicator dynamics for games with more than two pure strategies have been examined by Cabrales (2000) and Imhof (2005).

It is natural to investigate game dynamics in finite populations. In the present paper, we analyze a model of stochastic evolution in finite populations which incorporates frequency-dependent selection and weak mutation. The members of the population are programmed to a pure strategy of a symmetric two-player game. They are randomly matched to play the game and the expected

*Corresponding author.

E-mail addresses: dfudenberg@harvard.edu (D. Fudenberg), nowak@fas.harvard.edu (M.A. Nowak), christine_taylor@harvard.edu (C. Taylor), limhof@uni-bonn.de (L.A. Imhof).

payoff is interpreted as fitness. The composition of the population evolves as follows. At each time step, one individual is chosen for reproduction with probability proportional to its fitness. With probability close to 1, the offspring inherits the strategy of its parent, but we assume that there is a positive probability that the offspring is a mutant playing another strategy. After reproduction, one randomly chosen individual dies and is replaced by the offspring. Thus, both selection and mutation are stochastic. Our model is a generalization of the Moran process (Moran, 1962, Ewens, 2004) to frequency-dependent fitness. The classic Moran process corresponds to the special case where selection is constant and mutations are not allowed. The assumption that in our model the mutation rates are positive ensures that the process is ergodic and so has a unique stationary distribution. In the present paper we study the limit distribution, that is, the limit of the stationary distribution obtained by letting the mutation rates go to zero. We show that this limit distribution is always concentrated on the pure states where all agents use the same strategy.

Most of this paper analyzes the case of symmetric 2×2 games. After obtaining the formula for the limit distribution in terms of the parameters of the payoff matrix, we then classify games according to which strategy is favored, where a strategy is said to be favored if the limit distribution assigns probability greater than $\frac{1}{2}$ to the corresponding pure state. In some cases, one strategy is favored for any population size, N , and the strategy is selected in the sense that the probability assigned to the corresponding pure state converges to 1 as $N \rightarrow \infty$. In other cases, the favored strategy depends on the population size. This is true in particular for the case where B is the dominant strategy, but the payoff of playing A against B exceeds the payoff of playing B against A . Here because of “spite” effects, a dominated strategy can be favored in small populations. The possibility of a spite effect in finite populations was pointed out by Hamilton (1971), and led Schaffer (1988) to propose an alternative definition of evolutionary stability for finite populations. In our model, the favored equilibrium can also depend on the population size in games with two pure-strategy equilibria, and in the limit of large populations, the selected equilibrium may differ from the equilibrium selected by many other stochastic finite-population models. This stems from the fact that our process is stochastic even without mutations. We discuss the topic in detail in the last section, where we also relate our process to the deterministic and stochastic replicator dynamics.

In addition to our classification of the limit distribution in all 2×2 games, we analyze the limit distribution in 3×3 coordination games with large populations. In the spirit of the literature on large deviations we relate the long-run behavior of the process with mutations to that of the simpler process where mutations are absent. The classic reference for large deviations is Freidlin and Wentzell (1984), though they focus on the continuous time case;

Kifer (1990) gives a general result for discrete time. For any fixed population size, the limit distribution will not be concentrated on a single point, so it cannot be determined by the standard “mutation-counting” arguments in the style of Kandori et al. (1993) and Young (1993). Our analysis relies on a general result of Fudenberg and Imhof (2005), which yields, for every finite population size, an expression of the limit distribution in terms of certain absorption probabilities of the no-mutation process. The present paper takes advantage of the special structure of the Moran process to provide a characterization of the asymptotic behavior of the absorption probabilities as the population size goes to infinity. We use this characterization to show that for generic payoffs the limit distribution converges to a point mass on one of the three pure states, and we obtain a simple criterion to determine the selected strategy. We apply our results for 3×3 games to study selective language dynamics in finite populations (Nowak et al., 2001, Komarova and Nowak, 2003).

Frequency-dependent Moran processes have also been investigated by Taylor et al. (2004) and Nowak et al. (2004). Taylor et al. (2004) present analytic results for the case of fixed finite populations. They look at the process without mutations, and compare the probability that the process starting at the state where a single agent plays A (with the others playing B) is absorbed at the state “all A ” to the corresponding fixation probability in the neutral case where A gives the same payoff as B at each state. This neutral selection probability is $1/N$. Nowak et al. (2004) apply these results on the no-mutation process to the comparison of the strategies “Always defect” and “Tit for Tat” in the finitely repeated prisoner’s dilemma game. They also define evolutionary stability in finite populations and prove the “1/3 law”, which holds in the limit of weak selection for sufficiently large populations. The present paper differs in its assumption of recurrent mutations, and also in the nature of its conclusion: we provide a simple sufficient condition on the payoff matrix for the limit distribution to assign probability greater than $\frac{1}{2}$ to all agents playing A , and for it to assign probability 1 to this state when the population is large. A common assumption in the above models is that reproduction is asynchronous: at any time step, a single individual is chosen to reproduce. A related model of frequency-dependent selection in finite populations where all individuals reproduce at the same time but without mutations has been studied by Imhof and Nowak (2006).

2. The model

2.1. A frequency-dependent Moran process

Consider a population of N individuals playing a symmetric 2-player game with strategies A and B and

payoff matrix

	<i>A</i>	<i>B</i>
<i>A</i>	<i>a</i>	<i>b</i>
<i>B</i>	<i>c</i>	<i>d</i>

where *a, b, c* and *d* are positive. If *i* individuals play *A* and *N - i* play *B*, then the fitness of individuals using *A* is

$$f_i = \frac{a(i - 1) + b(N - i)}{N - 1}$$

and the fitness of those using *B* is

$$g_i = \frac{ci + d(N - 1 - i)}{N - 1}.$$

At every time step, one individual is chosen to reproduce. The chance of being chosen is proportional to fitness. That is, the probability that an individual using *A* is chosen is given by $if_i/[if_i + (N - i)g_i]$. We assume that with probability $\mu_{AB} > 0$, an *A*-offspring is a mutant which plays *B* instead of *A*, and with probability $\mu_{BA} > 0$, a *B*-offspring plays *A*. After reproduction, the offspring replaces a randomly chosen member of the population, so that the population size is constant. The process that describes the number of individuals that use *A* is a Markov process with state space $\{0, \dots, N\}$ and transition matrix (p_{ij}) , where

$$p_{ij} = 0 \text{ if } |i - j| > 1, \quad p_{01} = \mu_{BA} = 1 - p_{00},$$

$$p_{N,N-1} = \mu_{AB} = 1 - p_{NN},$$

and for $i = 1, \dots, N - 1$,

$$p_{i,i+1} = \frac{if_i(1 - \mu_{AB}) + (N - i)g_i\mu_{BA}}{if_i + (N - i)g_i} \frac{N - i}{N},$$

$$p_{i,i-1} = \frac{if_i\mu_{AB} + (N - i)g_i(1 - \mu_{BA})}{if_i + (N - i)g_i} \frac{i}{N},$$

$$p_{i,i} = 1 - p_{i,i-1} - p_{i,i+1}.$$

Because of the presence of mutations, this process is ergodic, with a unique invariant distribution that we denote by $\pi(k) = \pi(k; \mu_{AB}, \mu_{BA})$, $k = 0, \dots, N$.

We are interested in the small mutation limit of the invariant distribution:

$$\pi^*(k) = \lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \pi(k; \mu_{AB}, \mu_{BA}), \quad k = 0, \dots, N.$$

Throughout, the ratio of the mutation rates is kept fixed as they tend to zero. That is, for fixed numbers $\mu_{AB}^*, \mu_{BA}^* \in (0, 1)$, the limits are calculated as $\mu_{AB}, \mu_{BA} \rightarrow 0$ subject to

$$\frac{\mu_{AB}}{\mu_{BA}} = \frac{\mu_{AB}^*}{\mu_{BA}^*}.$$

π^* describes the long-run behavior of the Moran process when mutations are rare. As the Moran process is a birth–death process, the invariant distribution can easily be written down explicitly, see e.g. Ewens (2004, p. 91).

For $k = 0, \dots, N$,

$$\pi(k) = \frac{\lambda(k)}{\lambda(0) + \dots + \lambda(N)}, \quad \lambda(k) = \prod_{i=0}^{k-1} \frac{p_{i,i+1}}{p_{i+1,i}}.$$

The empty product is defined to be 1. Denoting the transition probabilities of the Moran process without mutations by \hat{p}_{ij} , we have $\hat{p}_{00} = \hat{p}_{NN} = 1$ and for $i = 1, \dots, N - 1$,

$$\hat{p}_{i,i+1} = \frac{if_i}{if_i + (N - i)g_i} \frac{N - i}{N},$$

$$\hat{p}_{i,i-1} = \frac{(N - i)g_i}{if_i + (N - i)g_i} \frac{i}{N},$$

$$\hat{p}_{ii} = 1 - \hat{p}_{i,i-1} - \hat{p}_{i,i+1}.$$

For $k = 1, \dots, N - 1$,

$$\begin{aligned} \lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \frac{\lambda(k)}{\mu_{BA}} &= \lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \frac{\prod_{i=1}^{k-1} p_{i,i+1}}{\prod_{i=0}^{k-1} p_{i+1,i}} \\ &= \frac{\prod_{i=1}^{k-1} \hat{p}_{i,i+1}}{\prod_{i=0}^{k-1} \hat{p}_{i+1,i}} < \infty, \end{aligned}$$

and so $\lambda(k) \rightarrow 0$ as $\mu_{AB}, \mu_{BA} \rightarrow 0$. Since $\lambda(0) = 1$, $\pi(k) \leq \lambda(k)$. It follows that $\pi(k) \rightarrow 0$ for $k = 1, \dots, N - 1$. Thus, for very small mutation rates, the Moran process spends nearly all the time at one of the states 0 and *N*. It remains to determine

$$\pi_A^* = \pi^*(N) \quad \text{and} \quad \pi_B^* = \pi^*(0),$$

the limits of the fractions of time that the Moran process spends at the states “all *A*” and “all *B*”, respectively. We have

$$\begin{aligned} \lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \lambda(N) &= \lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \frac{\mu_{BA} \prod_{i=1}^{N-1} p_{i,i+1}}{\mu_{AB} \prod_{i=0}^{N-2} p_{i+1,i}} \\ &= \frac{\mu_{BA}^*}{\mu_{AB}^*} \prod_{i=1}^{N-1} \frac{\hat{p}_{i,i+1}}{\hat{p}_{i,i-1}}. \end{aligned}$$

Observing that $\hat{p}_{i,i+1}/\hat{p}_{i,i-1} = f_i/g_i$, we obtain

$$\lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \lambda(N) = \frac{\mu_{BA}^*}{\mu_{AB}^*} \gamma,$$

where

$$\gamma = \prod_{i=1}^{N-1} \frac{f_i}{g_i}.$$

It follows that

$$\begin{aligned} \lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \lambda(0) + \dots + \lambda(N) &= \lim_{\mu_{AB}, \mu_{BA} \rightarrow 0} \lambda(0) + \lambda(N) = 1 + \frac{\mu_{BA}^*}{\mu_{AB}^*} \gamma, \end{aligned}$$

and, therefore,

$$\pi_A^* = \frac{\mu_{BA}^* \gamma}{\mu_{BA}^* \gamma + \mu_{AB}^*}, \quad \pi_B^* = \frac{\mu_{AB}^*}{\mu_{BA}^* \gamma + \mu_{AB}^*}. \tag{1}$$

The following assertions on the limit distribution and its connection with γ are now obvious. To emphasize the dependence on the population size, write $\pi_A^*(N) = \pi_A^*$, $\pi_B^*(N) = \pi_B^*$, and $\gamma(N) = \gamma$.

Lemma 1. For every population size N ,

$$0 < \pi_A^*(N) < 1, \quad 0 < \pi_B^*(N) < 1.$$

Moreover,

$$\pi_A^*(N) > \frac{1}{2} \quad \text{if and only if} \quad \gamma(N) > \frac{\mu_{AB}^*}{\mu_{BA}^*}.$$

As N tends to ∞ , $\pi_A^*(N)$ converges to 1 or 0 according as $\gamma(N)$ converges to ∞ or 0, respectively.

These results provide the basis for our main conclusions in Section 3.

In deriving representation (1) of the limit distribution we made use of the simple explicit formula for the invariant distribution π . If the underlying game has more than two pure strategies, the invariant distribution cannot be expressed in a similarly simple form, and a different approach is required to determine the limit distribution. We now give an alternative heuristic derivation of (1), which is perhaps more intuitive than the direct calculation and extends easily to games with any number of pure strategies. The approach is based on a limit theorem of Fudenberg and Imhof (2005) and will be used in Section 4 to solve equilibrium selection problems for 3×3 games.

Under the Moran process without mutations, the states 0 and N are absorbing and the other states are transient. Let ρ_{AB} denote the probability of absorption at state 0, if the process is initially at state $N - 1$. Thus, ρ_{AB} is the probability that a single individual that plays B takes over a population where everyone else plays A . Define ρ_{BA} analogously. Since, for small mutation rates, the Moran process spends nearly all the time at one of the states 0 and N , we consider the Markov chain obtained from the Moran process with mutations by ignoring all periods where it is in a state other than 0 or N . For this embedded chain the probability of a transition from 0 to N is given by the probability that a mutation from B to A occurs times the probability that a single A -player takes over the whole population. For small mutation rates, the second probability is close to the corresponding fixation probability of the no-mutation process, ρ_{BA} . This suggests that the limit of the ergodic distribution of the embedded chain, i.e. (π_A^*, π_B^*) , is given by the ergodic distribution of a Markov chain with transition matrix

$$\begin{pmatrix} 1 - \mu_{BA}^* \rho_{BA} & \mu_{BA}^* \rho_{BA} \\ \mu_{AB}^* \rho_{AB} & 1 - \mu_{AB}^* \rho_{AB} \end{pmatrix}.$$

Hence

$$\pi_A^* = \frac{\mu_{BA}^* \rho_{BA}}{\mu_{BA}^* \rho_{BA} + \mu_{AB}^* \rho_{AB}},$$

$$\pi_B^* = \frac{\mu_{AB}^* \rho_{AB}}{\mu_{BA}^* \rho_{BA} + \mu_{AB}^* \rho_{AB}}. \tag{2}$$

That these formulas are indeed correct follows from Theorem 1 of Fudenberg and Imhof (2005). Note that their Assumptions 1–4 are obvious from the definition of the transition probabilities at hand, and their Assumption 5 follows from the fact that $\rho_{AB}, \rho_{BA}, \mu_{AB}^*$ and μ_{BA}^* are positive. Since the Moran process is a birth–death process, the fixation probabilities have a simple explicit expression, see e.g. Ewens (2004, p. 90), and using that $\hat{p}_{j,j-1}/\hat{p}_{j,j+1} = g_j/f_j$, we obtain

$$\begin{aligned} \rho_{BA} &= \frac{1}{1 + \sum_{i=1}^{N-1} \prod_{j=1}^i (g_j/f_j)}, \\ \rho_{AB} &= \frac{\prod_{j=1}^{N-1} (g_j/f_j)}{1 + \sum_{i=1}^{N-1} \prod_{j=1}^i (g_j/f_j)}. \end{aligned} \tag{3}$$

Thus, $\rho_{BA}/\rho_{AB} = \gamma$, and inserting this relation into (2) gives representation (1). The reason that this approach can also be applied to games with more than two pure strategies is that it requires only the fixation probabilities of the restriction of the no-mutation process to the edges of the state space where at most two pure strategies are present.

2.2. A closely related process

Here is another model that gives rise to the same long-run outcome: each period, a single agent is randomly chosen from the population, and matched (without replacement) with a single second agent to play the underlying game. The first agent reproduces with probability equal to its realized payoff divided by a scale factor z that is larger than $\max\{a, b, c, d\}$. If the agent reproduces, it replaces a randomly chosen member of the entire population. Thus, for example the probability of an increase in the number of A -players is the probability an A is chosen times the expected payoff of an A times the probability a B is replaced divided by z , so that

$$p'_{i,i+1} = \frac{i}{N} \frac{(i-1)a + (N-i)b}{N-1} \frac{N-i}{N} \frac{1}{z},$$

and

$$\frac{p'_{i,i+1}}{p'_{i,i-1}} = \frac{f_i}{g_i} = \frac{p_{i,i+1}}{p_{i,i-1}},$$

so the new process differs from the original one only in its speed.

3. Favored and selected strategies

We now study how the limit distribution $(\pi_A^*(N), \pi_B^*(N))$ depends on the size of the population and on the payoff matrix. We will say that strategy i is *favored* by the Moran process if $\pi_i^*(N) > \frac{1}{2}$, and that strategy i is *selected* if $\lim_{N \rightarrow \infty} \pi_i^*(N) = 1$. To determine favored and selected

strategies by means of Lemma 1, we substitute the values of the payoff functions at each state into the equation for γ :

$$\begin{aligned} \gamma(N) &= \frac{[(N-1)b][(N-2)b+a] \dots [2b+(N-3)a][b+(N-2)a]}{[c+(N-2)d][2c+(N-3)d] \dots [(N-2)c+d][(N-1)c]} \\ &= \frac{[b+(N-2)a][2b+(N-3)a] \dots [(N-2)b+a][(N-1)b]}{[c+(N-2)d][2c+(N-3)d] \dots [(N-2)c+d][(N-1)c]} \end{aligned} \tag{4}$$

As a preliminary step, note that when $N = 2$, $\gamma(N) = b/c$, so that strategy A is favored when $b > c$ and both sorts of mutations are equally likely. Intuitively, every time there is a mutation, the system moves to the state with one A and one B , and at this state, the relative payoffs of the two strategies are determined by their payoffs when playing each other. Our results will follow from a more detailed analysis of the ratio in (4). Note that multiplying all of the payoffs by the same constant has no effect on $\gamma(N)$, and so has no effect on the limit of the ergodic distributions. However, $\gamma(N)$ may change when a constant is added to all of the payoff functions.

To cut down on cases, we assume now that $b > c$. This can be done by re-labeling the strategies except for the knife-edge case where $b = c$. We first deal with favored strategies and focus on the case where both sorts of mutations are equally likely. We return to the general case of possibly different mutation rates when considering the large population limit in Theorem 2.

Theorem 1. Suppose that $\mu_{AB}^* = \mu_{BA}^*$.

- (a) If $b > c$ and $a > d$, then $\pi_A^*(N) > \frac{1}{2}$ for all N .
- (b) If $b > c$ and $a < d$, then whether $\pi_A^*(N) > \frac{1}{2}$ may depend on the population size. A sufficient condition for $\pi_A^*(N) > \frac{1}{2}$ is $b - c > (N - 2)(d - a)$.

Proof. In case (a), the first term in the product in the numerator of the final ratio in (4) exceeds the corresponding term in the denominator, as does the second, etc., so that $\gamma(N) > 1$, and Lemma 1 implies that $\pi_A^*(N) > \frac{1}{2}$.

In case (b), if $b + (N - 2)a > c + (N - 2)d$, the pairwise comparison of terms again shows that $\gamma(N) > 1$, and $b + (N - 2)a > c + (N - 2)d$ is equivalent to $b - c > (N - 2)(d - a)$. \square

Theorem 1 gives results for any N , but here both strategies have positive weights, and the ratio of the weights depends on the ratio of the mutation probabilities. However, the effects of the payoff matrix overwhelm the effect of the ratio of the mutation probabilities when the population is sufficiently large, which is one reason for our interest in the case of population sizes tending to infinity. A second reason for studying this case is to see how well it is captured by the replicator dynamic, which corresponds to the behavior of the system in a continuum population. That is, the replicator is the mean field of this system.

Recall that in a 2×2 game, a strategy is risk dominant if it is the unique best response to the distribution $(\frac{1}{2}A, \frac{1}{2}B)$. Thus, strategy A is risk dominant if $a + b > c + d$, and B is

risk dominant if the reverse inequality holds. Strategy A is Pareto-dominant if $a > d$, and B is Pareto-dominant if $d > a$.

Theorem 2. (a) If $b > c$ and $a > d$, then $\lim_{N \rightarrow \infty} \pi_A^*(N) = 1$.

(b.1) If $b > d > a > c$, then $\lim_{N \rightarrow \infty} \pi_A^*(N) = 1$.

(b.2) If $d > b > c > a$, then $\lim_{N \rightarrow \infty} \pi_A^*(N) = 0$.

(b.3) If $d > b > a > c$, $d > a > b > c$, then there are two pure-strategy Nash equilibria, and $\lim_{N \rightarrow \infty} \pi_A^*(N)$ is either 1 or 0 as

$$\int_0^1 \ln[b + (a - b)x] dx$$

is greater or less than

$$\int_0^1 \ln[d + (c - d)x] dx.$$

The risk-dominant equilibrium need not be selected, even if it is also Pareto-dominant.

(b.4) If $b > c > d > a$ or $b > d > c > a$, then $\lim_{N \rightarrow \infty} \pi_A^*(N)$ is either 1 or 0 and which case obtains depends on the same integral condition as in (b.3).

Proof. In case (a), the ratio of each pair of terms in $\gamma(N)$ is bounded away from 1, so $\gamma(N) \rightarrow \infty$ as $N \rightarrow \infty$. Thus, Lemma 1 implies that $\lim_{N \rightarrow \infty} \pi_A^*(N) = 1$. In subcases (b.1) and (b.2), we examine the expression in the first line of (4): in subcase (b.1), every term in the numerator exceeds the corresponding term in the denominator, and in (b.2) the reverse is true provided that $N - 2 > (b - c)/(d - b)$ and $N - 2 > (b - c)/(c - a)$. The argument for large N in subcases (b.3) and (b.4) involves approximating $\gamma(N)$ by the ratio of two integrals, using

$$\begin{aligned} &\prod_{j=0}^{N-1} \left[a \frac{j}{N-1} + b \left(1 - \frac{j}{N-1} \right) \right] \\ &= \exp \left(N \int_0^1 \ln[a + (b - a)x] dx \right) + O(1); \end{aligned}$$

the details are in the Appendix. \square

The class of games in case (a) is composed of games where A is strictly dominant ($a > b > c > d$, $a > b > d > c$, $b > a > d > c$, and $b > a > c > d$), coordination games where A is both a Pareto-dominant equilibrium and risk-dominant ($a > d > b > c$), and “hawk–dove games” with two asymmetric equilibria and an equilibrium in mixed strategies ($b > c > a > d$). When A is strictly dominant, it is selected by the deterministic replicator dynamic from any initial position, and it is not surprising that the same thing happens here as N goes to infinity. It is similarly unsurprising that A is selected when it is both risk- and

Pareto-dominant: although both of the Nash equilibria are asymptotically stable in the deterministic replicator dynamic, past work on stochastic evolutionary models has always selected any strategy that is both risk- and Pareto-dominant.

In case (b), $b > c$, so we expect that A will be favored for small N . In subcase (b.1), A is the dominant strategy, so this tendency is reinforced for large N . In subcase (b.2), B is dominant, and is selected for large N , but A is favored for small N ; this is the “spite” effect. Subcase (b.3), where A and B are both pure-strategy equilibria, is more complex. Past work has concluded that the long-run distribution is concentrated on the risk-dominant equilibrium. However, in the present setting the long-run distribution need not be concentrated on the risk-dominant equilibrium. This is easiest to see by considering a game where $d > b > a > c$, so B is payoff dominant, and $a + b = c + d$ so that neither strategy is risk dominant. Then the two integrals are the expectations of the logarithm of two random variables with the same mean. Because the log is a concave function, the expected value of the log is reduced by a mean-preserving spread, and so A is selected because $b - a < d - c$. Intuitively, the condition $a + b = c + d$ implies that the two strategies are equally fit at the point $i = N/2$, but the support of the long-run distribution depends on the transition probabilities at every state, and these are not determined by the value of the payoff functions at the midpoint.

Both case (a) and case (b) include subcases that correspond to “hawk–dove games,” that is symmetric games with two asymmetric pure-strategy equilibria. Since we are working in a one-population model, the asymmetric equilibria in the hawk–dove case cannot arise. The conclusion that the process spends almost all of its time in homomorphous states is a consequence of our focus on the limits of the ergodic distribution for the case of vanishingly small mutation rates. The results of Benaïm and Weibull (2003a) suggest that for any fixed small mutation rate, the large population limit of the invariant distribution is concentrated near the mixed equilibrium.

We summarize the above discussion of our results for large populations by regrouping cases according to more traditional game-theoretic criteria.

Corollary. *If the game has a strictly dominant strategy, the probability assigned to this strategy by the limit distribution converges to 1 as N goes to infinity. If the game has two strict Nash equilibria, then except for knife-edge cases there is an equilibrium to which the limit distribution assigns probability converging to 1 as N goes to infinity, but the risk-dominant equilibrium need not be selected. Which equilibrium is selected depends on the conditions specified in Theorem 2.*

4. 3×3 coordination games

Now consider a 3×3 symmetric game with pure strategies 1, 2, 3, and strictly positive payoff matrix

$(a_{ij})_{i,j=1}^3$, where a_{ij} is the payoff for strategy i playing strategy j . As before, there is a population of N agents; the state $x = (x_1, x_2, x_3)$ of the system consists of the numbers of agents using each strategy. We denote the payoff of strategies 1, 2, and 3 by f_x , g_x , and h_x , respectively, where for example

$$f_x = \frac{(x_1 - 1)a_{11} + x_2a_{12} + x_3a_{13}}{N - 1}.$$

The no-mutation process is again constructed by assuming that each agent playing strategy i has a number of offspring equal to its payoff, with one offspring chosen at random to replace a random member of the current population, so that e.g. the probability that a 1-user replaces a 2-user is

$$\frac{x_1 f_x}{x_1 f_x + x_2 g_x + x_3 h_x} \frac{x_2}{N},$$

and the ratio of this probability to that of a move in the other direction is f_x/g_x , just as it was in the two-strategy case. Let $\rho_{ji}(N)$ be the probability that the no-mutation process is absorbed at the homogeneous state where all agents play i , starting from the state where only one agent plays i and all the rest play j . Since the no-mutation process never introduces an extinct strategy, $\rho_{ji}(N)$ is the same as if the game only had these two strategies.

We add mutations to the system by supposing that there is common probability μ that the offspring of an i -strategist is a j -strategist, $j \neq i$. Again, the process has a unique ergodic distribution for every $\mu > 0$. We let $\pi^*(N) = (\pi_1^*(N), \pi_2^*(N), \pi_3^*(N))$ denote the limit of the ergodic distributions as $\mu \rightarrow 0$. As before we represent this as a probability distribution over the homogeneous states instead of as a distribution over the entire state space, so that $\pi_i^*(N)$ is the probability that the limit distribution gives to the state “all i .”

The procedure for calculating the limit distribution based on the embedded Markov chain on the pure states, as outlined in Section 2, can also be applied in the present 3×3 case. This yields the following expression for the limit distribution:

$$\pi_i^*(N) = \frac{(\rho_{ji}(N)\rho_{ki}(N) + \rho_{ji}(N)\rho_{kj}(N) + \rho_{ki}(N)\rho_{jk}(N))_{i \neq j \neq k}}{\sum_{r=1}^3 (\rho_{jr}(N)\rho_{kr}(N) + \rho_{jr}(N)\rho_{kj}(N) + \rho_{kr}(N)\rho_{jk}(N))_{r \neq j \neq k}}, \tag{5}$$

see Example 2 of Fudenberg and Imhof (2005). Note that the formula given is that for the invariant distribution of a Markov chain on the three states 1, 2, 3, whose off-diagonal transition probabilities are given by the ρ 's. To find the selected strategy, we have to analyze the asymptotic behavior of the $\rho_{ij}(N)$ as $N \rightarrow \infty$.

We now specialize to the coordination-game case, where each of the pure strategies corresponds to a symmetric Nash equilibrium, so that $a_{ii} > a_{ji}$ for all i and $j \neq i$. From the 2×2 case, we know that the $\rho_{ij}(N)$ will all be positive, so (5) implies that the limit distribution will give positive probability to each strategy in a population of fixed size.

Our goal is to determine the behavior of the limit distribution as the population size N goes to infinity. To this end let

$$z_{ji}^* = \frac{a_{jj} - a_{ij}}{a_{jj} - a_{ij} + a_{ii} - a_{ji}} = 1 - z_{ij}^*$$

this is the weight given to strategy i in the mixed equilibrium in the 2×2 subgame corresponding to strategies i and j , and is strictly between 0 and 1 because we have specialized to coordination games. Define

$$\phi_{ji}(z) = \ln \left[\frac{a_{ji}z + a_{ij}(1 - z)}{a_{ii}z + a_{ij}(1 - z)} \right];$$

this is the logarithm of the ratio of the expected payoffs of j and i in an infinite population where fraction z plays i and all the rest play j . Finally define

$$\beta_{ji} = \int_0^{z_{ji}^*} \phi_{ji}(z) dz;$$

note that this is strictly positive. The following result characterizes the selected strategy in terms of the integrals β_{ji} , where only knife-edge cases are disregarded.

Theorem 3. *For $i = 1, 2, 3$ let $\alpha_i = \min\{\beta_{ji} + \beta_{ki}, \min_j(\beta_{ji} + \beta_{kj})\}_{i \neq j \neq k}$. Suppose i is the strategy for which $\alpha_i < \min(\alpha_j)_{j \neq i}$. Then $\lim_{N \rightarrow \infty} \pi_i^*(N) = 1$.*

The proof of Theorem 3 is in the Appendix. An important step is to characterize the limit on N of the $\rho_{ij}(N)$ in terms of integrals of logarithms of the relative payoff of strategies i and j , using an argument that is similar in spirit but more complex than that in the proof of Theorem 2, case (b.3).

We now apply our results to study selective language dynamics. Consider a 3×3 language game: strategies 1, 2, 3 correspond to three different languages, and a_{ij} is the payoff for someone who speaks language i communicating with someone who speaks language j . Thus, the payoffs measure how successful communication is, and we assume the received payoffs contribute to biological fitness (Nowak et al., 2001). Children learn the language of their parents, possibly with mistakes. The mutation probability μ is the probability that the child switches to a language that is different from its parents' language. Alternatively, one may assume that successful communication results in cultural fitness, so that individuals that have to choose a language are more likely to "imitate" speakers of a language with high current payoff. The interpretation of the Moran process as an imitation process will be discussed in detail in the next section.

Applying Theorem 3, we can determine which language will be spoken in a large population when mistakes are rare. We consider three examples, where the payoff

matrices are given by

$$\begin{aligned} \text{(i) } A_1 &= \begin{pmatrix} 10 & 1 & 1 \\ 1 & 9 & 1 \\ 1 & 1 & 8 \end{pmatrix}, & \text{(ii) } A_2 &= \begin{pmatrix} 10 & 1 & 1 \\ 1 & 9 & 7 \\ 1 & 8 & 8 \end{pmatrix}, \\ \text{(iii) } A_3 &= \begin{pmatrix} 10 & 1 & 1 \\ 1 & 9 & 1 \\ 2.5 & 2.5 & 8 \end{pmatrix}, \end{aligned}$$

respectively. In the first case, communication between individuals using different languages is equally inefficient, and for individuals using the same language, coherence is best if language 1 is used. The α_i from Theorem 3 are $\alpha_1 = 0.7997$, $\alpha_2 = 0.8862$, and $\alpha_3 = 0.9816$, so that language 1 will be spoken, as was to expected.

In the second case, users of language 2 communicate efficiently with users of language 3 and vice versa, but communication between speakers of language 1 and speakers of one of the other two languages is still inefficient. Now $\alpha_1 = 0.4041$, $\alpha_2 = 0.5448$, and $\alpha_3 = 0.5422$, so that again language 1 will be selected.

In the third case, speakers of language 3 receive a somewhat higher payoff when communicating with speakers of language 1 than they did in the first case. Now $\alpha_1 = 0.7005$, $\alpha_2 = 0.7870$, and $\alpha_3 = 0.6508$, so that language 3 will be selected, even though coherence would be higher if language 1 were chosen.

To compare the Moran process with the standard replicator dynamics note first that under the replicator dynamics, each homogeneous state is asymptotically stable. Which language is selected, that is, to which vertex the trajectory converges, depends on the initial state. For payoff matrix A_1 , language 1 has, of course, the largest basin of attraction. The basins of language 1, 2 and 3 cover, respectively, 39%, 33% and 28% of the area of the whole state space. For payoff matrix A_2 , the proportions are 31%, 35% and 34%. Note that the behavior of the replicator dynamics is determined by the vector field on the interior of the state space. Selection takes place as the population moves through states where each language is spoken. Thus, the fact that language 2 has the largest basin corresponds to the fact that users of this language communicate efficiently with users of the other languages, whereas users of language 1 communicate efficiently only with users of the same language. On the other hand, the behavior of the stochastic process is determined by the transition probabilities of moving from one pure state to another after a single mutation, where during a transition, only two languages are spoken. Here language 1 is selected as it profits from its high payoff in a population where only this language is used, resulting in a small probability of moving away from this state. If the payoffs are given by A_3 , the proportions of the basins of attraction under the replicator dynamics are 35.2%, 29.9% and 34.9%. Again,

the language that is selected by the Moran process, language 3, does not have the largest basin.

5. Related models

In this section we show that our frequency-dependent Moran process can also be viewed as an imitation process, and we compare our model with related economic imitation models. The interpretation as an imitation process is particularly natural in the context of language dynamics as discussed in Section 4. Rhode and Stegeman (2001) discuss the relevance of imitation from an economic perspective.

Consider again a population of N agents playing the 2×2 game of Section 2. Every agent is matched with every opponent from the population, so that when i agents play A , the average payoff of individuals using A is $f_i = [(i-1)a + b(N-i)]/(N-1)$ and the average payoff to individuals using B is $g_i = [ci + d(N-1-i)]/(N-1)$. Note that, as in Kandori et al. (1993) and many subsequent papers in this field, the two types of agent face a slightly different distribution of opponents' play. This is what underlies the difference between maximizing relative payoff and maximizing absolute payoff and generates "spite" effects. Rhode and Stegeman (1996) analyze the effect of spite in the "Darwinian" model of Kandori et al. (1993), which supposes that the unperturbed dynamic is deterministic; they show that even allowing for spite the risk-dominant equilibrium is selected in large populations. Schaffer (1988) suggests an alternative concept of evolutionary stability in finite populations that takes the possibility of spite effects into account. See Schaffer (1989), Crawford (1991) and Alós-Ferrer et al. (2000) for discussion of the spite effect in economic models, and Vega-Redondo (1997) for the connection to Walrasian behavior. Recent papers in this direction include Possajennikov (2003), Alós-Ferrer and Ania (2005), and Leininger (2005).

By the law of large numbers, f_i and g_i can be regarded as average payoffs when the agents are infinitely often randomly paired. Robson and Vega-Redondo (1996) analyze a model where at each period, agents are randomly matched in pairs for a finite number of rounds, and evolution is governed by the realized payoffs of the strategies, which depend on the outcome of the matching process. In their model, the limit as the number of rounds goes to infinity of the limit distributions (i.e. of the limit of the ergodic distributions as the mutation rate vanishes) can be different than the limit distribution with an infinite number of rounds, but that is not the case here.

Suppose the distribution of strategies in the population evolves as follows. Agents decide whose strategy to imitate based on the prevailing payoffs and also on the relative popularity of each strategy. The decision is described by a random variable whose distribution is given by an "updating function" that depends only on the current state, that is, the numbers of agents using each choice. The

updating function we study has two parts, the "base rate" updating rule and a lower-frequency probability of "mutation." In the base-rate or "unperturbed" process, each period, one randomly chosen individual re-evaluates its choice, with the probability of choosing a given strategy equal to the total payoff of players using that strategy divided by the total payoff in the population, so that the choice depends on both the payoff of each strategy and on the strategy's popularity, and there is non-zero probability of moving to the strategy with the lowest current payoff. However, there is a probability μ_{AB} that an agent who intends to play A plays B instead of A , and a probability μ_{BA} that an agent who intends to play B plays A . This derivation gives rise to the same Moran process as the one we defined in Section 2.

Giving weight to popularity as well as current payoffs is a rule of thumb that allows current choices to in effect depend on past states of the system. Ellison and Fudenberg (1993, 1995) show that such rules can be socially beneficial because popularity incorporates information about long-run performance. Even if such rules do not perform well in the stark setting of this model, boundedly rational agents might still use them because they perform well in general. Binmore and Samuelson (1997) consider a model in which agents adopt a new strategy according to the current popularities. In the process they consider, as in ours, one agent is chosen at random to re-evaluate its choice. The agent first decides whether it is "dissatisfied," where the probability of dissatisfaction depends on the state. When an agent is dissatisfied, it adopts the strategy of a randomly chosen member of the population (including themselves in the population for this purpose). As in our model, there is a positive probability that the agent switches to a strategy that gives a lower payoff. They give a general characterization of the limit distribution as a function of the "dissatisfaction functions," and then specialize to a particular "aspiration and imitation model" to relate the limit distribution in large populations to the payoff matrix of the game. Unlike us, they do not relate the payoff matrix to the limit distribution for finite populations. They determine the limit distribution by substituting into the detailed balance equation, but they note that it will require other methods to handle the 3×3 case. Our alternate derivation may be more revealing and also has the advantage of extending to the case of more than two strategies.

6. Discussion

In our model, the risk-dominant equilibrium need not be selected in a 2×2 coordination game, even in the limit of large populations. In contrast, the risk-dominant equilibrium is selected even for finite populations in the work of Kandori et al. (1993), Young (1993), Robson and Vega-Redondo (1996), Sandholm (1998), Maruta (2002) and Blume (2003). This is because those papers analyze "essentially deterministic" no-mutation processes, where

the agent selected to move typically plays a best response to the observed state, or copies the action with the highest payoff. The processes return to the nearby equilibrium with probability 1 after any “small” number of mutations. As a consequence, these models predict that the limit distribution is a point mass, at least for sufficiently large (but finite) N . In Kandori et al. (1993) and Young (1993), the risk-dominant equilibrium is selected as the mutation rate goes to 0.

In contrast, we analyze a stochastic no-mutation process. Both the fact that $0 < \rho_{AB}(N) < 1$ for finite N here, and the fact that the ratio of the mutation probabilities has an impact on the relative probability of the two states “all A ” and “all B ” even when this ratio is bounded away from zero and infinity, is related to the fact that a single mutation can lead to a transition from one absorbing state of the no-mutation process to the other. Thus, the equilibrium that is selected depends on the “expected speed of the flow” at every point in the state space, and two coordination games with the same mixed-strategy equilibrium (and hence the same basins for the best-response and replicator dynamic) can have systematically different speeds at other states. By replicator dynamic we mean the standard version with linear payoff functions, which is due to Taylor and Jonker (1978) and Hofbauer et al. (1979). Moreover, even a strategy that is both a risk-dominant and a payoff-dominant equilibrium need not be selected, in contrast to the result of Binmore and Samuelson (1997). Both their paper and Fudenberg and Harris (1992) discuss why one should expect stochastic stability to depend on the “speed of the flow” as well as on the expected direction of motion. The influence of spatial population structures on the size of the basin of attraction of a risk-dominant strategy in best-response models has been studied by Blume (1993, 1995).

Since in our model, the risk-dominant equilibrium need not be selected in the 2×2 case, there is no reason to expect a $\frac{1}{2}$ -dominant equilibrium (Morris et al., 1995) to be selected in a 3×3 coordination game. Thus, the equilibrium selection can be different than that of the models of Kandori and Rob (1998) and of Kandori et al. (1993), where in the 2×2 case the risk-dominant equilibrium is selected. The easiest examples are where strategy 3 risk-dominates both strategy 1 and strategy 2 in pairwise comparisons (so it is “pairwise risk dominant”) but where strategy 1 or 2 is selected over strategy 3 as $N \rightarrow \infty$ in the 2×2 games. This is the case for the following payoff matrix:

$$\begin{pmatrix} 16 & 5 & 8 \\ 1 & 16 & 8 \\ 1 & 1 & 24 \end{pmatrix}.$$

Here strategy 1 is selected, i.e. $\lim_{N \rightarrow \infty} \pi_1^*(N) = 1$.

To place our findings about large populations into perspective, it is helpful to note that the mean field of our

process (that is, expected motion) converges to that of the standard replicator dynamic as the population becomes infinite and time is rescaled appropriately. Thus, the results of Benaïm and Weibull (2003a, b) show that for large N the state of the process will, with high probability, remain close to the path of the deterministic replicator dynamic for some large time T . One can instead obtain a diffusion approximation if the payoffs in each interaction are scaled with the population size in the appropriate way, see e.g. Binmore et al. (1995) and Ewens (2004). Note that in small populations the mean field in our model can be very different from that of the replicator due to the difference between maximizing absolute and relative performance. However, even for large N the asymptotics of the stochastic system depend on the details of the stochastic structure, and can differ from those of the deterministic replicator dynamic. Moreover, our finding that the risk-dominant equilibrium need not be selected in a 2×2 game shows that the asymptotic behavior of our system can differ from that of the forms of the stochastic replicator dynamics studied by Foster and Young (1990, 1997) and Fudenberg and Harris (1992).

Fogel et al. (1998) and Ficici and Pollack (2000) report some simulations of the “frequency-dependent roulette wheel” selection dynamic, which is equivalent to the generalized Moran process we analyze. Fogel et al. emphasize that the finite population results can be very different than the predictions of the replicator equation, while Ficici and Pollack argue that the two models make fairly similar predictions in the hawk–dove game.

Acknowledgments

We thank the referees for several suggestions that helped to improve the presentation of the paper. The material is based on work supported by NSF grant SES-0426199 and by the John Templeton Foundation. The Program for Evolutionary Dynamics is sponsored by Jeffrey Epstein. L. A. I. thanks the program for its hospitality.

Appendix

Proof of Theorem 2, case (b.3). Our goal is to characterize the behavior of $\gamma(N)$ as $N \rightarrow \infty$, where

$$\begin{aligned} \gamma(N) &= \frac{[(N-1)b][(N-2)b+a] \dots [2b+(N-3)a][b+(N-2)a]}{[c+(N-2)d][2c+(N-3)d] \dots [(N-2)c+d][(N-1)c]} \\ &= \frac{d(N-1)a[b+(N-2)a] \dots [(N-2)b+a][(N-1)b]}{a(N-1)d[c+(N-2)d] \dots [(N-2)c+d][(N-1)c]} \\ &= \frac{d}{a} \prod_{j=0}^{N-1} \left[\frac{aj+b(N-1-j)}{cj+d(N-1-j)} \right]. \end{aligned}$$

We will approach this by comparing the numerator of the expression to the denominator. We rewrite the

numerator as

$$d \prod_{j=0}^{N-1} \left[a \frac{j}{N-1} + b \left(1 - \frac{j}{N-1} \right) \right] \\ = d \exp \left\{ N \sum_{j=0}^{N-1} \frac{1}{N} \ln \left[a \frac{j}{N-1} + b \left(1 - \frac{j}{N-1} \right) \right] \right\},$$

and the denominator as

$$a \prod_{j=0}^{N-1} \left[c \frac{j}{N-1} + d \left(1 - \frac{j}{N-1} \right) \right].$$

We then approximate the numerator by

$$d \exp \left\{ N \int_0^1 \ln[a + (b - a)x] dx \right\} + O(1)$$

and the denominator by

$$a \exp \left\{ N \int_0^1 \ln[c + (d - c)x] dx \right\} + O(1).$$

For large N , this comparison is determined by the comparison of the integrals, and $\gamma(N)$ converges to either 0 or infinity. Since

$$\int \ln[a + (b - a)x] dx = x \ln[a + (b - a)x] \\ - x + \frac{a}{b - a} \ln[a + (b - a)x],$$

$$\int_0^1 \ln[a + (b - a)x] dx = \frac{b}{b - a} \ln b - \frac{a}{b - a} \ln a - 1.$$

So the question becomes whether

$$\frac{b \ln b - a \ln b}{b - a} > \frac{d \ln d - c \ln c}{d - c}. \quad \square$$

Claim. *There are cases with $a + b < c + d$ but*

$$\frac{b \ln b - a \ln a}{b - a} > \frac{d \ln d - c \ln c}{d - c}.$$

Take $c = 1$, $b = 8$, $a = 16$, $d = 24$. Then the LHS = $5 \ln 2 \approx 3.465$; the RHS = $(24 \ln 24)/23 \approx 3.307$.

We now turn to the proof of Theorem 3. Let ρ_{ij} , ϕ_{ij} , z_{ij}^* , and β_{ij} be as defined in Section 4. As a preparation, we prove the following result which describes the asymptotic behavior of the fixation probabilities $\rho_{ji}(N)$ for $N \rightarrow \infty$ in terms of the integrals β_{ji} .

Lemma 2. *There are constants $0 < m < M < \infty$ such that for all $N \geq 2$, and all pairs of distinct strategies i, j , $m \leq N^{1/2} \exp(N\beta_{ji})\rho_{ji}(N) \leq M$.*

Proof. Fix a pair of distinct strategies i and j . From (3),

$$\frac{1}{\rho_{ij}(N)} = 1 + \sum_{v=1}^{N-1} \prod_{\kappa=1}^v \frac{\kappa a_{ij} + (N - 1 - \kappa)a_{ii}}{(\kappa - 1)a_{jj} + (N - \kappa)a_{ji}} \\ = 1 + \sum_{v=1}^{N-1} \frac{va_{jj} + (N - 1 - v)a_{ji}}{(N - 1)a_{ji}} \\ \times \prod_{\kappa=1}^v \frac{\kappa a_{ij} + (N - 1 - \kappa)a_{ii}}{\kappa a_{jj} + (N - 1 - \kappa)a_{ji}} \\ = 1 + \sum_{v=1}^{N-1} \frac{va_{jj} + (N - 1 - v)a_{ji}}{(N - 1)a_{ji}} \\ \times \exp \left\{ \sum_{\kappa=1}^v \phi_{ij} \left(\frac{\kappa}{N - 1} \right) \right\}. \quad (6)$$

Set $\psi(\xi) = \int_0^\xi \phi_{ij}(z) dz$ and $K_1 = \sup\{|\phi'_{ij}(z)| : 0 \leq z \leq 1\}$. Since all payoffs are positive, $K_1 < \infty$. Thus, by the mean value theorem,

$$\phi_{ij} \left(\frac{\kappa}{N - 1} \right) \leq (N - 1) \int_{(\kappa-1)/(N-1)}^{\kappa/(N-1)} \phi_{ij}(z) dz \\ + \frac{K_1}{N - 1} \quad \text{for } \kappa = 1, \dots, N - 1$$

and so, for $v = 1, \dots, N - 1$,

$$\exp \left\{ \sum_{\kappa=1}^v \phi_{ij} \left(\frac{\kappa}{N - 1} \right) \right\} \\ \leq \exp \left\{ (N - 1) \psi \left(\frac{v}{N - 1} \right) + K_1 v \right\}.$$

The positivity of the payoffs also implies that ψ is Lipschitz continuous with Lipschitz constant $K_2 = \sup\{|\phi_{ij}(z)| : 0 \leq z \leq 1\}$. Thus,

$$\exp \left\{ (N - 1) \psi \left(\frac{v}{N - 1} \right) \right\} \\ = (N - 1) \int_{(v-1)/(N-1)}^{v/(N-1)} \exp \left\{ (N - 1) \psi \left(\frac{v}{N - 1} \right) \right\} d\xi \\ \leq e^{K_2} (N - 1) \int_{(v-1)/(N-1)}^{v/(N-1)} \exp \{ (N - 1) \psi(\xi) \} d\xi$$

for $v = 1, \dots, N - 1$.

Note also that

$$\frac{va_{jj} + (N - 1 - v)a_{ji}}{(N - 1)a_{ji}} \leq \max \left\{ \frac{a_{jj}}{a_{ji}}, 1 \right\} \\ \text{for all } v = 1, 2, \dots, N - 1.$$

Substituting into (6), it follows that

$$\frac{1}{\rho_{ij}(N)} \leq 1 + \max \left\{ \frac{a_{jj}}{a_{ji}}, 1 \right\} e^{K_1 + K_2} (N - 1) \\ \times \int_0^1 \exp \{ (N - 1) \psi(\xi) \} d\xi.$$

A similar argument shows that

$$\frac{1}{\rho_{ij}(N)} \geq 1 + \min \left\{ \frac{a_{ij}}{a_{ii}}, 1 \right\} e^{-K_1 - K_2(N-1)} \times \int_0^1 \exp[(N-1)\psi(\xi)] d\xi.$$

To determine the asymptotic behavior of the integral note that $\phi_{ij}(z) > 0$ for $0 \leq z < z_{ij}^*$ and $\phi_{ij}(z) < 0$ for $z_{ij}^* < z \leq 1$. Thus, $\psi(\xi) < \psi(z_{ij}^*) = \beta_{ij}$ for all $\xi \neq z_{ij}^*$. Moreover, $\psi''(z_{ij}^*) = \phi'_{ij}(z_{ij}^*) < 0$. Thus, the Laplace method for approximating integrals of the form $\int \exp[Nh(x)] dx$ as $N \rightarrow \infty$ (see e.g. de Bruijn, 1958, Chapter 4) yields that

$$\lim_{N \rightarrow \infty} \frac{\int_0^1 \exp[N\psi(\xi)] d\xi}{N^{-1/2} \exp(N\beta_{ij})} = \sqrt{-\frac{2\pi}{\phi'_{ij}(z_{ij}^*)}}.$$

It follows that there exist constants $0 < m_{ij} < M_{ij} < \infty$ such that $m_{ij} \leq N^{1/2} \exp(N\beta_{ij})\rho_{ij}(N) \leq M_{ij}$ for all $N \geq 2$. The assertion is obtained by considering all pairs of distinct strategies and taking $m = \min_{i \neq j} m_{ij}$ and $M = \max_{i \neq j} M_{ij}$. \square

Proof of Theorem 3. From (5),

$$\frac{\pi_i^*(N)}{\pi_l^*(N)} = \frac{(\rho_{ji}(N)\rho_{ki}(N) + \rho_{ji}(N)\rho_{kj}(N) + \rho_{ki}(N)\rho_{jk}(N))_{i \neq j \neq k}}{(\rho_{ji}(N)\rho_{kl}(N) + \rho_{jl}(N)\rho_{kj}(N) + \rho_{kl}(N)\rho_{jk}(N))_{l \neq j \neq k}}.$$

From Lemma 2, $m \leq N^{1/2} \exp(N\beta_{ji})\rho_{ji}(N) \leq M$, so we can bound each of the six terms in this ratio, for example

$$\frac{m^2}{N \exp[N(\beta_{ji} + \beta_{ki})]} \leq \rho_{ji}(N)\rho_{ki}(N) \leq \frac{M^2}{N \exp[N(\beta_{ji} + \beta_{ki})]}.$$

Since all six terms are positive, if $\alpha_j > \min\{\alpha_1, \alpha_2, \alpha_3\}$, $\lim_{N \rightarrow \infty} \pi_j^*(N) = 0$. \square

References

Alós-Ferrer, C., Ania, A.B., 2005. The evolutionary stability of perfectly competitive behavior. *Econ. Theory* 26, 497–516.
 Alós-Ferrer, C., Ania, A.B., Schenk-Hoppé, K.R., 2000. An evolutionary model of Bertrand oligopoly. *Games Econ. Behav.* 33, 1–19.
 Benaïm, M., Weibull, J., 2003a. Deterministic approximation of stochastic evolution in games. *Econometrica* 71, 873–903.
 Benaïm, M., Weibull, J., 2003b. Deterministic approximation of stochastic evolution in games: a generalization. *Mimeo.*
 Binmore, K., Samuelson, L., 1997. Muddling through: noisy equilibrium selection. *J. Econ. Theory* 74, 235–265.
 Binmore, K., Samuelson, L., Vaughan, R., 1995. Musical chairs: modeling noisy evolution. *Games Econ. Behav.* 11, 1–35.
 Blume, L.E., 1993. The statistical mechanics of strategic interaction. *Games Econ. Behav.* 5, 387–424.
 Blume, L.E., 1995. The statistical mechanics of best-response strategy revision. *Games Econ. Behav.* 11, 111–145.
 Blume, L.E., 2003. How noise matters. *Games Econ. Behav.* 44, 251–271.
 de Bruijn, N.G., 1958. *Asymptotic Methods in Analysis*. North-Holland, Amsterdam.
 Cabrales, A., 2000. Stochastic replicator dynamics. *Int. Econ. Rev.* 41, 451–481.

Crawford, V.P., 1991. An “evolutionary” interpretation of Van Huyck, Battalio, and Beil’s experimental results on coordination. *Games Econ. Behav.* 3, 25–29.
 Ellison, G., Fudenberg, D., 1993. Rules of thumb for social learning. *J. Polit. Economy* 101, 612–643.
 Ellison, G., Fudenberg, D., 1995. Word of mouth communication and social learning. *Quart. J. Econ.* 110, 93–126.
 Ewens, W.J., 2004. *Mathematical Population Genetics*, second ed. Springer, New York.
 Ficici, S., Pollack, J., 2000. Effects of finite populations on evolutionary stable strategies. In: L. Darrell Whitley (Ed.), *Proceedings of the 2000 Genetic and Evolutionary Computation Conference*. Morgan-Kaufmann, Los Altos, CA.
 Fogel, G., Andrews, P., Fogel, D., 1998. On the instability of evolutionary stable strategies in small populations. *Ecol. Model.* 109, 283–294.
 Foster, D., Young, P., 1990. Stochastic evolutionary game dynamics. *Theor. Popul. Biol.* 38, 219–232.
 Foster, D., Young, P., 1997. A correction to the paper “Stochastic evolutionary game dynamics”. *Theor. Popul. Biol.* 51, 77–78.
 Freidlin, M., Wentzell, A., 1984. *Random Perturbations of Dynamical Systems*. Springer, New York.
 Fudenberg, D., Harris, C., 1992. Evolutionary dynamics with aggregate shocks. *J. Econ. Theory* 57, 420–441.
 Fudenberg, D., Levine, D.K., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
 Fudenberg, D., Imhof, L.A., 2005. Imitation processes with small mutations. *J. Econ. Theory*, in press, doi:10.1016/j.jet.2005.04.006.
 Hamilton, W., 1971. Selection of selfish and altruistic behavior in some extreme models. In: Eisenberg, J.F., Dillon, W.S. (Eds.), *Man and Beast: Comparative Social Behavior*. Smithsonian Institution, Washington, DC.
 Hofbauer, J., Sigmund, K., 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, England.
 Hofbauer, J., Sigmund, K., 2003. Evolutionary game dynamics. *Bull. Am. Math. Soc.* 40, 479–519.
 Hofbauer, J., Schuster, P., Sigmund, K., 1979. A note on evolutionarily stable strategies and game dynamics. *J. Theor. Biol.* 81, 609–612.
 Imhof, L.A., 2005. The long-run behavior of the stochastic replicator dynamics. *Ann. Appl. Probab.* 15, 1019–1045.
 Imhof, L.A., Nowak, M.A., 2006. Evolutionary game dynamics in a Wright-Fisher process. *J. Math. Biol.* 52, 667–681.
 Kandori, M., Rob, R., 1998. Bandwagon effects and long run technology choice. *Games Econ. Behav.* 22, 30–60.
 Kandori, M., Mailath, G., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 61, 29–56.
 Kifer, Y., 1990. A discrete-time version of the Wentzell–Freidlin theory. *Ann. Probab.* 18, 1676–1692.
 Komarova, N.L., Nowak, M.A., 2003. Language dynamics in finite populations. *J. Theor. Biol.* 221, 445–457.
 Leininger, W., 2005. Fending off one means fending off all: evolutionary stability in quasi-submodular aggregative games. *Econ. Theory*, in press, doi:10.1007/s00199-005-0027-9.
 Maruta, T., 2002. Binary games with state-dependent stochastic choice. *J. Econ. Theory* 103, 351–376.
 Moran, P.A.P., 1962. *The Statistical Processes of Evolutionary Theory*. Clarendon Press, Oxford.
 Morris, S., Rob, R., Shin, H., 1995. p-dominance and belief potential. *Econometrica* 63, 145–157.
 Nowak, M.A., Sigmund, K., 2004. Evolutionary dynamics of biological games. *Science* 303, 793–799.
 Nowak, M.A., Komarova, N.L., Niyogi, P., 2001. Evolution of universal grammar. *Science* 291, 114–118.
 Nowak, M.A., Sasaki, A., Taylor, C., Fudenberg, D., 2004. Emergence of cooperation and evolutionary stability in finite populations. *Nature* 428, 646–650.
 Possajennikov, A., 2003. Evolutionary foundations of aggregate-taking behavior. *Econ. Theory* 21, 921–928.
 Rhode, P., Stegeman, M., 1996. A comment on “Learning, mutation, and long run equilibrium in games”. *Econometrica* 64, 443–450.

- Rhode, P., Stegeman, M., 2001. Non-Nash equilibria of Darwinian dynamics with applications to duopoly. *Int. J. Ind. Organ.* 19, 415–453.
- Robson, A., Vega-Redondo, F., 1996. Efficient equilibrium selection in evolutionary games with random matching. *J. Econ. Theory* 70, 65–92.
- Sandholm, W.H., 1998. Simple and clever decision rules for a model of evolution. *Econ. Lett.* 61, 165–170.
- Schaffer, M.E., 1988. Evolutionarily stable strategies for a finite population and a variable contest size. *J. Theor. Biol.* 132, 469–478.
- Schaffer, M.E., 1989. Are profit-maximisers the best survivors? *J. Econ. Behav. Organ.* 12, 29–45.
- Taylor, P., Jonker, L., 1978. Evolutionary stable strategies and game dynamics. *Math. Biosci.* 40, 145–156.
- Taylor, C., Fudenberg, D., Sasaki, A., Nowak, M.A., 2004. Evolutionary game dynamics in finite populations. *Bull. Math. Biol.* 66, 1621–1644.
- Vega-Redondo, F., 1997. The evolution of Walrasian behavior. *Econometrica* 65, 375–384.
- Young, P., 1993. The evolution of conventions. *Econometrica* 61, 57–84.