

Evolutionary Dynamics with Aggregate Shocks*

D. FUDENBERG

*Massachusetts Institute of Technology, Cambridge,
Massachusetts 02139*

AND

C. HARRIS

Nuffield College, Oxford OX1 1NF, United Kingdom

Received March 22, 1991; revised March 3, 1992

We study the evolution of the continuous-time replicator dynamics when payoffs are subject to aggregate shocks that take the form of a Wiener process. In the absence of "mutation," the system need not have an ergodic distribution. With mutation, the system does have an ergodic distribution. In the limit as the mutation rate and the variance of the shocks converge to zero, this distribution concentrates on the risk-dominant equilibrium. This result is not, however, robust to changes in the underlying deterministic dynamics. *Journal of Economic Literature* Classification Numbers: C72, C73, C79. © 1992 Academic Press, Inc.

I. INTRODUCTION

Until recently, studies of evolutionary dynamics have used deterministic models, and in particular have focused on the adjustment process known as the replicator dynamics, in which the proportion of the population playing a particular pure strategy grows at a rate proportional to the difference between that strategy's current payoff and the average payoff obtained by the population as a whole. All Nash equilibria are fixed points of these dynamics, and all strict Nash equilibria are asymptotically stable fixed points.¹ Thus, the deterministic replicator model does not help in selecting between strict equilibria.

* We thank Dean Foster, Peyton Young, George Mailath, Raj Parekh, and the referee for helpful comments, Glenn Ellison for able research assistance, and the National Science Foundation (Grant SES-9008770) for financial support.

¹ A strategy profile is a strict Nash equilibrium if each player's strategy is a strict best response to the strategies of his opponents, i.e., all other strategies yield strictly lower payoffs.

Of course, results derived in deterministic models leave open the question of whether different conclusions obtain once stochastic influences are taken into account. This paper modifies the usual continuous-time replicator model by supposing that the payoff functions are subject to population-level or aggregate shocks that we model using Wiener processes. We exploit this specific stochastic system to make a number of points that we feel apply to the study of stochastic evolutionary models more generally. First, whether a stochastic evolutionary model has an ergodic distribution depends on apparently fine details about which it is difficult to have a very precise intuition. Second, in order to identify which stochastic models are most reasonable, it is better to introduce the stochastic elements at the level of the growth rates of individual populations than to introduce them directly at the level of the dynamics of population shares. Introducing the stochastic shocks at the more primitive level of the individual populations has the advantage of focusing attention on how the stochastic shocks are meant to be interpreted. In a model with a large population, such shocks are more naturally thought of as due to aggregate effects like the “weather” than as the results of individual-specific effects that might be expected to average out across the population. Another benefit of introducing the stochastic shocks at the level of the populations is that the implications of these shocks for the behavior of the population shares can be different than what intuition might suggest.

Our paper follows previous work by Foster and Young [4] and Kandori, Mailath, and Rob [8], which suggests that it may be possible to discriminate between strict Nash equilibria by considering evolutionary models with stochastic shocks. More precisely, both papers consider stochastic models that have ergodic distributions—asymptotic distributions over strategies that are independent of the initial position of the system—and consider the limits of the ergodic distributions over strategies as the variance of the noise term converges to 0. The papers then identify a class of 2×2 games in which the limit distribution is concentrated at one of the game’s two strict Nash equilibria, namely the equilibrium which is “risk dominant” in the sense of Harsanyi and Selten [6].

Our paper is closer to Foster and Young, who were, as far as we know, the first to consider a stochastic-differential-equation model of evolutionary dynamics. Our paper differs from theirs in the following ways. First, Foster and Young add the stochastic shocks directly to the deterministic replicator dynamics for population shares. As we will see, once stochastic effects are taken into account, it is not clear that the deterministic part of the evolution of the aggregate state should be the same as if stochastic effects are absent. Second, while Foster and Young only determine the limit of the ergodic distribution for the case where the variance of the stochastic process is independent of the population shares,

our analysis suggests that the most neutral assumption is that the variance depends on the population shares in a particular way. That is, if the variance is constant at the level of the growth rates of the individual populations, it will not be constant in the derived system for population shares. Third, Foster and Young suppose that the boundary of the strategy simplex is reflecting, which they justify as an approximation of the effects of immigration and mutation. (As we explain in the concluding section, we do not believe that their model can be interpreted in this way.) Fourth, our basic model does not always have an ergodic distribution, and the asymptotic behavior of the system can depend on its initial conditions. In particular, it is not possible to choose between two strict equilibria by letting the variance of the noise term go to zero, since the limiting distribution depends on the initial condition. (There is an ergodic distribution, however, when mutations are included in the model, and it converges to the risk-dominant equilibrium as the mutation rates and the variance of the noise term go to zero.)

Kandori, Mailath, and Rob differ in considering a discrete-time system with a finite population size, where each individual “mutates” from one strategy to the other strategy with a fixed probability. Not only does this model have an ergodic distribution, but the limit of this distribution as the probability of mutations shrinks is concentrated at the risk-dominant equilibrium for *any* deterministic dynamics that is a sign-preserving transformation of the replicator dynamics, i.e., so long as the population share playing a strategy tends to grow when that strategy does better than the population average. This striking conclusion is due to the fact that Kandori, Mailath, and Rob consider a system which can make discrete jumps: When the probability of mutation is low, the system spends most of its time at the two strict equilibria, with shifts from one equilibrium to the other occurring when enough individuals mutate simultaneously to shift the state of the system to the region where the other strategy is a best response. The limit of the ergodic distribution is determined by the relative probabilities of the shifts from one equilibrium to the other, which depends only on the relative sizes of the equilibria’s basins of attraction, and not on the speed of adjustment in each basin. In contrast, the limit of the ergodic distribution in our model (when one exists) does depend on the exact form of the deterministic dynamics. Intuitively, the likelihood that a Wiener process will be able to “swim upstream” k meters against a deterministic flow depends both on the distance k and on the strength of the flow, while the probability that a discrete-time system jumps k or more meters “over the flow” in a single period depends on k but not on the strength of the flow. This explains the differences in the generality of the models’ conclusions, and suggests that long-run behavior may depend on the precise form of the deterministic dynamics in any model with continuous sample paths.

II. THE MODEL

This section develops a stochastic version of the continuous-time replicator dynamics that we call our “preferred model.” We begin by reviewing the deterministic replicator dynamics. Consider a symmetric two-player game with n pure strategies s_1, \dots, s_n and payoff function $u(\cdot, \cdot)$. We suppose that there are n different populations or phenotypes of players, with every member of the i th population genetically coded to play s_i . Each population contains a continuum of infinitesimal individuals. The size of the i th population is r_i , $r = (r_1, \dots, r_n)$ is the vector of population sizes, and the total size of the population is $R = \sum_i r_i$.

The model evolves over time as follows. Individuals are repeatedly matched with a randomly chosen opponent to play the stage game. Let $u_i(r) = \sum_j u(s_i, s_j) r_j / R$ denote the expected payoff of an individual of type i when matched against a randomly chosen member of the population. Then we assume that

$$dr_i(t)/dt = r_i(t) \cdot u_i(r(t)), \quad (1)$$

where $r(t)$ is the vector of population sizes at time t . If we then define $\phi_i = r_i/R$ to be the proportion of type- i individuals, and note that u_i is homogeneous of degree 0, we can compute:

$$d\phi_i(t)/dt = \phi_i(t) \cdot \left[u_i(\phi(t)) - \sum_j \phi_j(t) u_j(\phi(t)) \right]. \quad (2)$$

Equation (2) is the deterministic replicator dynamics.

For future reference, we note that in the case of only two strategies, (2) simplifies to

$$d\phi_1(t)/dt = \phi_1(t) \phi_2(t) [u_1(\phi(t)) - u_2(\phi(t))]. \quad (3)$$

Before introducing the aggregate shocks, it may be helpful to review a few key features of the deterministic model. The standard interpretation of (1) is that individuals reproduce asexually, and that the number of their offspring is determined by their payoff in the game. Under this interpretation, (1) describes the *net* growth rate of the populations, i.e., the difference between the birth and death rates, so negative growth rates correspond to situations in which deaths outnumber births. Of course, actual populations are finite; (1) is intended to describe the limit behavior of a discrete-population process in which players from the various population are randomly matched to play the game. Intuitively, the randomness created by the random-matching process will “average out” by a law of large numbers, so that the continuum-of-players limit will be deterministic.² The

² Boylan [3] provides conditions under which the solution to (1) approximates the limit behavior of a finite-population, discrete-time, random-matching evolutionary model, where the time period shrinks to zero as the number of players tends to infinity.

same large-numbers intuition suggests that any stochastic effects that are i.i.d. across individuals will average out as well. In particular, i.i.d. "mutations" would not lead to a stochastic continuum-of-players limit, and any stochastic effects that persist in the limit must stem from shocks that are correlated across individuals.

Note finally that for any j the boundary point $s_j = 1$ is a steady state of (2) for *any* specification of the payoff functions.³ (This is because a population with zero members cannot grow by reproduction.) These boundary points will be steady states in our stochastic model as well, which is one way of explaining why we will not encounter difficulties in describing the behavior of the system near the boundary.

Now suppose that we wish to introduce the possibility of stochastic shocks. One way to do this is to add a stochastic disturbance directly to (2), which is the method used by Foster and Young. In our opinion, such an approach has two drawbacks. The first one is primarily technical: If, as in Foster and Young, one adds a Wiener process with constant variance to Eq. (2), the resulting solutions can have negative population shares, so that some additional changes are required for the model to be well defined; our method avoids these complications. Second, under the biological interpretation of the replicator dynamics, it is (1) that is fundamental, so that it is easier to interpret and evaluate changes to it than changes to the derived (2).

For these reasons, we prefer to add the disturbance to (1), and then derive the analog of (2). Let W be an n -dimensional Wiener process with unit variance and 0 covariance. Then suppose that r is an Ito process, with

$$dr_i(t) = r_i(t) \cdot [u_i(r(t)) dt + \sigma_i dW_i(t)]. \quad (1')$$

Note that the term $\sigma_i dW_i$ is added to the expected payoff $u_i(r)$ of an individual of type i and is *not* added directly to dr_i . For this reason, the population size r_i will remain positive for any realization of W_i . The interpretation of (1') is that the payoff to playing strategy i is subject to an *aggregate* shock, say due to the "weather." The equation incorporates the restrictions that the shock on each strategy i is independent of which other strategy it is matched with, and that the shocks on different strategies are independent.⁴

³ Strategies that are not Nash equilibria are unstable steady states.

⁴ To emphasize this point, consider for the moment a more general model in which the payoff to playing strategy i against strategy j is $u(s_i, s_j) + dV_{ij}(t)/dt$, where $dV(t)/dt$ is a matrix of possibly correlated payoff shocks. The model of this section is then the special case in which $dV_{ij}(t)/dt = dV_{ik}(t)/dt$ for all i, j and k , and $dV_{ij}(t)/dt$ and $dV_{km}(t)/dt$ are independent for all $i \neq k$. We believe that main points would still emerge without these restrictions, but we have not verified that this is the case.

Now we want to use (1') to compute the analog of (2). To do this, we write $\varphi_i(t) = f_i(r(t)) = r_i(t)/R(t)$, and apply Ito's lemma to obtain

$$d\varphi_i = \sum_j [\partial f_i(r)/\partial r_j] dr_j + \frac{1}{2} \sum_{j,k} [\partial^2 f_i(r)/\partial r_j \partial r_k] dr_j dr_k, \quad (2')$$

where $dr_j dr_k = r_j^2 \sigma_j^2 dt$ if $j = k$, and 0 otherwise.

When there are only two pure strategies, substituting in for the partial derivatives of f_i in (3) and simplifying yields

$$\begin{aligned} d\varphi_1 &= \varphi_1 \varphi_2 [(u_1(\varphi) - u_2(\varphi)) dt + (\sigma_2^2 \varphi_2 - \sigma_1^2 \varphi_1) dt + \sigma_1 dW_1 - \sigma_2 dW_2] \\ &= \varphi_1 \varphi_2 [(u_1(\varphi) - u_2(\varphi)) dt + (\sigma_2^2 \varphi_2 - \sigma_1^2 \varphi_1) dt + \sigma d\tilde{W}], \end{aligned} \quad (3')$$

where $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$, and $\tilde{W} = (W_1 - W_2)/\sigma$ is a standard Wiener process.

Let us discuss this equation before proceeding. The first term inside the square brackets, $(u_1(\varphi) - u_2(\varphi)) dt$, says that φ_1 tends to increase when strategy 1 does better than strategy 2; this term is familiar from the deterministic dynamics. The deterministic term $(\sigma_2^2 \varphi_2 - \sigma_1^2 \varphi_1) dt$ is not present in Eq. (2), and does not arise if the stochastic shocks are added directly to that equation. This term, which says roughly that if φ_1 gets large the stochastic shocks in the system will tend to make it smaller, arises for the following reason: For fixed r_2 , the function $g_1(r_1) = r_1/(r_1 + r_2)$, which maps $[0, \infty)$ to $[0, 1]$, is concave for $r_1 > r_2$, so by Jensen's inequality the expected value of $g_1(r_1)$ is less than g_1 applied to the expected value of r_1 . Likewise, g_1 is convex for $r_1 < r_2$.

The third term in square brackets represents the direct effect of the current shock. Finally, note that, as in (3), each term inside the square brackets in (3'), including the stochastic term $\sigma d\tilde{W}$, is multiplied by $\varphi_1 \varphi_2$. Thus, the boundary points are steady states of the stochastic process, and, if the process begins with positive shares for each strategy, the boundaries will never be reached in finite time.

III. ASYMPTOTIC BEHAVIOR

This section analyzes the asymptotic behavior of the system represented by (2'), with particular emphasis on when the model has an ergodic distribution, and when instead the system is eventually absorbed at a steady state.

a, a	b, c
c, b	d, d

FIGURE 1

Suppose that the payoff matrix is as in Fig. 1. We distinguish three cases:

(i) The case of a strictly dominant strategy. If $a > c$ and $b > d$ then strategy 1 strictly dominates strategy 2. Similarly, if $a < c$ and $b < d$ then strategy 2 strictly dominates strategy 1.

(ii) The coordination case. If $a > c$ and $b < d$ then the game has three Nash equilibria: two pure-strategy equilibria at $s_1 = 1$ and $s_1 = 0$, and a mixed-strategy equilibrium at $s_1^* = (d - b)/((d - b) + (a - c))$.

(iii) The case of a mixed-strategy equilibrium. If $a < c$ and $b > d$ then the unique symmetric equilibrium is a mixed-strategy equilibrium at $s_1^* = (d - b)/((d - b) + (a - c))$. (There are also two asymmetric pure-strategy equilibria, but these are not relevant in our setting.)

Aside from borderline cases, these cases are exhaustive.

Our main finding for these cases are as follows. In the case of a strictly dominant strategy, the deterministic replicator dynamics converge to this strategy. This is true of the stochastic dynamics (3') as well, provided that the variances σ_1^2 and σ_2^2 of the shocks are sufficiently small. The basic reason for this is that when σ_1^2 and σ_2^2 are small, the contribution that σ_1^2 and σ_2^2 make to the deterministic part of the dynamics is swamped by the contribution made by the payoffs, resulting in an unambiguous flow in the direction of the strictly dominant strategy. Since the variance of the stochastic part of the dynamics tends to zero as the strictly dominant strategy is approached, the stochastic dynamics converge to this strategy. (Because the deterministic part of the dynamics also tends to zero as the strictly dominant strategy is approached, the exact argument is somewhat more complex. The result depends, roughly speaking, on the ratio of the deterministic and stochastic parts.)

In the coordination case, the deterministic replicator dynamics converge to $s_1 = 1$ if the initial state $s_1(0) > s_1^*$, and to $s_1 = 0$ if the initial state $s_1(0) < s_1^*$. If $s_1(0) = s_1^*$, then the system remains at the unstable steady state s_1^* . The stochastic dynamics are similar: Provided that σ_1^2 and σ_2^2 are sufficiently small, the solution to (2') converges to one of the pure-strategy equilibria with probability one. The main difference with the deterministic dynamics is that each of the pure-strategy equilibria is reached with positive probability. This difference disappears, however, as σ_1^2 and σ_2^2 go to zero. For example, if $s_1(0) > s_1^*$ then the probability that our stochastic dynamics converge to $s_1 = 1$ goes to 1 as σ_1^2 and σ_2^2 go to zero.

The intuition for this case builds on that for the case of a strictly dominant strategy. In that case there is a single strict equilibrium that acts as a stable attractor for the stochastic dynamics. In the coordination case there are two strict equilibria, each of which acts as a stable attractor. Either attractor may "capture" the system; which one does will depend on

the stochastic shocks. Finally, as σ_1 and σ_2 go to zero, the shocks become relatively unimportant in determining which attractor captures the system. Our conclusions here are in contrast to those of Foster and Young, and Kandori, Mailath and Rob, who both obtain an ergodic distribution in this case.^{5,6}

In the case of a mixed-strategy equilibrium, the deterministic replicator dynamics converges to σ_1^* from all starting points. Our stochastic dynamics possess an ergodic distribution, to which the system settles down from all starting points. Moreover the ergodic distribution collapses to a point mass at σ_1^* as σ_1 and $\sigma_2 \rightarrow 0$. Our dynamics possess an ergodic distribution for two reasons. First, at the boundaries $\sigma_1 = 1$ and $\sigma_1 = 0$ the contribution of the payoffs to the deterministic part of the dynamics and the contribution of σ_1^2 and σ_2^2 both point towards the interior of $[0, 1]$. Hence the system cannot get stuck at a boundary. Second, the stochastic part of the dynamics has positive variance at all interior points, so the system cannot get stuck at an interior point either.

In order to establish these claims, we shall exploit results of Gihman and Skorohod [5] and Skorohod [9]. In order to state these results, let $d\sigma_1 = \alpha(\sigma_1) dt + \beta(\sigma_1) d\tilde{W}$ be a stochastic differential equation on the interval $(0, 1)$, and let $\sigma_1(0)$ be its initial position. Fix an arbitrary $z \in (0, 1)$, and introduce the quantities

$$I_1 = \int_0^{\sigma_1(0)} \exp \left[- \int_z^x [2\alpha(y)/\beta^2(y)] dy \right] dx,$$

$$I_2 = \int_{\sigma_1(0)}^1 \exp \left[- \int_z^x [2\alpha(y)/\beta^2(y)] dy \right] dx,$$

and the function

$$M(x) = \frac{2}{\beta^2(x)} \exp \left[\int_z^x [2\alpha(y)/\beta^2(y)] dy \right].^7$$

According to Theorem 16.1 of Gihman and Skorohod [5, p. 119], the main features of the asymptotic behavior of the system can be deduced from the

⁵ Foster and Young compute the ergodic distribution only for the case $b = c$.

⁶ Both Foster and Young, and Kandori, Mailath and Rob, consider only the coordination case.

⁷ Note that replacing z by z' in the formulae for I_1 , I_2 , and M has the effect of multiplying the exponential term by a constant which is independent of x . The arguments below depend on whether I_1 and I_2 are finite, on their ratio, and on the relative size of M at various x , and all of these properties are invariant when z is changed.

properties of I_1 and I_2 .⁸ Specifically, if I_1 is infinite and I_2 is finite, then the system converges to 1 almost surely; if I_1 is finite and I_2 is infinite, then the system converges to 0 almost surely; if I_1 and I_2 are both finite, then the system converges to 1 with probability $I_1/(I_1 + I_2)$ and to 0 with probability $I_2/(I_1 + I_2)$; and if I_1 and I_2 are both infinite, then the system oscillates forever, with

$$\text{Prob}\{\liminf_{t \rightarrow \infty} \phi_1(t) = 0\} = \text{Prob}\{\limsup_{t \rightarrow \infty} \phi_1(t) = 1\} = 1.$$

Also, according to Theorem 1.17 of Skorohod [9, p. 48], the description of the asymptotic behavior of the system can be refined in the case where I_1 and I_2 are both infinite by using the properties of M .⁹ Specifically, if I_1 and I_2 are both infinite and $\int_0^1 M(x) dx$ is finite, then the system has a unique ergodic distribution with density $M(x)/\int_0^1 M(w) dw$, and the distribution of $\phi_1(t)$ converges to this ergodic distribution as $t \rightarrow \infty$.

We can now establish

PROPOSITION 1. (i) *If $a - c > (\sigma_1^2 - \sigma_2^2)/2$ and $d - b < (\sigma_2^2 - \sigma_1^2)/2$ then $\phi_1(t) \rightarrow 1$ as $t \rightarrow \infty$ with probability 1.*

(ii) *If $a - c < (\sigma_1^2 - \sigma_2^2)/2$ and $d - b > (\sigma_2^2 - \sigma_1^2)/2$ then $\phi_1(t) \rightarrow 0$ as $t \rightarrow \infty$ with probability 1.*

(iii) *If $a - c > (\sigma_1^2 - \sigma_2^2)/2$ and $d - b > (\sigma_2^2 - \sigma_1^2)/2$ then $\phi_1(t) \rightarrow 1$ as $t \rightarrow \infty$ with probability $I_1/(I_1 + I_2)$ and $\phi_1(t) \rightarrow 0$ as $t \rightarrow \infty$ with complementary probability.*

⁸ The theorem quoted is actually developed for processes on the whole real line, whereas our process is confined to the unit interval. That a parallel theorem applies to our process can be seen in at least two ways. First it can be checked that Theorem 15.4 of Gihman and Skorohod [5, p. 110] applies to the unit interval, and that the proof of Theorem 16.1 therefore extends as well. Secondly, one can consider the process $\tilde{\phi}_1$ given by the formula $\tilde{\phi}_1 = \ln(\phi_1/(1 - \phi_1))$. By Ito's lemma, the process $\tilde{\phi}_1$ satisfies the stochastic differential equation

$$d\tilde{\phi}_1 = \left[\frac{\alpha(\phi_1)}{\phi_1(1 - \phi_1)} + \frac{\beta^2(\phi_1)(2\phi_1 - 1)}{2\sigma_1^2(1 - \phi_1)^2} \right] dt + \frac{\beta(\phi_1)}{\phi_1(1 - \phi_1)} d\tilde{W}.$$

But Theorem 16.1 of Gihman and Skorohod applies to $\tilde{\phi}_1$ in the case in which we are interested, namely Eq. (2'): $\alpha(\phi_1)$ and $\beta(\phi_1)$ both include the factor $\phi_1(1 - \phi_1)$, and so the standing regularity conditions of Gihman and Skorohod, namely linear boundedness and Lipschitz continuity, are satisfied. Hence the relevant features of the asymptotic behavior of $\tilde{\phi}_1$ can be determined by reference to integrals \tilde{I}_1 and \tilde{I}_2 . Finally, it turns out that there is a scalar $k > 0$ such that $I_1 = k\tilde{I}_1$ and $I_2 = k\tilde{I}_2$. (This can be verified by changing the variables of integration in \tilde{I}_1 and \tilde{I}_2 twice in the natural way.)

⁹ Once again, the theorem quoted is actually developed for processes on the whole real line. This difficulty can be overcome in exactly the same way as in the preceding footnote.

(iv) If $a - c < (\sigma_1^2 - \sigma_2^2)/2$ and $d - b < (\sigma_2^2 - \sigma_1^2)/2$ then

$$\text{Prob}\{\liminf_{t \rightarrow \infty} s_1(t) = 0\} = \text{Prob}\{\limsup_{t \rightarrow \infty} s_1(t) = 1\} = 1.$$

Moreover, the system possesses a unique ergodic distribution, to which the distribution of $s_1(t)$ converges as $t \rightarrow \infty$.

Part (i) of the proposition shows that, if strategy 1 is strictly dominant and if σ_1 and σ_2 are sufficiently small, then our stochastic dynamics converge with probability one to strategy 1. Part (ii) establishes the analogous result for the case when strategy 2 is strictly dominant. Part (iii) shows that, if the coordination case obtains and if σ_1 and σ_2 are sufficiently small, then our stochastic dynamics converge with probability one to one of the two strict equilibria, and that each equilibrium is reached with positive probability. Part (iv) shows that, in the case of a mixed-strategy equilibrium our stochastic dynamics possesses an ergodic distribution.

More generally, the proposition shows how the qualitative features of the asymptotics of our stochastic system, such as the existence or non-existence of an ergodic distribution, depend in relatively subtle ways on the interaction between the payoffs of the game and the size of the shocks to the system. It also shows that such qualitative features may be sensitive to the choice of stochastic system. For example, the systems of Foster and Young, and Kandori, Mailath and Rob both possess ergodic distributions irrespective of the exact configuration of payoffs.

Proof. Substituting the payoff matrix into (3') and eliminating $s_2 = 1 - s_1$ yields

$$\alpha(s_1) = s_1(1 - s_1)[\sigma_2^2(1 - s_1) - \sigma_1^2 s_1 + s_1(a - c) + (1 - s_1)(b - d)],$$

and

$$\beta(s_1) = s_1(1 - s_1) \sigma.$$

Substituting these values of α and β into the definitions of I_1 and I_2 yields

$$\begin{aligned} I_1 &= \int_0^{s_1(0)} \exp \left[- \int_z^x 2[\sigma_2^2(1 - y) - \sigma_1^2 y + (a - c) y \right. \\ &\quad \left. - (d - b)(1 - y)] / [\sigma^2 y(1 - y)] dy \right] dx \\ I_2 &= \int_{s_1(0)}^1 \exp \left[- \int_z^x 2[\sigma_2^2(1 - y) - \sigma_1^2 y + (a - c) y \right. \\ &\quad \left. - (d - b)(1 - y)] / [\sigma^2 y(1 - y)] dy \right] dx. \end{aligned}$$

To evaluate these integrals, note that the inner integral can be integrated exactly to obtain

$$\frac{-2(a-c-\sigma_1^2)}{\sigma^2} \ln\left(\frac{1-x}{1-z}\right) - \frac{2(d-b-\sigma_2^2)}{\sigma^2} \ln\left(\frac{x}{z}\right). \quad (4)$$

Hence

$$I_1 = \int_0^{\varphi_1(0)} (x/z)^{2[d-b-\sigma_2^2]/\sigma^2} [(1-x)/(1-z)]^{2[a-c-\sigma_1^2]/\sigma^2} dx. \quad (5)$$

$$I_2 = \int_{\varphi_1(0)}^1 (x/z)^{2[d-b-\sigma_2^2]/\sigma^2} [(1-x)/(1-z)]^{2[a-c-\sigma_1^2]/\sigma^2} dx. \quad (6)$$

But (5) is finite if and only if $2[d-b-\sigma_2^2]/\sigma^2 > -1$, i.e., if and only if $d-b > (\sigma_2^2 - \sigma_1^2)/2$. Similarly, (6) is finite if and only if $a-c > (\sigma_1^2 - \sigma_2^2)/2$. This establishes parts (i), (ii), and (iii), and the first part of (iv). (Note that we only need $a-c \leq (\sigma_1^2 - \sigma_2^2)/2$ and $d-b \leq (\sigma_2^2 - \sigma_1^2)/2$ for the first part of (iv).)

Proceeding as for I_1 and I_2 , one obtains

$$M(x) = \frac{2}{\sigma^2} \left(\frac{x}{z}\right)^{2[-\sigma_1^2 - d + b]/\sigma^2} \left(\frac{1-x}{1-z}\right)^{2[-\sigma_2^2 - a + c]/\sigma^2}.$$

So $\int_0^1 M(x) dx < \infty$ if and only if $2[-\sigma_1^2 - d + b]/\sigma^2 > -1$ and $2[-\sigma_2^2 - a + c]/\sigma^2 > -1$, i.e., if and only if $d-b < (\sigma_2^2 - \sigma_1^2)/2$ and $a-c < (\sigma_1^2 - \sigma_2^2)/2$. This establishes the second part of (iv). ■

Proposition 1 has established the basic asymptotics of our stochastic dynamics. However, in the coordination case it tells us only that the system converges to one of the two strict equilibria with a probability depending on the initial condition, and in the mixed-strategy case it tells us only that the system possesses an ergodic distribution. Proposition 2 which follows examines the behavior of the probability with which a given equilibrium is reached in the coordination case, and the behavior of the ergodic distribution in the mixed-strategy case, as σ_1 and σ_2 go to zero.

PROPOSITION 2. *Suppose that $a > c$ and $b < d$ (the coordination case). Then*

(i) *If $\varphi_1(0) > (d-b)/[d-b+a-c]$, the probability that the system converges to strategy 1 converges to 1 as $(\sigma_1, \sigma_2) \rightarrow (0, 0)$.*

(ii) *If $\varphi_1(0) < (d-b)/[d-b+a-c]$, the probability that the system converges to strategy 2 converges to 1 as $(\sigma_1, \sigma_2) \rightarrow (0, 0)$.*

Suppose, on the other hand, that $a < c$ and $b > d$ (the mixed-strategy case). Then:

(iii) *The ergodic distribution converges to the degenerate distribution that assigns unit mass to the unique mixed strategy equilibrium.*

These conclusions show that, as σ_1 and σ_2 become small, the behavior of the stochastic system comes to resemble that of the deterministic replicator dynamics more and more closely.

Proof. We begin with cases (i) and (ii). Note first that the integrands in (5) and (6) for I_1 and I_2 are the same, and are proportional to

$$f(x; \sigma_1, \sigma_2) = [x^{d-b-\sigma_2^2}(1-x)^{a-c-\sigma_1^2}]^{2/\sigma^2} = g(x; \sigma_1, \sigma_2)^{2/\sigma^2},$$

say. Second, because $a > c$ and $b < d$, I_1 and I_2 are both finite for σ_1, σ_2 sufficiently small. (Cf. the proof of Proposition (iii)). It follows that $f(\cdot; \sigma_1, \sigma_2)$ can be thought of as the unnormalized density of a random variable on $[0, 1]$. The ratios $I_1/(I_1 + I_2)$ and $I_2/(I_1 + I_2)$ are then the probability that this random variable lies in the intervals $[0, \phi_1(0)]$ and $[\phi_1(0), 1]$ respectively. Hence, to establish parts (i) and (ii), it will suffice to show that the distribution of this random variable converges to a unit mass at $\phi_1^* = (d-b)/(d-b+a-c)$. For this implies that if $\phi_1(0) < \phi_1^*$ then the probability that the random variable lies in $[0, \phi_1(0)]$ converges to zero, and if $\phi_1(0) > \phi_1^*$ then the probability that the random variable lies in $[0, \phi_1(0)]$ converges to one.

To this end, fix $\varepsilon > 0$. Since $g(\cdot; 0, 0)$ has a unique maximum at ϕ_1^* , $\zeta = g(\phi_1^*; 0, 0) - \max\{g(\phi_1^* - 2\varepsilon; 0, 0), g(\phi_1^* + 2\varepsilon; 0, 0)\} > 0$. Since $g(\cdot; \sigma_1, \sigma_2) \rightarrow g(\cdot; 0, 0)$ uniformly as $\sigma_1, \sigma_2 \rightarrow 0$, we may find $\delta > 0$ such that: $\delta \leq \varepsilon$; $g(x; \sigma_1, \sigma_2) \geq g(\phi_1^*; 0, 0) - \zeta/3$ for all $x \in [\phi_1^* - \delta, \phi_1^* + \delta]$ and all $|(\sigma_1, \sigma_2)| \leq \delta$; and $g(x; \sigma_1, \sigma_2) \leq g(\phi_1^*; 0, 0) - 2\zeta/3$ for all $x \in [0, \phi_1^* - 2\varepsilon] \cup [\phi_1^* + 2\varepsilon, 1]$. It follows that

$$\frac{\min_{x \in [\phi_1^* - \delta, \phi_1^* + \delta]} f(x; \sigma_1, \sigma_2)}{\max_{x \in [0, \phi_1^* - 2\varepsilon] \cup [\phi_1^* + 2\varepsilon, 1]} f(x; \sigma_1, \sigma_2)} \geq \left(\frac{g(\phi_1^*; 0, 0) - \zeta/3}{g(\phi_1^*; 0, 0) - 2\zeta/3} \right)^{2/\sigma^2},$$

for all $|(\sigma_1, \sigma_2)| \leq \delta$. Since the right-hand side of this inequality tends to infinity as $\sigma_1, \sigma_2 \rightarrow 0$, we conclude that all the mass of the distribution corresponding to $f(x; \sigma_1, \sigma_2)$ concentrates in the interval $[\phi_1^* - 2\varepsilon, \phi_1^* + 2\varepsilon]$ as $\sigma_1, \sigma_2 \rightarrow 0$. Since ε was arbitrary, this completes the demonstration of parts (i) and (ii).

A similar argument applied to the unnormalized density M establishes part (iii). ■

Note that Proposition 2 and the discussion immediately preceding it shows that the limiting behavior as the variance shrinks does not depend on the ratio of σ_1 to σ_2 .

IV. DETERMINISTIC "MUTATION" RATES

We call the system represented by (2') our "preferred" model because it adds aggregate stochastic shocks to the replicator dynamics in a way that we find natural. However, stochastic shocks are far from the only change to the replicator dynamics that one might want to consider. This section further alters the replicator dynamics to capture the possibility of mutation. We think of mutations as being independent between the many individuals in the population. This being the case, mutation is best modelled as a deterministic flow between strategies, as in Boylan [1].

More precisely, suppose that each individual of phenotype i has a finite flow probability of mutating into each phenotype $j \neq i$. Then, because there is an infinite number of individuals of phenotype i , there will be for each $j \neq i$ a deterministic flow from phenotype i into phenotype j equal to the flow probability for mutations of i into j times the total number of individuals in phenotype i . In the case of two phenotypes, if λ_i is the flow probability that an individual of phenotype i mutates into an individual of phenotype $3-i$, then the i th population satisfies the stochastic differential equation

$$dr_i = r_i(u_i(r) dt + \sigma_i dW_i) - \lambda_i r_i dt + \lambda_j r_j dt. \quad (7)$$

Note that this specification is appropriate for reproduction-independent mutations.¹⁰ With reproduction-dependent mutations it would be better to model mutation by a deterministic flow from phenotype i to phenotype j proportional to the *gross* rate of reproduction of phenotype i .¹¹ Such a model would be interesting, but complex, as it would require that separate processes be introduced to model births and deaths, both of which would have to be positive-valued. (Note that the above does *not* suggest that the mutation rate should apply to the *net* growth rate dr_i/dt .)

Returning to the system (7) we define $\varphi_1 = r_1/(r_1 + r_2)$ to be the population fraction playing strategy 1, and use Ito's lemma to compute

$$\begin{aligned} d\varphi_1 = & (-\varphi_2 \varphi_1^2 \sigma_1^2 + \varphi_1 \varphi_2^2 \sigma_2^2) dt \\ & + \left[\frac{\varphi_2}{R} (r_1 u_1 - \lambda_1 r_1 + \lambda_2 r_2) \right] - \left[\frac{\varphi_1}{R} (r_2 u_2 - \lambda_2 r_2 + \lambda_1 r_1) \right] dt \\ & + \varphi_1 \varphi_2 \sigma_1 dW_1 - \varphi_1 \varphi_2 \sigma_2 dW_2 \end{aligned}$$

¹⁰ Reproduction-independent mutations include those caused by mutagen damage (including chemical and spectral mutagens) and those caused by mobile genetic elements (including those associated with plasmids, transposons, viruses, and phages).

¹¹ Reproduction-independent mutations occur because replication of DNA is error-prone.

$$\begin{aligned}
&= s_1 s_2 [\sigma_2^2 s_2 - \sigma_1^2 s_1 + u_1(s) - u_2(s) + \lambda_2 - \lambda_1] dt \\
&\quad + [\lambda_2 s_2^2 - \lambda_1 s_1^2] dt + s_1 s_2 \sqrt{\sigma_1^2 + \sigma_2^2} d\tilde{W} \\
&= \alpha(s_1) dt + \beta(s_1) d\tilde{W},
\end{aligned} \tag{8}$$

say, where $d\tilde{W}$ is a standard Wiener process.

Inspection of (8) shows that the mutation terms λ_1 and λ_2 enter in two ways, once scaled by $s_1 s_2$ along with the payoffs, and once not so scaled, so that, e.g., the term $\lambda_2 s_2^2$ does not vanish at the boundary $s_1 = 1$, but rather points inwards. Thus the population will never converge to the boundaries. The deterministic part of the system may vanish in the interior of the interval, but here the variance is bounded away from zero. So we would not expect the population to converge to an interior point either, and hence we expect that (8) will have an ergodic distribution for any specification of the stage-game payoffs. Proposition 3 shows that this is the case.

PROPOSITION 3. *The process corresponding to (8) has an ergodic distribution for all $\sigma_1, \sigma_2 > 0$ and all $\lambda_1, \lambda_2 > 0$. Moreover the distribution of $s_1(t)$ converges to this ergodic distribution as $t \rightarrow \infty$.*

Proof. As above, we need to show that I_1 and I_2 are both infinite, and that $\int_0^1 M(x) dx$ is finite. To this end, note that

$$\begin{aligned}
\alpha(y)/\beta^2(y) &= [y(1-y)[\sigma_2^2(1-y) - \sigma_1^2 y + (a-c)y \\
&\quad - (d-b)(1-y) + \lambda_2 - \lambda_1] + \lambda_2(1-y)^2 - \lambda_1 y^2] / [\sigma_2^2 y^2(1-y)^2].
\end{aligned}$$

The largest contribution to $\alpha(y)/\beta^2(y)$ in the neighborhood of zero comes from the term $\lambda_2(1-y)^2$ in the numerator, and is of order $1/y^2$. It follows that the integral of $\alpha(y)/\beta^2(y)$ is of order $-1/x$, and therefore that I_1 is of order $\int_0^{s_1(0)} \exp(1/x) dx$. That is, I_1 is infinite. Similarly, I_2 is infinite.

Arguing in the same way we conclude that $M(x)$ is of order $\exp(-1/x)/x^2$ in the neighborhood of 0 and of order $\exp(-1/(1-x))/(1-x)^2$ in the neighborhood of 1. So $\int_0^1 M(x) dx$ is finite. ■

Proposition 1 characterized the long-run behavior of the stochastic dynamics (3'). Proposition 3 characterizes the long-run behavior of the stochastic dynamics with mutations (8). Comparison of the two propositions shows how a small change in the dynamics—in this case the introduction of arbitrarily small mutation rates λ_1 and λ_2 —can have a significant impact on long-run behavior. The contrast between the two propositions is greatest in the coordination case. In that case the long-run behavior of (3') depends on the initial state, whereas the long-run behavior of (8) does not.

One way of understanding why, in the coordination case, there is an ergodic distribution for any positive λ_1 and λ_2 but for $\lambda_1 = \lambda_2 = 0$ is to note that for a fixed small σ , the expected transit time from one basin of attraction to the other grows without bound as λ_1 and λ_2 go to zero. Thus the behavior of the system until a fixed finite time T is continuous at the limit $\lambda_1, \lambda_2 = 0$; it is the asymptotic behavior that changes discontinuously.

The reason that the expected transit times increase at the λ 's shrink is that this causes the stable steady states of the deterministic part of (8) to approach the boundary, and the process is likely to travel to the neighborhood of the "nearby" steady state before transiting to the basin of the other one. Thus for small λ_1 and λ_2 the process is likely to approach very near the "nearer" boundary before reaching the other basin, and near the boundary the replicator dynamics evolve very slowly. This raises the question posed by Ellison [3] in his critique of the work of Kandori *et al.*: For plausible parameter values, does the system converge to its ergodic distribution fast enough for that distribution to be relevant? We will return to this question after Proposition 4 below, which determines the limit of the ergodic distribution for the coordination case as the mutation rates and the variance of the shocks both converge to zero.

Propositions 1 and 3 show that obtaining an ergodic distribution in a modified version of the replicator dynamics requires both some noise at the aggregate level and a force that keeps the system away from the boundaries.¹² Aggregate shocks to payoffs do not satisfy the second requirement, as we have seen. Aggregate shocks with constant variance at the level of the population shares, a case considered by Foster and Young, provide a *probability* of moving in from the boundaries, but they also imply a positive probability of reaching the boundary in finite time, which poses difficulties we discuss in the concluding section.

It is interesting to note that in the finite-population model of Kandori, Mailath, and Rob, stochastic mutations satisfy both requirements: They add aggregate noise to the system, and keep it from being absorbed at the boundaries. This highlights an important difference between shocks to payoffs and shocks that operate directly on the strategies players use: The former will have little impact at states in which all players are currently using the same strategy, while the latter need not be expected to vanish at such states.¹³

Returning to the analysis of (8), one can compare the asymptotic behavior as σ_1 and σ_2 go to zero with that of system (3') when σ_1 and σ_2

¹² Actually it suffices that the system be kept away from one of the boundaries: If the other one is absorbing, the ergodic distribution will concentrate there.

¹³ Kandori *et al.* specify the mutation process at the individual level, and then derive the associated aggregate process, so that they do not encounter the boundary problems that complicate Foster and Young's model.

go to zero, as summarized in the remarks following Proposition 1 and in Proposition 2. Fortunately, the ergodic distribution of system (8) always has a limit as σ_1 and σ_2 go to zero, so the behavior of system (8) in this limit is well defined. The limiting distribution does, however, depend on λ_1 and λ_2 in general. The most concise way in which to make the comparison is therefore to compare the iterated limit in which first σ_1 and σ_2 and then λ_1 and λ_2 go to zero in system (8), with the limit in which σ_1 and σ_2 go to zero in system (3').

The iterated limit of the ergodic distribution of system (8) is easy to calculate in the cases of a strictly dominant strategy and of a mixed-strategy equilibrium: In the former case it is the strictly dominant strategy, and in the latter it is the mixed-strategy equilibrium. In the coordination case the general iterated limit defined above is not well defined. If, however, one requires that λ_1 and λ_2 go to zero in such a way that the ratio between them remains fixed, then the iterated limit does exist, and corresponds to the risk-dominant equilibrium.

From this perspective, then, the behavior of the system (3') as σ_1 and σ_2 to zero is very close to that of system (8) when σ_1 and σ_2 go to zero in the cases of a strictly dominant strategy and of a mixed-strategy equilibrium. In the coordination case, on the other hand, the two systems behave very differently. Letting σ_1 and σ_2 go to zero in system (3') leads to the selection of an equilibrium based on the initial condition, whereas letting σ_1 and σ_2 go to zero in system (8) (and then letting λ_1 and λ_2 go to zero, holding their ratio fixed) leads to the selection of the risk-dominant equilibrium independently of the initial condition. So we see once again how a small change in the dynamics can have a significant impact on the analysis.

We do not give precise statements or proofs of these results concerning the iterated limit of the ergodic distribution of system (8). The results for the cases of a strictly dominant strategy and of a mixed-strategy equilibrium are unsurprising; the results for the coordination case are a corollary of Proposition 4 below.

We pursue instead a slightly different line of thought. The overall effect of adding noise and mutations to the deterministic replicator dynamics is to ensure that the long-run behavior of the system can be described by an ergodic distribution independent of the initial condition. How does this distribution behave as the perturbation which gave rise to it is made small? The answer is contained in Proposition 4.

PROPOSITION 4. *Suppose that $a > c$ and $d > b$ (the coordination case). Fix $\rho_1, \rho_2 > 0$, and assume that $\lambda_1 = \rho_1 \lambda$ and $\lambda_2 = \rho_2 \lambda$, where $\lambda > 0$. Let $\Pi(\cdot | \sigma_1, \sigma_2, \lambda_1, \lambda_2)$ denote the ergodic distribution of σ_1 . Then*

$$\lim_{\sigma_1, \sigma_2, \lambda \rightarrow 0} \Pi(\cdot | \sigma_1, \sigma_2, \lambda_1, \lambda_2)$$

is

- (i) a point mass on $s_1 = 1$ if $a - c > d - b$; and
- (ii) a point mass on $s_1 = 0$ if $a - c < d - b$.

That is, the ergodic distribution collapses to a point mass on the risk-dominant equilibrium when we let σ_1 , σ_2 , λ_1 , and λ_2 go to zero, provided that the ratio of λ_1 to λ_2 is held fixed. (Actually, as the proof will make clear, all that we really require is that the ratio of λ_1 to λ_2 remains bounded away from zero and infinity.)

It is interesting to compare Proposition 4 with the result obtained by Kandori, Mailath, and Rob for the coordination case in their discrete-time, finite-population model. They found that the ergodic distribution collapses to the risk-dominant equilibrium as the probability of individual mutation goes to zero. Proposition 4 shows that the ergodic distribution collapses to the risk-dominant equilibrium as the probability of individual mutation and the size of aggregate shocks go to zero; moreover, this is true irrespective of the order in which limits are taken. Proposition 4 therefore supports their conclusion by obtaining a closely related result in a different model. The support lent by Proposition 4 is, however, qualified: The result of Kandori, Mailath, and Rob is robust to sign-preserving changes in the underlying deterministic replicator dynamics, but this is not true of Proposition 4. For example, it is easy to see that if we replace $u_1(s) - u_2(s)$ by $f(u_1(s) - u_2(s), s)$ in (8), where $f(\cdot, s)$ is a strictly increasing function such that $f(0, s) = 0$ for all s , then the equilibrium selected will depend on f as well as the payoffs a , b , c , and d .

Proof. We treat the case in which $a - c > d - b$. Recall that

$$M(x) = \frac{2}{\beta^2(x)} \exp \left[\int_z^x \frac{2\alpha(y)}{\beta^2(y)} dy \right],$$

with $z \in (0, 1)$ fixed. Noting that

$$\begin{aligned} \frac{1}{\beta^2(x)} &= \frac{1}{\sigma^2 x^2 (1-x)^2} \\ &= \frac{1}{\sigma^2} \exp \left[2 \int_z^x \left[\frac{1}{1-y} - \frac{1}{y} \right] dy - 2 \ln[z(1-z)] \right], \end{aligned}$$

it follows that $M(x)$ is proportional to

$$\exp \left[\int_z^x \frac{2\tilde{\alpha}(y)}{\sigma^2 \beta^2(y)} dy \right] = \exp \left[\frac{2}{\sigma^2} h(x) \right],$$

say, where

$$\begin{aligned}\tilde{\alpha}(y) &= \alpha(y) + \sigma^2[y^2(1-y) - y(1-y)^2] \\ &= y(1-y) \left[(a-c + \sigma_2^2)y - (d-b + \sigma_1^2)(1-y) + \frac{\lambda_2}{y} - \frac{\lambda_1}{1-y} \right],\end{aligned}$$

and $\tilde{\beta}(y) = y(1-y)$.

Evaluating the integral defining h , and choosing z appropriately, we may take it that

$$\begin{aligned}h(y) &= -(a-c + \sigma_2^2 + \lambda_2 - \lambda_1) \ln(1-y) \\ &\quad - (d-b + \sigma_1^2 + \lambda_1 - \lambda_2) \ln y - \lambda_2/y - \lambda_1/(1-y).\end{aligned}$$

It can also be shown that, for σ_1 , σ_2 , λ_1 , and λ_2 sufficiently small, $\tilde{\alpha}(y)/\tilde{\beta}^2(y)$ has precisely three zeros: $0 < Y_1 < Y_2 < Y_3 < 1$. So h is strictly increasing on $[0, Y_1]$ and on $[Y_2, Y_3]$, and strictly decreasing on $[Y_1, Y_2]$ and on $[Y_3, 1]$. Moreover Y_1 , Y_2 , and Y_3 are asymptotically equivalent to $\lambda_2/(d-b)$, $(d-b)/(d-b+a-c)$, and $1-\lambda_1/(a-c)$, respectively, when σ_1 , σ_2 , λ_1 and $\lambda_2 \rightarrow 0$. So $h(Y_3)$ and $h(Y_1)$ are asymptotically equivalent to $(a-c) \log(1/\lambda_1)$ and $(d-b) \log(1/\lambda_2)$ when σ_1 , σ_2 , λ_1 , and $\lambda_2 \rightarrow 0$. And these expressions in turn are asymptotically equivalent to $(a-c) \log(1/\lambda)$ and $(d-b) \log(1/\lambda)$ when σ_1 , σ_2 , and $\lambda \rightarrow 0$ (in other words, when we take account of the additional assumption that the ratio of λ_1 to λ_2 is fixed). Since $a-c > d-b$, we conclude that Y_3 is the global maximum of h . (Note that this maximum becomes more pronounced as λ gets smaller.)

Let X_1 be the unique point in (Y_2, Y_3) such that $h(X_1) = h(Y_1)$. Let $X_2 = 1 - \lambda_1 Y_1 / \lambda_2$. Since Y_1 is asymptotically equivalent to $\lambda_2/(d-b)$, X_2 is asymptotically equivalent to $1 - \lambda_1/(d-b)$. Hence we may take it that $X_2 \in (Y_2, Y_3)$. Next, from the formula for h we obtain

$$h(X_1) = h(Y_1) = (d-b) \ln \left(\frac{1}{Y_1} \right) - \frac{\lambda_2}{Y_1} + O(\varepsilon \log \varepsilon)$$

as $\varepsilon = \max\{\sigma_1^2, \sigma_2^2, \lambda_1, \lambda_2\} \rightarrow 0$. Also,

$$\begin{aligned}h(X_2) &= (a-c) \ln \left(\frac{1}{1-X_2} \right) - \frac{\lambda_1}{1-X_2} + O(\varepsilon \log \varepsilon) \\ &= (a-c) \ln \left(\frac{\lambda_2}{\lambda_1 Y_1} \right) - \frac{\lambda_2}{Y_1} + O(\varepsilon \log \varepsilon)\end{aligned}$$

as $\varepsilon \rightarrow 0$. Hence, bearing in mind that the ratio λ_2/λ_1 is fixed, $h(X_2) > h(X_1)$ for all ε sufficiently small, and we may take it that $X_2 > X_1$. It can be shown, further, that X_1 is asymptotically equivalent to $1 - \exp((d-b)/(a-c)) Y_1^{(d-b)/(a-c)}$. In particular, $X_1 \rightarrow 1$ as $\varepsilon \rightarrow 0$.

Now the *ratio* of the probability mass that the ergodic distribution assigns to the interval $[X_2, Y_3]$ to the probability mass that it assigns to the interval $[0, X_1]$ is at least

$$\frac{(Y_3 - X_2) \exp(2h(X_2)/\sigma^2)}{X_1 \exp(2h(X_1)/\sigma^2)}. \quad (9)$$

But $X_1 \rightarrow 1$ as $\varepsilon \rightarrow 0$, and $Y_3 - X_2$ behaves asymptotically like

$$\left(1 - \frac{d-b}{a-c}\right) \frac{\lambda_1}{\lambda_2} Y_1.$$

Hence (9) behaves asymptotically like

$$\begin{aligned} &\left(1 - \frac{d-b}{a-c}\right) \frac{\lambda_1}{\lambda_2} \exp \left[-\ln \left(\frac{1}{Y_1} \right) + \frac{2}{\sigma^2} \left[((a-c) - (d-b)) \right. \right. \\ &\quad \left. \left. \times \ln \left(\frac{1}{Y_1} \right) + (a-c) \ln \left(\frac{\lambda_2}{\lambda_1} \right) \right] \right]. \end{aligned}$$

Since λ_1/λ_2 is fixed, we conclude that (9) converges to infinity as $\sigma \rightarrow 0$; so the probability mass in the interval $[0, X_1]$ tends to zero. Since $X_1 \rightarrow 1$, we are finished. ■

As promised, we will now return to the question of the relevance of the ergodic distribution for plausible parameter values. Suppose that the payoff matrix is given by $a=2$, $b=c=0$, and $d=1$, so that strategy 1 is the risk-dominant equilibrium. Supposing that the system starts in the basin of strategy 2, say at $\sigma_1 = \frac{1}{4}$, how long will it take to reach the point $\sigma_1 = \frac{2}{3}$? If we let $v(x)$ be the expected waiting time to reach $\frac{2}{3}$ from initial condition x , a standard argument shows that v satisfies the differential equation

$$\beta^2(x) v''(x)/2 + \alpha(x) v'(x) = -1, \quad (10)$$

where α and β , respectively, are the deterministic and stochastic parts of the system (8). This second-order equation requires two boundary conditions. Clearly $v(\frac{2}{3}) = 0$; and the solution should also satisfy the condition $\lim_{x \rightarrow 0} v'(x) = -1/\lambda_2$. (Remember that at $\sigma_1 = 0$, (8) simplifies to $d\sigma_1 = \lambda_2 dt$.)

It can be verified that the formula

$$v(x) = \int_x^{2/3} \left(\int_0^y 2\beta^{-2}(s) \exp \left[- \int_s^y (2\alpha(t)/\beta^2(t) dt \right] ds \right) dy \quad (11)$$

satisfies (10) and both boundary conditions.

For given parameter values, (11) can be numerically integrated to obtain the expected transition time. But what parameter values are reasonable? Our specification of the deterministic replicator dynamics suggests that one unit of time is of the order of the average lifetime of the individuals in the population.¹⁴ Thus the mutation rates should correspond to the percentage of mutations per generation. Hofbauer and Sigmund [7] cite evidence that the frequency of mutations at the level of individual genes is of the order 10^{-6} . Of course, even for animals, the rate of mutations at the level of the organism can differ from that for the individual genes. Also, if the replicator dynamics is interpreted as a rough metaphor for a non-biological process, the appropriate mutation rate may be much higher. For this reason we computed the transition times for $\lambda = 10^{-4}$ and $\lambda = 10^{-2}$ as well. Finally, we choose the value $\sigma = 1$ for the variance of the payoff shocks. For all three parameter values, the expected transition times from $\frac{1}{4}$ to $\frac{2}{3}$ are of the order $10/\lambda$. Thus the transition times are “reasonably small” if $\lambda = 10^{-2}$, and perhaps for $\lambda = 10^{-4}$ as well.

V. RELATED WORK

The paper closest to ours is Foster and Young [4], who were, so far as we know, the first to model evolutionary dynamics with a stochastic differential equation.¹⁵ Our work differs from theirs in several ways. First, we begin with equations for population sizes, and then derive the equations for population shares using Ito’s lemma. This leads us to a different specification of the deterministic part of the system than they used. It also results in the variance term shrinking to zero at the boundaries. The processes Foster and Young considered did not have this property, and hence their analog of our (3’) can have solutions that reach the boundary in finite time. In particular, this is the case for their (3), which is the one example to which they apply their main theorem.

¹⁴ If a strategy’s payoff corresponds to its net reproduction rate, and, as in Boylan’s [2] construction, each individual reproduces once per unit of time, then the size of that strategy’s population should double in one time period. With our specification (1), the strategy’s population would actually increase by a factor of e ; this could be corrected by multiplying the right-hand side of (1) by the constant $\ln(2)$.

¹⁵ We thank George Mailath for making us aware of their work.

Foster and Young are aware that the state can reach the boundary in finite time. They respond to this problem by specifying the dynamics of their process only on the subset of the state space in which all population shares are at least $\Delta > 0$, and by arguing that immigration and mutation—appropriately modelled—will ensure that the process remains within this subset almost all of the time. They argue further that the resulting process (which they do not specify precisely) will be well approximated by a process that undergoes reflection at the boundary.¹⁶ This response raises a number of difficulties. First, the conclusions obtained from the analysis are likely to be sensitive to the specification of the dynamics for population shares less than Δ . Second, if the process undergoes reflection at the boundary, or is well approximated by such a process, then it is as though immigration of individuals of a given phenotype, or mutations resulting in individuals of a given phenotype, occur *only* when the population share of that phenotype is small. One might argue that the immigration process will be small compared to other effects in the interior of S_Δ , and hence can safely be ignored, but this conflicts with the fact that the rate of mutation and immigration must be infinite in order to generate instantaneous reflection in systems, like Foster and Young's examples, where the variance of the process is constant.¹⁷

Another closely related paper is Young and Foster [10], which reports simulations of a stochastic discrete-time evolutionary system with a finite population on a 3×3 two-player game with two pure-strategy equilibria, one of which is a strict equilibrium, and the other is weak. Each period, each individual plays every other individual, so that the matching process is deterministic. However, the payoff to each player given the strategies played is a random variable. In the simulations, as the population size grows and stochastic effects become correspondingly small, the system spends most of its time in the neighborhood of the strict equilibria. They suggest that this limit behavior can be proved by first applying their [4] results to a continuous-time, continuous-state-space system, and then arguing that these results can be extended to large-population limits of the kind they are considering. (Note that this will require a different limit argument than that of Boylan [2], who considers the case in which the

¹⁶ They state on page 223 that: "Mutation and immigration keep the process within the space S_Δ almost all of the time, and the boundary acts, to a good approximation, as if it were reflecting." The space S_Δ corresponds in our notation to $\{s | s_i \geq \Delta > 0 \text{ for all } i\}$. They have indicated in correspondence with us that they mean the boundary of S_Δ to be reflecting (as opposed to the boundary of the original strategy space), and that they have in mind instantaneous, as opposed to delayed, reflection.

¹⁷ A third difficulty with Foster and Young's use of a model with instantaneous reflection is that such models are not covered in the section of Freidlin and Wentzell that they cite. Foster and Young acknowledge this mistake and are preparing a note to correct it.

fraction of the population that is matched in a single time period goes to zero as the population grows.)

Kandori, Mailath, and Rob [8] consider a discrete-time, finite-population system playing a 2×2 game. As in Young and Foster, each period, each individual plays every other individual one at a time. At the end of the period, the current population shares $s_1(t)$ and $s_2(t)$ of the two strategies are updated to "intended" shares $\hat{s}_1(t+1)$ and $\hat{s}_2(t+1)$; the actual shares at $(t+1)$ can differ from the intended ones due to "mutations" from one strategy to the other. In this model, when the mutation rate is low the most probable way for the population to shift from one strict equilibrium to another is through a sufficiently large number of *simultaneous* mutations, as opposed to a succession of single mutations. Thus the ergodic distribution is determined by how many mutations are required to jump from the basin of attraction of one equilibrium to the basin of attraction of the other, i.e., on which equilibrium is risk-dominant. Because the transitions are made by "large" jumps, as opposed to a succession of small ones, the exact form of the process that maps the date- t state to the intended state at date $(t+1)$ is not important. All that matters is the size of the two basins. Since the exact form of the deterministic process does matter in models with continuous sample paths, their work combined with ours raises the question of which type of model is more appropriate for studying the evolution of large but finite populations.

REFERENCES

1. R. BOYLAN, Equilibria resistant to mutation, mimeo, California Institute of Technology, 1990.
2. R. BOYLAN, Continuous approximation of dynamical systems with randomly matched individuals, mimeo, Olin School of Business, Washington University, St. Louis, 1991.
3. G. ELLISON, Learning, local interaction, and coordination, mimeo, Department of Economics, Massachusetts Institute of Technology, 1991.
4. D. FOSTER AND P. YOUNG, Stochastic evolutionary game dynamics, *Theor. Pop. Biol.* **38** (1990), 219–232.
5. I. GIHMAN AND A. V. SKOROHOD, "Stochastic Differential Equations," Springer-Verlag, Berlin, 1972.
6. J. HARSANYI AND R. SELTEN, "A General Theory of Equilibrium Selection in Games," MIT Press, Cambridge, 1988.
7. J. HOFBAUER AND C. SIGMUND, "The Theory of Evolution and Dynamical Systems," London Mathematical Society, London, 1988.
8. M. KANDORI, G. MAILATH, AND R. ROB, Learning, mutation, and long-run equilibrium in games, mimeo, 1991.
9. A. V. SKOROHOD, "Asymptotic Methods in the Theory of Stochastic Differential Equations," American Mathematical Society, Providence, RI, 1989. [Translated by H. H. McFaden]
10. P. YOUNG AND D. FOSTER, Cooperation in the short and in the long run, *Games Econ. Behav.* **3** (1991), 145–156.