GAMES AND ECONOMIC BEHAVIOR **11**, 1–35 (1995)

Musical Chairs: Modeling Noisy Evolution*

KENNETH G. BINMORE

Department of Economics, University College London, Gower Street,
London WC1E 6BT, England

LARRY SAMUELSON

Department of Economics, University of Wisconsin, 1180 Observatory Drive,
Madison, Wisconsin 53706

AND

RICHARD VAUGHAN

Department of Economics, University College London, Gower Street,
London WC1E 6BT, England

Received October 1, 1993

We examine an evolutionary process based on an explicit model of choice in which agents occasionally make mistakes in choosing their strategies. If the population size is sufficiently large, then the deterministic replicator dynamics provides a good approximation of the behavior of the system over finite time periods. The limiting behavior of the process is captured by combinatorial techniques introduced by Freidlin and Wentzell and popularized by Young and Kandori, Mailath, and Rob. We find conditions under which the limiting distribution selects the risk-dominant equilibrium in a 2×2 game and conditions under which the payoff-dominant equilibrium is selected. *Journal of Economic Literature* Classification Numbers: C70, C72.

This paper is part of the special issue devoted to the Nobel Symposium on Game Theory published in *Games and Economic Behavior* **8,** No. 1 (Jan. 1995).

* Financial support from National Science Foundation Grants SES-9122176 and SBR-9320678, the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303, and British Economic and Social Research Council Grant L122-25-1003 is gratefully acknowledged. This paper was prepared for the Nobel Symposium at Björkborn, Sweden, June 18–20, 1993. Part of this work was done while the authors were visiting the University of Bonn and the Institute for Advanced Studies at Hebrew University; we are grateful to both for their hospitality and support. We are thankful for the helpful comments of Larry Blume, Josef Hofbauer, Yuri Kaniovski, Robert Seymour, three referees, and an associate editor.

0899-8256/95 $12.00

	dove	*hawk*
dove	5 5	4 0
hawk	0 4	2 2

a

	dove	*hawk*
dove	a(b-5) a(b-5)	a(b-4) ab
hawk	ab a(b-4)	a(b-2) a(b-2)

b

	dove	*hawk*
dove	1 1	2 0
hawk	0 2	-1 -1

c

FIG. 1. Some canonical games: (a) Stag-Hunt Game, (b) Death Game, and (c) Chicken Game.

1. INTRODUCTION

The Stag-Hunt Game of Fig. 1a is a canonical example in the literature on equilibrium selection in games. The game has two Nash equilibria in pure strategies, (*dove*, *dove*) and (*hawk*, *hawk*). It also has a mixed Nash equilibrium in which both players use *hawk* with probability $\frac{1}{3}$. Harsanyi and Selten (1988) call (*dove*, *dove*) the payoff-dominant equilibrium because it is Pareto-superior to the other Nash equilibria of the game. They call (*hawk*, *hawk*) the risk-dominant equilibrium. In a 2 × 2 symmetric game with two symmetric, strict Nash equilibria, the risk-dominant equilibrium is characterized by having the larger basin of attraction under the best-reply dynamics, as illustrated in Fig. 2b.

Which of the Nash equilibria in the Stag-Hunt Game should be selected? The mixed-strategy equilibrium is commonly dismissed, but the choice between the two pure strategy equilibria has provoked much debate. This debate has largely been conducted from the standpoint of evolutionary game theory, as the literature on refinements of Nash equilibrium is usually silent when a choice has to be made between two strict Nash equilibria.[1]

---

[1] Amongst nonevolutionary theories that do make a choice, Harsanyi and Selten's (1988) equilibrium selection theory proposes the payoff-dominant equilibrium as the rational choice in the Stag-Hunt Game. (This Nobel Symposium on Game Theory contains papers in which both Harsanyi and Selten separately offer new and different versions of their theory.) Anderlini's (1990) theory of cheap-talking Turing machines also selects the payoff-dominant equilibrium. Carlsson and Van Damme (1993), in contrast, offer the risk-dominant equilibrium as the rational choice.

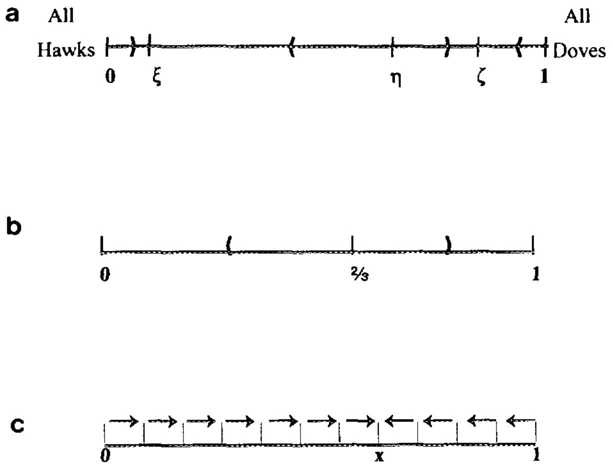

FIG. 2. Population states: (a) replicator dynamics with mutation, (b) best reply dynamics, and (c) the significant x-tree.

Biological theories of evolution in games often take replicator dynamics as their basic tool of analysis. But such an approach does not readily distinguish between the two equilibria of the Stag-Hunt Game. Both (*dove, dove*) and (*hawk, hawk*) are asymptotically stable[2] with respect to the replicator dynamics. The same conclusion also follows when various other dynamic adjustment processes are examined. However, Young (1993) and Kandori *et al.* recently have studied the effect of introducing small random shocks (or mutations) into the adjustment dynamics. They considered versions of a myopic best-reply dynamic, but their methodology is more generally applicable. If the amount of noise that perturbs the adjustment process is sufficiently small, then the system is highly likely to begin by finding its way to an equilibrium of the game and remaining in the vicinity of that equilibrium for an appreciable time. However, the small random shocks built into the process make it inevitable that the system will eventually bounce out of the basin of attraction of whichever equilibrium has captured it and find its way to another. Such transitions give rise to a probability distribution that describes the fraction of the time that the system spends at the various equilibria. In both Young (1993) and Kandori *et al.* (1993) we find that, as the noise level becomes negligible, this limiting distribution places all its probability on the risk-dominant equilibrium.

---

[2] We use the terminology of Hofbauer and Sigmund (1988, p. 51). A *stable point* is a rest point of the dynamics with the property that orbits that start nearby stay nearby. An *asymptotically stable point* is a stable point with the additional property that orbits originating in some neighborhood of the point converge to the point in question.

Two concerns arise in connection with this result. First, might not the noise generated by mutations of arbitrarily small probability be overwhelmed by other sources of randomness currently excluded from the model? For example, Foster and Young (1990), Fudenberg and Harris (1992), and Cabrales (1993) use stochastic differential equations to examine models in which either payoffs or the process by which relative payoffs are translated into strategy adjustments are subject to continual perturbation. These perturbations may be small enough to be negligible in a model where mutation probabilities are fixed, but can less obviously be ignored in a model driven by arbitrarily small mutation probabilities. Second, how long must one wait to reach the stationary distribution? With random shocks of very low probability, the expected waiting time to get from the payoff-dominant equilibrium to the risk-dominant equilibrium may be very long indeed.[3]

In this paper, we study a "musical chairs" model of evolution. The effective source of noise in this model is located in the selection process itself. Even when mutations are absent, the selection process may still make "mistakes" in that agents need not always adjust their strategies in the direction of the current best reply. In Binmore and Samuelson (1994) it is argued that such models have a potentially wide application in economics. However, in the current paper we find it convenient to work with a biological interpretation of the model.

We have four aims in studying the musical chairs model. The first is to expose the relationship between (1) an explicit selection model; (2) the classical replicator dynamics; and (3) the combinatorial methodology of Freidlin and Wentzell (1984) upon which the equilibrium selection theories of Young (1993) and Kandori *et al.* (1993) are based. Making the necessary linkages turns out to depend primarily on the time span over which one chooses to study the behavior of the selection model. In particular, if the population is sufficiently large, then the replicator dynamics provide a good description of the behavior of the population along any finite interval of time. However, the replicator dynamics do not provide a good description of the model's limiting behavior, which is precisely captured by the methods of Freidlin and Wentzell.

The second aim is to contrast the equilibrium selection results of the musical chairs model with the equilibrium selection theories of Young (1993) and Kandori *et al.* (1993). Our attempt to incorporate more realistic sources of noise in the model yields shorter waiting times but generates

---

[3] To jump out of the basin of attraction of the payoff-dominant equilibrium in the Stag-Hunt game of Fig. 1, given the Kandori *et al.* model with a population size of 1,000, requires 333 simultaneous mutations. Ellison's (1992) work shows that the long expected waiting times in these models can be abbreviated if agents are matched only with their "neighbors" rather than with all other agents.

different equilibrium selection results. In particular, the musical chairs model sometimes chooses the payoff-dominant rather than the risk-dominant equilibrium.

The third aim is to emphasize that the ultralong-run outcomes of an evolutionary process depend on *fitnesses* which in turn depend both on the rewards available in the game *and* on the nature of the selection process. In consequence, when one game is obtained from another by making a strictly increasing, affine transformation of its rewards, the two games need not be strategically equivalent, since the derived fitnesses in the two games may differ.

The fourth aim is to comment on the use of continuous-time techniques, often involving stochastic differential equations, in evolutionary analyses. We find that these techniques do not always provide good approximations of the stationary distribution.

Our method in pursuing these aims is to focus on an explicit model of selection. Given that any such model can be criticized as arbitrary and unrealistic, why bother? Why not work directly with the replicator dynamics, or with some other abstractly formulated deterministic or stochastic equations of motion?

In response to such questions, we note that at least four variants of the discrete replicator dynamics have been studied (Cabrales and Sobel, 1992; Taylor and Jonker, 1978; Maynard Smith, 1982; Dekel and Scotchmer, 1992; van Damme, 1991; and Hofbauer and Sigmund, 1988). Limits of these can be taken to give at least two variants of a continuous-time replicator dynamics. Hofbauer and Sigmund (1988, p. 133), Samuelson and Zhang (1992), Dekel and Scotchmer (1992), Cabrales and Sobel (1992), Boylan (1992), Bomze and Eichberger (1992), and Seymour (1993) all contain examples showing how such different dynamics can lead to different behavior. Things become more complicated if one adds perturbations to the replicator dynamics, with Foster and Young (1990), Fudenberg and Harris (1992), and Cabrales (1993) advancing various formulations with different properties.

Like Boylan (1992, 1995), we see no refuge from such arrays of alternatives other than a close examination of explicit models of selection. The ultimate aim is to classify different types of models in terms of their equilibrium selection properties. The models studied in making a start on this program are admittedly arbitrary and unrealistic to some degree, but they must give pause to critics who argue for a more "general" theory of equilibrium selection, as different types of model demonstrably select different equilibria.

The following section presents the model. Section 3 examines the sample paths of the musical chairs model and their relationship to replicator dynamics. Section 4 studies the asymptotic distribution in the musical chairs model and its relationship to the techniques of Freidlin and Wentzell (1984). Section 5 examines equilibrium selection. Section 6 discusses the implications of interpreting payoffs as fitnesses. Section 7 seeks to relate our analysis to work based on stochastic differential equations. Section 8 summarizes our conclusions.

2. The Musical Chairs Model

Births. Time is indexed by t. It is divided into periods of length τ, the first of which begins at time 0. When time t is reached, $[t/\tau]$ periods will have been completed.[4]

At the beginning of each period there will be a fixed number N of rabbits who have survived from the previous period. Rabbits are genetically programmed either to play "hawk" or "dove." In every period, each rabbit gives birth with probability $\beta\tau$ to an instantly mature child who is programmed just like her mother unless a mutation takes place. Without loss of generality, we take the birth rate to be $\beta = 1$.

How long is a period? This depends upon what kind of process we are trying to model. One possibility is to let $\tau = 1$, so that in every period every rabbit gives birth. We can then think of a period as a breeding season. We might then also assume that parents die after giving birth, though this is not necessary. From such a model, we could derive a difference equation describing how the proportions of hawks and doves change from period to period.

In contrast, we examine a model in which generations overlap and births are rare, isolated events. Most periods pass without births occurring, and the probability of multiple births in a period is negligible. To this end, we study the limiting case as τ goes to zero. We will also be interested in examining the limit when the population size N gets large and the mutation probability λ gets small—primarily because these limits sharpen the results and make convenient approximations available. However, we shall always take the limit $\tau \rightarrow 0$ first.

An alternative approach would be to work "directly" in continuous time. This would provide immediate access to some elegant mathematical techniques of the general type discussed in Section 7. However, we avoid this shortcut because we have doubts about the correct interpretation of continuous-time models that have not been obtained explicitly as the limits of discrete-time models.

Deaths. In addition to births being rare, competition for survival in our world is fierce. It is assumed that, after a birth, there are too many rabbits

---

[4] Where $[x]$ is the largest integer smaller than x.

for the environment to support. One may perhaps imagine that there are only N rabbit holes into which rabbits can retreat when the fox comes by—just as there are only N chairs but $N + 1$ children in the game Musical Chairs when the music stops.

Whenever a birth occurs and there are then $N + 1$ rabbits, two randomly chosen rabbits will be forced to compete for the same hole. As a result of this competition either both rabbits may escape or one may perish. If both escape the process is timelessly repeated, with a new, randomly chosen pair of rabbits each time, until a rabbit dies.[5]

The competition between rabbits Mopsy and Flopsy for the same rabbit hole will be described by the Death Game of Fig. 1b. When a cell in the payoff table contains the pair (m, f), then m is the probability that Mopsy fails to survive the encounter, f is the probability that Flopsy fails to survive, and $1 - m - f$ is the probability that both survive. We might accordingly describe the payoffs in the Death Game as *antifitnesses,* since they describe probabilities of death. The payoff pair (m, f) in the Death Game is obtained from the corresponding payoff pair (M, F) in the Stag-Hunt Game of Fig. 1a by making the strictly *decreasing,* affine transformations $m = a(b - M)$ and $f = a(b - F)$, where the constants $a > 0$ and $b > 5$ are chosen so that $m > 0$, $f > 0$, and $m + f < 1$.

We think of the payoffs in the Stag-Hunt Game as being physical rewards. In a biological system, they might be measured in calories. In an economic model they might be monetary rewards. The payoffs in the Death Game are the antifitnesses or death probabilities that describe the implications of the physical rewards for survival. The nature of the relationship between rewards and antifitnesses will depend upon the details of the physical selection process by which some agents are chosen to die while others survive.

As Binmore and Samuelson (1994) point out, there is obviously no particular reason why antifitnesses should be affine functions of the physical rewards. If antifitnesses *are* affine functions of rewards, then traditionally minded game theorists would regard the Death Game, with players *minimizing* antifitness, as strategically equivalent to the Stag-Hunt Game in which players seek to *maximize* their expected payoffs in the usual way. For our purposes, the current affine relationship has the virtue of making the risk-dominant equilibrium in the Death Game the same as the risk-dominant equilibrium in the Stag-Hunt Game. Because we always work with the Death Game, our analysis would continue to hold for more general transformations between rewards and antifitnesses, though interpreting risk dominance then becomes more complex.

---

[5] In the negligible event that more than two rabbits are born in a period, such "death games" continue until the population size has been reduced to N. The death process need not be timelessly repeated as long as, in the continuous-time limit, it occurs arbitrarily rapidly relative to the birth rate.

Mutation. Selection is a noisy process in the musical chairs model. Given the current population mix, it may be optimal to be a dove. But it does not follow that the number of doves in the population will necessarily increase, because it remains possible that a dove competing for survival may simply be unlucky. If the basin of attraction of an equilibrium E is taken to be the set of population states from which it is possible to move to E without mutations, then the basin of attraction of the payoff-dominant equilibrium (in which all rabbits are doves) includes all population states in $(0, 1]$, where a state identifies the proportion of rabbits who are doves. Similarly, the basin of attraction of the risk-dominant equilibrium (in which all rabbits are hawks) includes all population states in $[0, 1)$.

To ensure that the system does not get trapped forever in one of the two external states in which all agents play the same strategy, mutations are necessary. When a rabbit is born, it is assumed that she is identical to her mother with probability $1 - \lambda$. But with probability λ, she is a mutant and hence comes programmed with the strategy that her mother does not use. Unlike the models of Young (1993) and Kandori *et al.* (1993), the musical chairs model requires only a single mutation to move the system out of one equilibrium into the basin of attraction of the other. When mutations are rare, expected waiting times are therefore much shorter than for a transition that requires large numbers of simultaneous mutations.

How Long is Long? How do we analyze the model? Two approaches are examined in this paper.

Consider first the traditional story that is told when the classical replicator dynamics are introduced. By first dispensing with mutations ($\lambda \rightarrow 0$) and then letting the population size grow large ($N \rightarrow \infty$), a model is obtained in which large numbers of simultaneous births occur at discrete intervals. This model is described by a deterministic difference equation known as the discrete replicator dynamic. This difference equation is commonly approximated by a differential equation ($\tau \rightarrow 0$). The limiting behavior of the differential equation is then studied ($t \rightarrow \infty$). We can summarize this derivation by representing the parameters of the model as having been taken to their limits in the following order: first $\lambda \rightarrow 0$, next $N \rightarrow \infty$, next $\tau \rightarrow 0$, then $t \rightarrow \infty$.[6]

In contrast, we are concerned in this paper with an overlapping generations model with isolated births. As we have noted, this dictates that the limit $\tau \rightarrow 0$ be taken before $N \rightarrow \infty$. We are not convinced that the limit $\lambda \rightarrow 0$ should be taken at all, and in some cases it makes no difference when this limit is taken. However, we argue in Section 5 that, where it

---

[6] The replicator dynamics are generally not rigorously derived from precise limiting operations and we view these limits as providing a convenient summary of the common informal derivation.

makes a difference, the limit $N \rightarrow \infty$ should be taken before $\lambda \rightarrow 0$. Hence, the order of the limits $\tau \rightarrow 0$, $N \rightarrow \infty$, and $\lambda \rightarrow 0$ is set.

We are then left with one key choice. Should we first let the parameters τ, N, and λ go to their limits and then calculate the limiting outcome of the model, or should we first calculate the limiting outcome of the model and then let τ, N, and λ go to their limits? In considering this question, we let "$t \rightarrow \infty$" be shorthand for "calculating the limiting outcome of the model." Note that using $t \rightarrow \infty$ in this way is at this point simply a convention. One of the topics we will investigate is whether the calculations to be performed in finding the limiting outcome of the model can be interpreted as taking the limit $t \rightarrow \infty$ in some precise way.

In Section 3, we first let the parameters τ, N, and λ go to their limits and then examine the outcome of the model. We find that the important relationship here is that τ and N approach their limits before examining the model, and that the treatment of λ does not matter. The limiting operations can thus be described as one of the following two cases.[7]

Case 1. First $\tau \rightarrow 0$, next $N \rightarrow \infty$, next $\lambda \rightarrow 0$, then $t \rightarrow \infty$.

Case 2. First $\tau \rightarrow 0$, next $N \rightarrow \infty$, next $t \rightarrow \infty$, then $\lambda \rightarrow 0$.

This analysis leads us to a version of the classical replicator dynamics (with an extra term added to take account of the existence of mutations). Studies of the replicator dynamics usually examine the asymptotics of the system with the simplifying assumption that $\lambda = 0$. This corresponds to our Case 1, in which $\lambda \rightarrow 0$ first and then $t \rightarrow \infty$. The more realistic Case 2, in which mutations are eliminated after the asymptotics of the system have been studied, leads to the same conclusions. Section 3 shows that the replicator dynamics we derive provide a good approximation of the behavior of our model over arbitrarily long finite periods. The replicator dynamics do not necessarily provide a good approximation over infinite periods of time, so that our usage of $t \rightarrow \infty$ must be interpreted with care here, even though most analyses based on the replicator dynamics simply assert that they are examining the behavior of the underlying system in the limit as $t \rightarrow \infty$ and literally take such a limit in the replicator model.

Section 4 turns to the case in which the outcome of the model is examined before the parameters τ, N, and λ are taken to their limits. We find that it is crucial here that the model be examined before $N \rightarrow \infty$, and hence we have the following two cases:

Case 3. First $\tau \rightarrow 0$, next $t \rightarrow \infty$, next $N \rightarrow \infty$, then $\lambda \rightarrow 0$.

Case 4. First $t \rightarrow \infty$, next $\tau \rightarrow 0$, next $N \rightarrow \infty$, then $\lambda \rightarrow 0$.

---

[7] The precise operation here is that $(\tau, N) \rightarrow (0, \infty)$ at rates such that $N^2\tau \rightarrow 0$.

After showing that the two cases give identical results, we concentrate on the more convenient Case 4. Here, the first limit is $t \rightarrow \infty$, meaning that the first step is to derive the stationary distribution (also called the limiting distribution or asymptotic distribution) of the Markov process for fixed values of the parameters. In this analysis, $t \rightarrow \infty$ is a literal description. The study of this case leads to limiting predictions for the musical chairs model that could be obtained using the techniques of Freidlin and Wentzell (1984) but which we prefer to derive using standard techniques of birth–death processes developed, for example, in Gardiner (1985). The remaining limits are comparative static exercises that allow the derivation of particularly sharp approximations of the limiting distribution. Sections 5 and 6 discuss some of the implications of these results.

Section 7 considers an alternative approach to characterizing the stationary distribution of the Markov process. It is shown that, if a commonly used approximation is introduced after taking the limit $\tau \rightarrow 0$ when following the program specified by Case 3, we are led to equations familiar in diffusion theory, which economics usually study with the aid of the theory of stochastic differential equations. We find that the approximation sometimes yields inaccurate predictions concerning the stationary distribution.

Which of these modes of analysis is appropriate? This depends upon the period of time over which we are interested in the behavior of the model. Following Binmore and Samuelson (1993), we find it helpful to speak of the *short run, medium run, long run,* and *ultralong run*. By the short run, we refer to the initial conditions that prevail when one begins one's observation or analysis. By the ultralong run, we mean a period of time long enough for the asymptotic distribution to be a good description of the behavior of the system. The long run refers to the time span needed for the system to reach the vicinity of the first equilibrium in whose neighborhood it will linger for some time. We speak of the medium run as the time intermediate between the short run and the long run, during which the adjustment to equilibrium is occurring. The analysis of Cases 1 and 2 thus applies if one is interested in the medium or long run, while the analysis of Cases 3 and 4 is relevant to the ultralong run. Medium- and long-run behavior is often referred to as "sample path" behavior, meaning behavior over time spans too short for the asymptotic distribution to be a good description.

Why do we use both medium run and long run to refer to sample path behavior, hijacking the latter from its more conventional use in describing the asymptotic distribution, and then invent the new term "ultralong run" to refer to the asymptotic distribution? There are two reasons.

First, consider the behavior one might observe in an experimental dynamic system. Initial behavior in experiments is driven to a large extent by rules and norms of behavior that are triggered by the framing of the problem. This is short-run behavior. Given adequate incentives and sufficient time, experimental subjects begin to learn. It is the importance of this learning phase that prompts us to give it a name. Following Roth and Erev (1995), we refer to this as medium-run behavior.[8] This learning process may converge on an equilibrium, with this equilibrium behavior persisting apparently indefinitely. This is the long run.

Second, suppose the system has settled on a long-run equilibrium. The occurrence of extremely unlikely realizations of the random components of the learning process may jolt the system from one equilibrium to the basin of attraction of another. Given sufficient time, these jumps will occur often enough to produce a well-defined asymptotic probability distribution over states. But, as Ellison in (1992) has noted, the "sufficient time" required here can be extraordinarily long, much longer than is often meant by the long run. Such concern about the length of time required for the asymptotic distribution to be relevant suggests the need for a new term. This is the ultralong-run distribution.

The important point is that sample path properties include both learning (the medium run) and apparent convergence to stable behavior (the long run). The ultralong-run phenomenon embodied in the asymptotic distribution may be very misleading as a guide to such sample path properties. We believe that in most applications to which economic theory is relevant, the time span of interest is likely to be the long run. If we are correct, then theories of equilibrium selection based on the ultralong run must bear the burden of showing that the ultralong run is not so long as to be irrelevant.

3. CASES 1 AND 2: LONG RUN DYNAMICS

Our aim in this section is to show how the musical chairs model is related to the replicator dynamics. In particular, we show that when the population size is large, finite segments of the sample paths of our musical chairs model are well approximated by the deterministic replicator dynamics.

The variant of the replicator dynamics to be studied has the form[9]

$$\dot{y} = y(1-y)\frac{g_h - g_d}{\bar{g}} + \lambda(1-2y), \qquad (1)$$

---

[8] See also, for example, the discussions of learning in Andreoni and Miller (1993), Crawford (1991, 1992), and Miller and Andreoni (1991).

[9] A more familiar form of the replicator dynamics can be written as $\dot{y} = y(\pi_d - \bar{\pi}) = y(1-y)(\pi_d - \pi_h)$, where π_d and π_h are the fitnesses attached to doves and hawks respectively. To reduce (1) to this form, begin by setting $\lambda = 0$. Next rescale time, as in Maynard Smith (1982, p. 183), to remove the denominator $\bar{g}$ from the right side of (1). Finally, definite fitnesses as the negative of death probabilities, or $\pi_d = -g_d$ and $\pi_h = -g_h$.

where y is the fraction of doves in an infinite population, $\bar{g} = yg_d + (1 - y)g_h$, and the quantities

$$g_{\mathrm{d}}(y) = a(b - 5y)$$
$$g_{\mathrm{h}}(y) = a(b - 2 - 2y)$$

are the death probabilities faced by Mopsy in the Death Game of Fig. 1b when Mopsy plays dove and hawk respectively and when the probability that Flopsy plays dove is y and the probability that Flopsy plays hawk is $1 - y$.

When the population size is N and the length of a time period is τ, the musical chairs model is a Markov process with $N + 1$ population states $z \in \{0, \nu, 2\nu, \dots, 1\}$, where $\nu = 1/N$ and zN is the number of doves in state z. We can calculate the probability $r(z, \nu, \tau)$ of moving one step to the right in a single period, from a state z with $0 \le z < 1$ to the state $z + \nu$:

$$\begin{aligned} r(z, \nu, \tau) &= \tau N(1 - \tau)^{N-1}\{z(1 - \lambda) + (1 - z)\lambda\}G_{\mathrm{h}}(z) + O(N^2\tau^2) \\ &= \tau N(1 - \tau)^{N-1}R(z, \nu) + O(N^2\tau^2), \end{aligned} \tag{2}$$

where $\tau N(1 - \tau)^{N-1} = \tau N + O(N^2\tau^2)$ is the probability that there will be exactly one birth in the period, $z(1 - \lambda) + (1 - z)\lambda$ is the probability that a dove is born given that there is a birth, and $G_{\mathrm{h}}(z)$ is the probability that a hawk will die given that a dove is born.[10] Similarly, the probability $\ell(z, \nu, \tau)$ of moving one step to the left in a single period from a state z with $0 < z \le 1$ to the state $z - \nu$ is given by

$$\begin{aligned} \ell(z, \nu, \tau) &= \tau N(1 - \tau)^{N-1}\{(1 - z)(1 - \lambda) + z\lambda\}G_{\mathrm{d}}(z) + O(N^2\tau^2) \\ &= \tau N(1 - \tau)^{N-1}L(z, \nu) + O(N^2\tau^2), \end{aligned} \tag{3}$$

where $G_{\mathrm{d}}(z)$ is the probability that a dove will die given that a hawk is born. Note that (2) and (3) implicitly define $R(z, \nu)$ and $L(z, \nu)$ by

$$R(z, \nu) = \{z(1 - \lambda) + (1 - z)\lambda\}G_{\mathrm{h}}(z) \tag{4}$$

$$L(z, \nu) = \{(1 - z)(1 - \lambda) + z\lambda\}G_{\mathrm{d}}(z). \tag{5}$$

To expose the connection between this model and the replicator dynamics, we must calculate $G_{\mathrm{h}}(z)$. For this purpose, we need the probabilities

---

[10] The $O(N^2\tau^2)$ term in (2) arises out of the possibility that there are multiple births in a period. The probability that two or more births occur and move the state one step to the right is less than or equal to the probability of two or more births, which is given by $\sum_{k=2}^{N} C_k^N \tau^k (1 - \tau)^{N-k}$, where C_k^N is the relevant binomial coefficient. This is in turn bounded above by $\sum_{k=2}^{N} N^k\tau^k \le 2N^2\tau^2 = O(N^2\tau^2)$, where the penultimate inequality requires $N\tau < \frac{1}{2}$, which is ensured by taking the limit $\tau \to 0$ before $N \to \infty$.

that various pairs of rabbits will be chosen to play the Death Game from a population containing $Nz + 1$ doves and $N(1 - z)$ hawks. The probability that a hawk and a dove will be drawn is $2z(1 - z) + O(\nu)$.[11] Similarly, the probability of drawing two hawks is $(1 - z)^2 + O(\nu)$. The probability of drawing two doves is $z^2 + O(\nu)$.

The probability that a hawk will die at the first draw can then be calculated as

$$F_{\mathrm{h}}(z) = a(b - 4)2z(1 - z) + 2a(b - 2)(1 - z)^2 + O(\nu),$$

while the probability that a dove dies at the first draw is

$$F_{\mathrm{d}}(z) = 2z(1 - z)ab + 2z^2a(b - 5) + O(\nu).$$

The probability that the rabbit who *eventually* dies will be a hawk is[12]

$$G_{\mathrm{h}}(z) = \frac{F_{\mathrm{h}}(z)}{F_{\mathrm{d}}(z) + F_{\mathrm{h}}(z)} = (1 - z)\frac{g_{\mathrm{h}}(z)}{\bar{g}} + O(\nu).$$

Similar considerations lead to $G_{\mathrm{d}}(z) = zg_{\mathrm{d}}(z)/\bar{g} + O(\nu)$.

We can now estimate the expectation and variance of $z(t + \tau) - z(t)$ conditional on $z(t) = z$. To this end, we introduce the quantities

$$\begin{aligned} \mu(z, \nu) &= R(z, \nu) - L(z, \nu) \\ &= z(1 - z)\frac{g_{\mathrm{h}} - g_{\mathrm{d}}}{\bar{g}} + \lambda(1 - 2z) + O(\nu) \qquad (6) \\ \sigma^2(z, \nu) &= R(z, \nu) + L(z, \nu) \\ &= (1 - 2\lambda)z(1 - z)\frac{g_{\mathrm{h}} + g_{\mathrm{d}}}{\bar{g}} + \lambda + O(\nu). \qquad (7) \end{aligned}$$

It is important to note that the replicator equation (1) can now be written in the form

$$\dot{y} = \mu(y, 0). \qquad (8)$$

In Theorem 1, we shall meet the perturbed form $\dot{y} = \mu(y, \nu)$ of this equation, to which a diffusion term involving $\sigma^2(z, \nu)$ will be added in Section 7.

---

[11] Given that there are $Nz + 1$ doves, $N(1 - z)$ hawks, and $N + 1$ rabbits in all, the probability of drawing a dove and a hawk is given by $2(Nz + 1)N(1 - z)/(N + 1)N = 2Nz(1 - z)/(N + 1) + 2(1 - z)/(N + 1) = 2z(1 - z) + O(\nu)$, where the final equality uses the facts that $N/(N + 1) = 1 - 1/(N + 1)$ and $1/(N + 1) = 1/N - 1/N(N + 1)$.

[12] Here, we use the fact that $(\alpha + O(\nu))/(\beta + O(\nu)) = \alpha/\beta + O(\nu)$.

Let $z(t+\tau) - z(t) \equiv \delta z(t)$. Then since $\mathscr{E}\{z(t)|z(t) = z\} = z$ and the probability of moving more than one step in a single period is $O(N^2\tau^2)$, Eqs. (6) and (7) lead to the following estimates:

$$\begin{aligned}\mathscr{E}\{\delta z(t)|z(t) = z\} &= \nu r(z, \nu, \tau) - \nu\ell(z, \nu, \tau) + O(N^2\tau^2)\\ &= \tau\mu(z, \nu) + O(N^2\tau^2), \qquad (9)\\ \mathscr{E}\{(\delta z(t))^2|z(t) = z\} &= \nu^2 r(z, \nu, \tau) + \nu^2\ell(z, \nu, \tau) + O(N^2\tau^2)\\ &= \tau\nu\sigma^2(z, \nu) + O(N^2\tau^2), \qquad (10)\end{aligned}$$

$$\begin{aligned}\operatorname{var}\{\delta z(t)|z(t) = z\} &= \tau\nu\sigma^2(z, \nu) - \{\tau\mu(z, \nu)\}^2 + O(N^2\tau^2)\\ &= \tau\nu\sigma^2(z, \nu) + O(N^2\tau^2). \qquad (11)\end{aligned}$$

Rewrite Eq. (9) as

$$\frac{\mathscr{E}\{z(t+\tau)|z(t) = z\} - \mathscr{E}\{z(t)|z(t) = z\}}{\tau} = \mu(z, \nu) + O(N^2\tau). \qquad (12)$$

When $(\tau, N) \rightarrow (0, \infty)$ in (12) so that $N^2\tau \rightarrow 0$, the right side converges to $\mu(z, 0)$, which is the right side of the replicator equation (8) with y replaced by z.

When $\tau \rightarrow 0$ and then $N \rightarrow \infty$, the left side of (12) converges to a derivative which informal derivations of the replicator dynamics often identify with $z'(t)$ without further ado. Equation (12) then reduces to the replicator equation in the limit. But, as Borgers and Sarin (1993) note, the identification of the left side of (12) with $z'(t)$ requires a formal justification. Börgers and Sarin (1993) appeal to Theorem 1.1 of Norman (1972, p. 118) in deriving the replicator dynamics for their model (in which the "population size" is infinite from the outset). We construct a proof of the following theorem along the lines of Boylan (1992, 1995).[13] In this theorem, it is important to distinguish between the state $z(t)$ of the system described by the musical chairs model and the state $y(t)$ of the system governed by the replicator equation (8). The former is a random variable, the latter is deterministic.

THEOREM 1. *Let $y(t)$ be the solution of the replicator equation $\dot{y} = \mu(y, 0)$ subject to the boundary condition $y(0) = z(0)$. Then for any $\varepsilon > 0$ and any integer $T > 0$, there exist positive constants N_0 and τ_0 such that if $N > N_0$ and $N^2\tau < \tau_0$, then*

$$\operatorname{prob}\{|y(t) - z(t)| \geq \varepsilon\} < \varepsilon,$$

for any t satisfying $0 \leq t \leq T$ at which $z(t)$ is defined.

---

[13] We thank Rob Seymour for his help in aiding our understanding of this proof.

Proof. Fix a value of $T > 0$. Unless otherwise stated, t will be assumed to be admissible, i.e., to be of the form $t = k\tau$ for some integer k. Bounds that are not given explicitly hold for all admissible t satisfying $0 \le t \le T$.

The first step is to observe that because $\mu(y, \nu)$ satisfies appropriate Lipschitz conditions, we can replace $y(t)$ by $Y(t)$, where $Y(t)$ is the solution of $\dot{Y} = \mu(Y, \nu)$ subject to the boundary condition $Y(0) = z(0)$. In particular, $|Y(t) - y(t)| < \frac{1}{2}\varepsilon$ provided that N is sufficiently large.[14] It then remains to establish the theorem with y replaced by Y and ε replaced by $\frac{1}{2}\varepsilon$.

On integrating, we obtain the following expression for $Y(t)$:

$$Y(t) - z(0) = \int_0^t \mu(Y(s), \nu)\, ds. \tag{13}$$

The next step is to find a corresponding expression for $z(t) - z(0)$. To this end, we follow Boylan (1995, p. 16) in defining

$$m(k\tau) = z(k\tau) - z(0) - \sum_{j=1}^{k} \mathscr{E}\{z(j\tau) - z(j\tau - \tau) | z(j\tau - \tau)\}. \tag{14}$$

Boylan notes that $m(k\tau)$ is a martingale, which is a fact that will be needed later. For the moment, we note only that $\mathscr{E}\{m(k\tau)\} = \mathscr{E}\{m(0)\} = m(0) = 0$.

On rearranging (14) and making use of (9), we obtain that

$$\begin{aligned} z(t) - z(0) &= m(t) + \sum_{j=1}^{k} \mathscr{E}\{z(j\tau) - z(j\tau - \tau) | z(j\tau - \tau)\} \\ &= m(t) + \sum_{j=1}^{k} \mu(z(j\tau - \tau), \nu) + O(N^2\tau^2 k) \\ &= m(t) + \int_0^t \mu(z([s/\tau]\tau), \nu)\, ds + O(N^2\tau). \end{aligned} \tag{15}$$

The next step is to subtract (15) from (13) to obtain

$$\begin{aligned} |Y(t) - z(t)| &\le |m(t)| + \int_0^t |\mu(Y(s), \nu) - \mu(z([s/\tau]\tau), \nu)|\, ds + O(N^2\tau) \\ &\le |m(t)| + C\int_0^t |Y(s) - z([s/\tau]\tau]), \nu)|\, ds + O(N^2\tau), \end{aligned} \tag{16}$$

where C has been chosen independently of ν so that $|\mu(a, \nu) - \mu(b, \nu)| \le C|a - b|$. Since $t = k\tau$, we have that $t = [t/\tau]\tau$. Hence, on writing

---

[14] From (6), there is a $C > 0$ such that for all x and y in $[0, 1]$, $|\mu(x, 0) - \mu(y, 0)| \le C|x - y|$ and $|\mu(y, \nu) - \mu(y, 0)| \le C\nu$. It is then standard to observe that $|y(t) - Y(t)| \le \int_0^t |\mu(y(s), 0) - \mu(Y(s), \nu)|\, ds \le C\{\nu t + \int_0^t |y(s)|\, ds$. It follows from Gronwall's lemma (Revuz and Yor, 1991, p. 499) that $|y(t) - Y(t)| \le C\nu t e^{Ct} < \varepsilon/2$, provided that N is sufficiently large.

$$\phi(s) = |Y(s) - z([s/\tau]\tau)|,$$

inequality (16) implies that

$$\phi(t) \le M(t) + C\int_0^t \phi(s)\,ds, \tag{17}$$

where

$$M(t) = \sup_{0\le s\le t} |m([s/\tau]\tau)| + O(N^2\tau) + O(\nu).$$

The final error term $O(\nu)$ has been added to the expression for $M(t)$ so that (17) holds for all t satisfying $0 \le t \le T$ and not just t of the form $k\tau$. The supremum in the expression for $M(t)$ ensures that M increases in t. We can then appeal directly to Gronwall's lemma, from which we deduce that

$$\phi(t) \le M(t)e^{Ct}$$

for all t satisfying $0 \le t \le T$.

Restricting ourselves again to the case where time variables take only admissible values of the form $k\tau$, we recall that $\mathscr{E}\{m(t)\} = 0$ and use Chebychev's inequality to obtain

$$\begin{aligned}
\text{prob}\{|Y(t) - z(t)| \ge \tfrac{1}{2}\varepsilon\} &\le \text{prob}\{M(t) \ge \tfrac{1}{2}\varepsilon e^{-Ct}\}\\
&\le t \max_{0\le s\le t} \text{prob}\{|m(s)| \ge \tfrac{1}{4}\varepsilon e^{-Ct}\}\\
&\le T \max_{0\le s\le T} \text{prob}\{|m(s)| \ge \tfrac{1}{4}\varepsilon e^{-CT}\} \qquad (18)\\
&\le T \max_{0\le s\le T} \frac{16e^{2CT}}{\varepsilon^2}\text{var}\{m(s)\}\\
&\le T \max_{0\le s\le T} \frac{16e^{2CT}}{\varepsilon^2}\mathscr{E}\{m(s)\}^2, \qquad (19)
\end{aligned}$$

provided that N is sufficiently large and $N^2\tau$ is sufficiently small.[15]

To make further progress, it is necessary to estimate $\mathscr{E}\{m(t)\}^2$. In so doing, we follow Boylan (1995) closely in exploiting the fact that $m(k\tau)$ is a martingale. Define

$$\Delta_i = m(i\tau) - m(i\tau - \tau).$$

---

[15] Inequality (18) is the statement that the probability of the largest of t random variables exceeding a constant is no greater than t times the largest probability that one of the random variables exceeds the constant. To obtain (18), we need N large enough to ensure that the error term $O(N^2\tau) + O(\nu)$ in the definition of $M(t)$ is less than $\varepsilon e^{-Ct}/4$.

Then, if $i > j$, we have

$$\mathscr{E}\{\Delta_i\Delta_j\} = \mathscr{E}\{\mathscr{E}\{\Delta_i\Delta_j|z(j\tau)\}\} = \mathscr{E}\{\Delta_j\mathscr{E}\{\Delta_i|z(j\tau)\}\} = 0,$$

because $\mathscr{E}\{\Delta_i|z(j\tau)\} = \mathscr{E}\{m(i\tau)|z(j\tau)\} - \mathscr{E}\{m(i\tau - \tau)|z(j\tau)\} = m(j\tau) - m(j\tau) = 0$. It follows that

$$\mathscr{E}\{m(k\tau)\}^2 = \mathscr{E}\left\{\sum_{i=1}^{k}\Delta_i\right\}^2 = \mathscr{E}\left\{\sum_{i=1}^{k}\sum_{j=1}^{k}\Delta_i\Delta_j\right\} = \sum_{i=1}^{k}\mathscr{E}\{\Delta_i^2\}. \tag{20}$$

To estimate $\mathscr{E}\{\Delta_i^2\}$, we use (11). Since $\mathscr{E}\{\Delta_i\} = \mathscr{E}\{m(i\tau)\} - \mathscr{E}\{m(i\tau - \tau)\} = 0$, we have

$$\begin{aligned}\mathscr{E}\{\Delta_i^2\} = \operatorname{var}\Delta_i &= \operatorname{var}\{z(i\tau) - z(i\tau - \tau)|z(i\tau - \tau)\}\\ &= \tau\nu\sigma^2(z(i\tau - \tau), \nu) + O(N^2\tau^2)\\ &\le \tau\nu S + O(N^2\tau^2),\end{aligned} \tag{21}$$

where S is an upper bound for $\sigma^2(z, \nu)$ that is independent of z and ν. Inserting the estimate (21) into (20), we obtain that

$$\begin{aligned}\mathscr{E}\{m(k\tau)\}^2 &\le k\tau\nu S + O(N^2\tau^2 k)\\ &= \frac{tS}{N} + O(N^2\tau)\\ &\le \frac{\varepsilon^3}{16T}e^{-2CT},\end{aligned} \tag{22}$$

provided that N is sufficiently large and $N^2\tau$ is sufficiently small.

On using the estimate (22) in (19), we find that

$$\begin{aligned}\operatorname{prob}\{|y(t) - z(t)| \ge \varepsilon\} &\le \operatorname{prob}\{|Y(t) - z(t)| \ge \tfrac{1}{2}\varepsilon\}\\ &\le T\max_{0\le s\le T}\frac{16e^{2CT}}{\varepsilon^2}\frac{\varepsilon^3}{16T}e^{-2CT}\\ &= \varepsilon,\end{aligned}$$

which yields the conclusion of the theorem. ■

Theorem 1 is a *weak* convergence result that applies when $(\tau, N) \to (0, \infty)$ so that $N^2\tau \to 0$. Note that λ is not mentioned and the limit $\lambda \to 0$ is not even notionally taken in Theorem 1. The result holds for both $\lambda > 0$ and $\lambda = 0$, and hence applies to both Cases 1 and 2 of Section 2.

Since Theorem 1 holds for finite values of t, it establishes sample-path properties of the musical chairs Markov process. These sample path properties have long-run implications. To explore these implications, suppose that

$y(t)$ is governed by the replicator dynamics (1) (or (8)) with the boundary condition $y(0) = z(0)$. Figure 2a shows a phase diagram for the replicator dynamics (1) when λ is small compared with b. There are three rest points ξ, η, and ζ satisfying $0 < \xi < \eta < \zeta < 1$. The inner rest point is unstable, but ξ and ζ are asymptotically stable with basins of attraction $[0, \eta)$ and $(\eta, 1]$ respectively. As $\lambda \rightarrow 0$, the rest points ξ, η, and ζ converge to the three Nash eqilibria of the Stag-Hunt Game. The rest point ξ converges to the risk-dominant equilibrium in which the whole population is hawks. The rest point ζ converges to the payoff-dominant equilibrium in which the whole population are doves. The rest point η converges to the mixed equilibrium in which two-thirds of the population are doves and one-third is hawks.

Suppose that $y(0) = z(0)$ happens to lie in the basin of attraction of ζ. The replicator dynamics will then move close to ζ and remain in the vicinity of ζ forever. For sufficiently large N, our continuous-time musical chairs model is arbitrarily likely to duplicate this behavior over any finite time period. More precisely, take T large enough so that, by time T, the process $y(t)$ will have spent a long period of time near ζ. Take ε to be arbitrarily small. Then for sufficiently large N and small $N^2\tau$ we have from Theorem 1 that $|y(t) - z(t)| < \varepsilon$ with probability at least $1 - \varepsilon$ for all $t \leq T$. With high probability, the musical chairs model thus also moves close to ζ and remains in the vicinity of ζ for an extended period of time. Upon observing such behavior, we would be tempted to say that the process has "selected" the equilibrium ζ. This is a long-run result in our terms.

Note, however, that when the limit $t \rightarrow \infty$ is actually taken to obtain the limiting outcome $y(\infty)$ of the replicator dynamic, the limit is applied to the process $y(t)$ and *not* to the process $z(t)$. As the next section on the ultralong-run behavior of $z(t)$ demonstrates, $z(t)$ will definitely not remain in the vicinity of $y(\infty)$ forever.

4. CASES 3 AND 4: ASYMPTOTIC DISTRIBUTION

In this section we study the asymptotic distribution over states given by the musical chairs model.

For fixed values of τ, ν, and λ, the musical chairs model is a Markov process on the state space $\{0, \nu, 2\nu, \ldots, 1\}$ with a single ergodic set consisting of the entire state space. Hence, the following standard result for Markov processes is immediate:

THEOREM 2. *Fix τ, ν, and λ. Then there exists a unique probability distribution $P(z, \nu, \tau)$ on $\{0, \nu, 2\nu, \ldots, 1\}$ that is a stationary distribution for the musical chairs Markov process. With the exception of a set of sample* *paths of measure zero, the relative frequencies of the various states along a sample path approach the distribution* $P(z, \nu, \tau)$. *For any initial condition* $z(0)$, *the probability distribution describing the likely state of the process at time* t *converges to* $P(z, \nu, \tau)$ *as* t *grows large.*

Proof. See Kemeny and Snell (1960, Theorems 4.1.4, 4.1.6, and 4.2.1). ■

For the case of small values of τ, we can derive a particularly convenient representation for $P(z, \nu, \tau)$. Let $P(z, \nu, \tau, t)$ be the probability attached to state z by the Markov process at time t, given ν and τ. Then for small values of τ, this satisfies

$$\begin{aligned} P(z, \nu, \tau, t + \tau) = {} & P(z + \nu, \nu, \tau, t)\ell(z + \nu, \nu, \tau) \\ & + P(z - \nu, \nu, \tau, t)r(z - \nu, \nu, \tau) \\ & + P(z, \nu, \tau, t)\{1 - \ell(z, v, \tau) - r(z, v, \tau)\} + O(N^2\tau^2), \end{aligned} \tag{23}$$

where $\ell(z, \nu, \tau)$ and $r(z, \nu, \tau)$ are defined by (2) and (3) in Section 2, $\ell(0, \nu, \tau) = r(1, \nu, \tau) = 0$, and $P(1 + \nu, \nu, \tau, t) = P(-\nu, \nu, \tau, t) = 0$. The error $O(N^2\tau^2)$ term in (23) arises from the possibility that there may be multiple births in a single period. The stationary distribution $P(z, \nu, \tau)$ is characterized by the following version of (23):

$$\begin{aligned} P(z, \nu, \tau) = {} & P(z + \nu, \nu, \tau)\ell(z + \nu, \nu, \tau) + P(z - \nu, \nu, \tau)r(z - \nu, \nu, \tau) \\ & + P(z, \nu, \tau)\{1 - \ell(z, v, \tau) - r(z, v, \tau)\} + O(N^2\tau^2). \end{aligned} \tag{24}$$

We find it helpful to think of this stationary distribution as corresponding to taking the limit $t \rightarrow \infty$ in $P(z, \nu, \tau, t)$. This is the first limit to be taken in Case 4. On rearranging (24), we obtain

$$\begin{aligned} 0 = {} & P(z + \nu, \nu, \tau)\ell(z + \nu, \nu) - P(z, \nu, \tau)\ell(z, v, \tau) \\ & + P(z - \nu, \nu, \tau)r(z - \nu, \nu, \tau) - P(z, \nu, \tau)r(z, v, \tau) + O(N^2\tau^2). \end{aligned} \tag{25}$$

The next limit to take in Casea 4 is $\tau \rightarrow 0$. As we let $\tau \rightarrow 0$, the probability of more than one birth in a single period becomes negligible. Transitions then consist of either a single step to the left or a single step to the right. The result is a *birth–death chain* (cf. Gardiner, 1985). The following description of the stationary distribution is standard for such chains.

THEOREM 3. *Let* $P(z, \nu) = \lim_{\tau \to 0} P(z, \nu, \tau)$. *Then for* $z \in \{0, \nu, 2\nu, \ldots, 1\}$, *we have*

$$\frac{P(z + \nu, \nu)}{P(z, \nu)} = \frac{R(z, \nu)}{L(z + \nu, \nu)}. \tag{26}$$

Proof. Summing (25) for values z and $(z + \nu)$ and taking the limit as $\tau \rightarrow 0$ gives (26). ■

An alternative route to this conclusion, which isolates the nature of the limiting argument as $\tau \rightarrow 0$, uses the method of Freidlin and Wentzell (1984), as employed by Young (1993) and Kandori *et al.* (1993). One first constructs, for each state z in the finite Markov process, the collection of all z-trees, where a z-tree is a collection of transitions between states with the property that each state other than z is the origin of one and only one transition, z is the origin of no transition, and there is a path of transitions to z from every state other than z. One then computes the products of the transition probabilities in each z-tree and adds these products (for all of the z-trees) to obtain a number $Z(z)$. The exact ultralong-run probabilities $P(z, \nu)$ and $P(z', \nu)$ of being at states z and z' are then proportional to $Z(z)$ and $Z(z')$. In our model, products of order greater than τ^N can be neglected, since the limit $\tau \rightarrow 0$ is to be taken.[16] Only the single z-tree of order $O(\tau^N)$ illustrated in Fig. 2c then needs to be retained. When the products for z and $z + \nu$ are divided, only one factor from each does not cancel, leaving (26).

Yet another path to this conclusion is provided by an analysis that we represent with limits in the order given by Case 3. First, we rewrite Eq. (23) in the form

$$\begin{aligned}\nu\left\{\frac{P(z,\nu,\tau,t+\tau)-P(z,\nu,\tau,t)}{\tau(1-\tau)^{N-1}}\right\} &= P(z+\nu,\nu,\tau,t)L(z+\nu,\nu)\\ &\quad -P(z,\nu,\tau,t)L(z,v)\\ &\quad + P(z-\nu,\nu,\tau,t)R(z-\nu,\nu)\\ &\quad - P(z,\nu,\tau,t)R(z,v) + O(N^2\tau^2).\end{aligned}$$

Now take the limit $\tau \rightarrow 0$ to obtain a continuous-time process. The error term on the right disappears and we are left with the probability $P(z, \nu, t)$ that satisfies

$$\begin{aligned}\nu\frac{\partial P(z,\nu,t)}{\partial t} &= P(z+\nu,\nu,t)L(z+\nu,\nu) - P(z,\nu,t)L(z,v)\\ &\quad + P(z-\nu,\nu,t)R(z-\nu,\nu) - P(z,\nu,t)R(z,v).\end{aligned} \tag{27}$$

Now we take the limit $t \rightarrow \infty$, which is to say that we examine the stationary distribution of this process, if one exists. A stationary distribution is defined by the property that $\partial P(z, \nu, t)/\partial t = 0$. Writing $\partial P(z, \nu, t)/\partial t = 0$ in (27) leads us to Eq. (25) with a zero error term. Theorem 3 therefore still applies.

---

[16] For similar reasons, Young (1993) and Kandori *et al.* (1993) can neglect products that do not minimize the order of λ because the limit $\lambda \rightarrow 0$ is taken.

It follows that the stationary distribution for Case 3 exists and is the same as for Case 4. We record this formally as

THEOREM 4. *Cases 3 and 4 yield identical stationary distributions.*

Waiting Times. What is the relationship between the replicator approximation of the musical chairs model and the asymptotic distribution? To bring this question into sharper focus, suppose that the initial condition lies in the basin of attraction of ζ relative to the replicator dynamics. As a result, the system will initially be highly likely to approach ζ and spend an extended period of time in the neighborhood of ζ. The larger N gets, the more likely this will be a good description of the behavior of the system and the longer the "long run" over which it will be valid. Yet the asymptotic distribution may put almost all of its probability masses on ξ. Hence, the ultralong-run behavior described by the asymptotic distribution need not provide useful information about the long-run behavior of the system.

To reconcile these observations, we note that a transition from a state near ζ to a state in the basin of attraction of ξ requires a rather special realization of random events. These events require a much larger proportion of hawk births and dove deaths than the system is expected to produce. During any bounded time interval, the probability of such a combination of hawk births and dove deaths, and hence the probability that the musical chairs population will reach the basin of attraction of ξ, is extraordinarily small and one must wait an extraordinarily long time for such an event to become likely. However, "extraordinarily long" is not the same as "forever." The probability that such a combination of births and deaths will *eventually* occur is unity. As a result, the system will, with certainty, visit states near ξ and will spend virtually all of its time in such states.

Should we be interested in the long run or the ultralong run? The answer to this question turns on the length of time required to reach the stationary distribution.

We gain some insight into this question by comparing the expected waiting times in our model with those obtained from the model of Kandori *et al.* (1993) for the Stag-Hunt Game given in Fig. 1a. When examining our model, we work in the limiting case as $\tau \to 0$, meaning that we work with the birth–death process whose stationary distribution is described by (3). We let $N = 100$ and $\lambda = 0.001$. To make the models comparable, we let each agent in the Kandori *et al.* model revise his strategy once in each unit of time, while in our model the probability that an agent gives birth in a period of length τ is τ. Hence, in both models, changes in an agent's strategy are expected to come at the rate of one per unit of time.

In the Kandori *et al.* model (1993), the expected waiting time to be estimated is the same as that for escaping from the basin of attraction of the payoff-dominant equilibrium.[17] In each period, such an escape requires at least 33 simultaneous mutations, an event whose probability is $q = (0.001)^{33}(0.999)^{67}(100!/33!67!)$. This leads to the approximation $1/q = 1.7 \times 10^{72}$ units of time for the expected number of periods before the risk-dominant equilibrium is reached.

Following Gardiner (1985, Chap. 7), one easily derives the following formula for the corresponding waiting time in our model, where $\nu = 1/100$ and $R(y, \nu)$ is defined by (5),[18]

$$\sum_{y=0}^{99} \frac{1}{100R(y, \nu)P(y, \nu)} \sum_{z=y+1}^{100} P(z, \nu). \tag{28}$$

When $b = 10$, the expected waiting time is approximately 5,000 units of time.

Our expected waiting time is much shorter than that of Kandori *et al.* (1993). The driving force behind this difference is that our selection process is noisy while that of Kandori *et al.* is not. In particular, suppose the system begins in the basin of attraction of the payoff-dominant equilibrium. Because the Kandori *et al.* learning process always calls for agents to choose best replies, the only hope for escaping this equilibrium is that enough simultaneous mutations occur to reach the basin of attraction of the risk-dominant equilibrium. In our example, this required 33 mutations, a very unlikely event when λ is small. In the Musical Chairs model, a mutation is again required to introduce the hawk strategy into the population. But once this mutation has occurred, the noisy learning process can cause more agents to switch to hawk. It is still an unlikely event that noisy learning causes enough agents to switch to hawk as to reach the risk-dominant equilibrium, but if the mutation probability is small this event is much more likely than the number of simultaneous mutations required to reach the risk-dominant equilibrium in the Kandori *et al.* model.

It is interesting to note that the model of Young (1993) differs from that of Kandori *et al.* in allowing learning to be noisy. In particular, before agents play in Young's model they observe k of the most recent m plays of the game, where $k < m$. The randomness involved in drawing this sample of size k produces noisy learning. This suggests that Young's model should give faster convergence than does that of Kandori *et al.* and suggests that convergence will be faster in Young's model as k is smaller relative to m.

On the one hand, we are encouraged by the relatively short waiting time

---

[17] Kandori *et al.* report much shorter waiting times which apply to small values of N and larger values of λ (e.g., $\lambda = 0.1$), or are averages over the waiting times for initial conditions that are randomly selected according to probabilities given by the stationary distribution.

[18] This formula corrects a typo in formula (7.4.13) of Gardiner (1985).

of our model, viewing the use of a noisy selection process as both a step toward realism and also a path to shorter expected waiting times. However, our expected waiting times are still sufficiently large that it will often be the long run rather than the ultralong run that is relevant to the equilibrium observed in practice.[19] We think this conclusion is likely to be robust whenever the inevitable noise in the selection process is realistically modeled. It is disappointing that the sharp, history-independent predictions of an ultralong-run analysis should have a limited domain of application. However, the fact has to be faced that, for practical purposes, one will usually need access to historical data in order to conduct a long-run analysis of the model.

5. EQUILIBRIUM SELECTION

Equilibrium. In this section, we consider the ultralong-run equilibrium selection implications of the musical chairs model. In particular, we examine $P(z, \nu)$ as the population gets large and then the mutation rate gets small. We say that an equilibrium E is selected in the ultralong run if, given any $\varepsilon > 0$, the stationary distribution converges weakly to a distribution that places unit mass on E as first $N \to \infty$ and then $\lambda \to 0$.

The key to our investigation is the characterization of the asymptotic distribution $P(z, \nu)$ given by (26). Figure 3 shows $P(z, \nu)$ and $\mu(z, \nu)$ for a fixed but small value of ν (so that the error term $O(\nu)$ is small).[20] Note first (from (26)) that $P(z, \nu)$ increases exponentially where $\mu(z, \nu) \geq \theta > 0$ and decreases exponentially where $\mu(z, \nu) \leq \theta < 0$. It follows that as long as N is reasonably large, the probability mass of $P(z, \nu)$ must be concentrated in the neighborhood of the points z at which $\mu(z, \nu) = 0$. The solutions of this equation approximate the rest points ξ, η, and ζ of the replicator dynamics (1). As previously, only ξ and ζ will be serious candidates for consideration, and in the limiting case when $N \to \infty$ all of the probability mass of the limiting distribution $P(z, \nu)$ will be arbitrarily close to one of these points.

Which of ξ and ζ will be selected? Let ξ_ν and ζ_ν be the local maxima of $P(z, \nu)$. Continuity ensures that $\xi_\nu \to \xi$ and $\zeta_\nu \to \zeta$ as $N \to \infty$. From (26),

---

[19] How can we assert that 5000 is long without having an idea of the units in which time is measured? Recall that agents are expected to change their strategies at the rate of once per unit of time. A waiting time of 5000 is then sufficiently long to expect each agent to change strategies 5000 times, which strikes us as long.

[20] The function $\mu(z, \nu)$ is meaningful only for values of $z \in \{0, \nu, 2\nu, \ldots, 1\}$, but we find it convenient to draw it as a continuous function on $[0, 1]$.

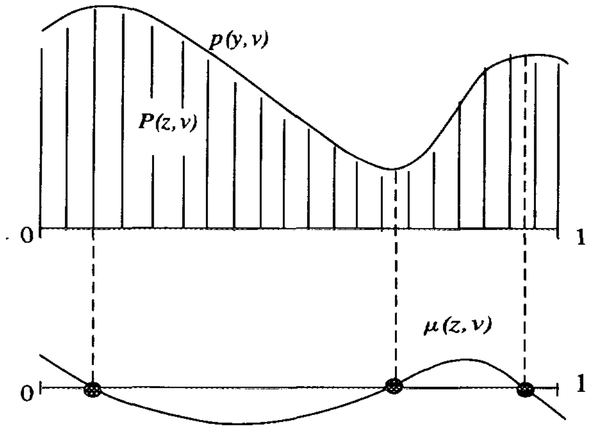

FIG. 3. Probability distributions.

the ratio

$$\frac{P(\zeta_\nu,\nu)}{P(\xi_\nu,\nu)} = \prod_{\xi_\nu \le z < \zeta_\nu} \frac{R(z,\nu)}{L(z+\nu,\nu)} \tag{29}$$

then converges (as $N \to \infty$) to the ratio of the probability masses placed on ξ and ζ. After taking the limit $N \to \infty$, we let $\lambda \to 0$ so that ξ and ζ converge to the equilibria of the Stag-Hunt Game.

THEOREM 5. *Let $b = b^* \approx 5.08$ solve*

$$\left(\frac{b-2}{b-4}\right)^{(1/2)b}\left(\frac{b-5}{b}\right)^{(1/5)b} = \frac{(b-2)(b-5)}{(b-4)^2}.$$

Then

$$\lim_{\lambda\to 0}\lim_{\nu\to 0}\frac{P(\zeta_\nu,\nu)}{P(\xi_\nu,\nu)} = \begin{cases} \infty, & \text{if } b < b^*, \\ 0, & \text{if } b > b^*. \end{cases}$$

Proof. Let $\delta > 0$. If $0 < \delta \le x \le 1 - \delta < 1$, then we have from Section 3 that

$$\lim_{\nu\to 0}\frac{R(x,\nu)}{L(x+\nu,\nu)} = \frac{g_h(x)}{g_d(x)} + O(\lambda).$$

After taking logarithms in (29), we express the sum that appears on the right as the product of N and an integral. Taking the limit $\lambda \to 0$ in the integral, we obtain

$$\int_0^1 \log \frac{g_h(x)}{g_d(x)}\, dx = \int_0^1 \log\left(\frac{b - 2x - 2}{b - 5x}\right) dx. \tag{30}$$

The integral is positive when $5 < b < b^*$ and negative when $b > b^*$. ■

Because ξ corresponds to the risk-dominant equilibrium when $\lambda = 0$, and is selected when (30) is negative, we have

COROLLARY 1. *The risk-dominant equilibrium is selected when $b > b^*$, but the payoff-dominant equilibrium is selected when $b < b^*$.*

Why does this result differ from that of Kandori *et al.*(1993), who uses the techniques of Freidlin and Wentzell to show that the risk-dominant equilibrium is always selected in their model in the ultralong run? Kandori *et al.* first show that in their model, the deterministic best-reply learning dynamics ensure that the system spends virtually all of its time at one of the two strict Nash equilibria. How are the relative probabilities attached to these two equilibria determined?

To move from one equilibrium to the other in the Kandori *et al.* model requires enough simultaneous mutations for the system to jump from an equilibrium to the basin of attraction of the other equilibrium. Such a jump is less likely to happen for the equilibrium with the larger basin of attraction, i.e., the risk-dominant equilibrium, because it requires more mutations. As the limit in which the mutation probability becomes arbitrarily small is examined, the system becomes arbitrarily less likely to switch from the risk-dominant equilibrium than from the other equilibrium, and the asymptotic distribution accordingly attaches virtually all of its limiting probability to the risk-dominant equilibrium.

In our musical chairs model, only the transitions away from the endpoints require mutations. All other transitions can occur as a result of the noisiness of the selection model. In the limiting case of rare mutations, transitions arising out of noisy selection become much more likely than transitions driven by mutations, and the latter play a very small role in equilibrium selection. Equilibrium selection results are then driven by the relative ease with which noisy selections can lead the system from one equilibrium, through its basin of attraction, to the basin of attraction of the other equilibrium. In essence, the system "swims upstream" out of the basin of attraction of an equilibrium, rather than jumping over the basin.[21] Both the sizes of basins of attraction and the magnitude of the "current" against which one must swim in making an escape are then relevant (where the strength of the current in a state depends on the relative magnitudes of the probabilities that the system moves to the left or right). As b increases, transitions have less and less to do with the payoffs in the game, and the ratio of left and right transition probabilities at any state approaches unity. For sufficiently large b, only the sizes of the basins of attraction then matter, and forces analogous to those appearing in Kandori *et al.* lead us to the risk-dominant equilibrium. For smaller values of b, however, the strengths of the currents within basins can overwhelm the different sizes of the basins, leading to the selection of the payoff-dominant equilibrium.

Making Mutation Negligible Prematurely. In Theorem 5, the limits are taken by first allowing the population size to approach infinity and then allowing the mutation probability to approach zero. What if we reversed the order of these limits? In 2 × 2 games with two strict Nash equilibria, such as the Stag-Hunt Game, reversing the order of limits has no effect, and the order is arbitrary. This is not always the case, however. The order does matter in 2 × 2 games with a single, mixed-strategy Nash equilibrium.

To illustrate this point, we briefly consider the Chicken Game of Fig. 1c, which is a special case of the Hawk–Dove game of Maynard Smith (1982). Chicken has a unique symmetric equilibrium in which each player plays dove and hawk with equal probability. The replicator dynamics lead to this equilibrium and the asymptotic distribution concentrates all of its probability near this equilibrium given that $N \to \infty$ before $\lambda \to 0$. Since there are no rival symmetric equilibria, its selection should come as no surprise—given that an equilibrium is to be selected at all.[22]

Suppose that the order of the limits $N \to \infty$ and $\lambda \to 0$ is reversed. Before the limit $N \to \infty$ is taken, the system will then have a positive probability of reaching any state z from any initial state $z(0)$ satisfying $0 < z(0) < 1$. But if $\lambda = 0$, the boundary states $z = 0$ and $z = 1$, in which the population

---

[21] Similar considerations, including the swimming upstream analogy, appear in Fudenberg and Harris (1992).

[22] The two asymmetric pure Nash equilibria are not candidates for selection because the musical chairs model has only one population from which both players are always drawn. To select an asymmetric equilibrium, it would be necessary to draw the two players from different populations that evolve separately.

consists either of all hawks or all doves, are absorbing states. Once the population enters such a state, it cannot leave without a mutation. In the ultralong run, the system will therefore necessarily lie in one of these two nonequilibrium states, with all rabbits playing a payoff-inferior strategy. This limit persists as $N \to \infty$. The outcome of the system is then determined by accidental extinctions from which the population cannot recover. This contrasts with the selection of the mixed equilibrium when $N \to \infty$ before $\lambda \to 0$.

The latter strikes us as the more appropriate ultralong-run model. As a result, our belief is that if one is interested in large populations and small mutation rates, the limit $\lambda \to 0$ should be taken last (or not taken at all).

6. PAYOFFS VERSUS FITNESSES

In conventional analyses, affine transformations of payoffs have no effect on the analysis. This is not the case in the musical chairs model. In particular, changes in the parameter b affect both the sample-path properties of the model and the asymptotic distribution. As b or λ increases, the system gets noisier in the sense that relative rewards become less important in driving transitions.

Consider first the replicator dynamics and the three rest points given by ξ, η, and ζ. As b or λ increase, ξ and η increase and ζ decreases. When $\lambda(9b - 40)$ gets to about 2, ξ is approximately $\frac{1}{9}$ and η and ζ coincide near $\frac{7}{9}$. When b or λ increases further, only ξ survives as a rest point. As $b \to \infty$ or $\lambda \to 1$, $\xi \to \frac{1}{2}$. It follows that the asymptotically stable state ξ nearest the risk-dominant equilibrium of the Stag-Hunt Game is more likely to include $z(0)$ in its basin of attraction for the replicator dynamics (1) as the system gets noisier.

Alternatively, fix $b > 5$ and let $\lambda \to 0$, so that the system gets less noisy. We then have the following estimates:

$$\xi = 0 + \tfrac{1}{2}(b - 2)\lambda + O(\lambda^2)$$

$$\eta = \tfrac{2}{3} + \tfrac{1}{6}(3b - 10)\lambda + O(\lambda^2)$$

$$\zeta = 1 - (b - 5)\lambda + O(\lambda^2).$$

Note that, since $\xi \to 0$, $\eta \to \frac{2}{3}$, and $\zeta \to 1$ as $\lambda \to 0$, the long-run predictions for models approximated by the replicator dynamics are close to those that would be obtained from the naive best-response dynamics for the Stag-Hunt Game, which are illustrated in Fig. 2b.

Now consider the asymptotic distribution. From Theorem 5, we can again note that it is unsafe to treat affine transformations of the payoffs in a game as being strategically irrelevant. Instead, the risk-dominant equilibrium is more likely to be selected when b is large, and hence selection is relatively noisy.

In the classical replicator dynamics without mutation, absolute fitness levels are irrelevant. The entries in the game matrices in an evolutionary context are therefore often taken to be "incremental fitnesses" above some arbitrary background level, and long-run predictions are unaffected by transformations that preserve these incremental fitnesses. Our results show that the same is not true for ultralong-run behavior in the musical chairs model and, from Eq. (1) of Section 2, is not true for the classical replicator dynamics, and hence long-run behavior, when the mutation rate is positive.

The parameter b can be seen as a measure of the extent to which deaths are due to causes that have nothing to do with the strategic behavior of the rabbits. As such, it is a measure of *background* antifitness. The effect of increasing b is to increase the noise with which the rabbits must cope. As b increases, their lives become increasingly hazardous. It is therefore appropriate that they end up at the risk-dominant equilibrium in the ultralong run when b is sufficiently large. But this is only one side of the coin. As b decreases, so that the manner in which the rabbits play becomes increasingly important to their survival prospects, it is the payoff-dominant equilibrium that eventually becomes the ultralong-run choice. One might say that the more important the game is to the rabbits, the more likely it is that they will be found playing the payoff-dominant equilibrium.

Biologists often escape the necessity of studying the relationship between payoffs and fitnesses by simply taking payoffs to *be* fitnesses. In economic applications, this luxury is seldom available. The payoffs with which we work commonly represent monetary rewards. Fitnesses refer to the relationship between these monetary rewards and the learning process by which agents switch strategies. This relationship is unlikely to be so straightforward as to allow us simply to equate rewards and fitnesses, and a more detailed study of the relationship between the two is essential. Binmore and Samuelson (1994) begins such a study.

7. Diffusion Theory

The literature contains a host of techniques involving continuous-time models based on stochastic differential equations that allow the stationary distribution to be conveniently approximated. We show in this section that one such approximation is not always sufficiently sharp.[23] This provides one reason to be wary when models in which all the variables are continuous are written down directly and standard techniques applied without being derived from more primitive discrete models.[24]

We consider a standard textbook procedure from diffusion theory (Risken, 1984, pp. 77–81). As specified in Case 3, we first take the limit $\tau \to 0$ to obtain Eq. (27). Following Risken, we then replace the right side of Eq. (27) by an approximation *before* taking the limit $t \to \infty$. To obtain this approximation, we expand the functions on the right of (27) around their first variables in a Taylor series in powers of ν and then discard terms of order higher than ν^2. This leads to a version of the classical Fokker–Planck equation

$$\frac{\partial p}{\partial t} = -\frac{\partial}{\partial y}\{\mu p\} + \tfrac{1}{2}\nu\frac{\partial^2}{\partial y^2}\{\sigma^2 p\}, \tag{31}$$

where z has been replaced by y to signal that it will now be treated as a continuous variable and $P(z, \nu, t)$ has been replaced by $p(y, \nu, t)$ to emphasize that we are now working only with an approximation. Recall that the functions μ and σ^2 are defined by (9) and (10).

The diffusion problems that physicists study with the aid of the Fokker–Planck equation are not new to economics, but they have usually been studied in the past using the theory of stochastic differential equations. The conclusions in this section are the same as would be obtained had we begun with the stochastic differential equation

$$dy = \mu\, dt + \nu\sigma^2\, dW, \tag{32}$$

with reflecting boundaries at $y = 0$ and $y = 1$, where dW is the standard Wiener process.[25] If we write $\nu = 0$ in (32), then the stochastic term $\nu\sigma^2$ dW disappears and the equation reduces to the replicator equation in the form (8).

---

[23] Evolutionary models involving stochastic differential equations are examined by Cabrales (1993), Foster and Young (1990), and Fudenberg and Harris (1992). Direct comparisons are difficult because our model has been constructed with a view to studying different phenomena. For example, the principle source of noise in the Fudenberg and Harris model consists of shocks to the payoff matrix, whereas our payoff matrix remains fixed. We will use Case 3 as our scheme for taking limits in this section. We know from Theorem 4 that Cases 3 and 4 yield identical stationary distributions, characterized by (26) and Corollary 1. The fact that we shall be led to conclusions different from Corollary 1 does not imply that Theorem 4 is false—it implies that the approximation to be used here is inadequately sharp.

[24] Physicists seldom do this, and when they do they are guided by physical intuitions that seldom have reliable analogues in social science applications.

[25] Note that our σ^2 is not constant as in Foster and Young (1990). When $\lambda = 0$, for example, $\sigma^2(z, \nu) \to 0$ as $z \to 0$ or $z \to 1$.

We now take the limit $t \to \infty$ in the Fokker–Planck equation (31), meaning that we look for a stationary distribution $p(y, \nu)$ that satisfies (31) with $\partial p(y, \nu)/\partial t = 0$. We thereby derive the following second-order, ordinary differential equation:

$$0 = -\frac{d}{dy}\{\mu p\} + \tfrac{1}{2}\nu\frac{d^2}{dy^2}\{\sigma^2 p\}. \tag{33}$$

The solution to this equation, $p(y, \nu)$, is an approximation of the stationary distribution of the musical chairs Markov process.[26]

The limits $\tau \to 0$ and $t \to \infty$ have now been taken, as specified by Case 3 (albeit with an intervening approximation). We have yet to take the limits $N \to \infty$ and $\lambda \to 0$. If the limit $N \to \infty$ is taken before solving (33), then we neglect the second-order term on the right side of (33), yielding a first-order differential equation with solution $\mu p = 0$.[27] The resulting approximation $p(y, 0)$ to the ultralong-run probability of finding the system near y would then be zero except at the zeros ξ, η, and ζ of $\mu(y, 0)$ (the rest points of the replicator dynamic (8)). At $y = \xi$, $y = \eta$, or $y = \zeta$, the values of $p(y, 0)$ would be indeterminate. Taking the limit $N \to \infty$ *before* solving the differential equation (33) therefore provides no useful information about the ultralong-run behavior of the system. In essence, it returns us to the long-run case studied in Section 3. If we are to have any hope of gaining information about the ultralong run, it is then essential to delay taking the limit $N \to \infty$ in (33) until *after* the equation has been solved. The reason may be more transparent in the stochastic differential equation formulation (32). It is the diffusion term $\nu\sigma^2\, dW$ that does all the work in shifting the system between equilibria in the ultralong run, and it cannot be neglected (by taking $N \to \infty$ before solving the equation) in an ultralong-run analysis. In particular, it is this term that provides information about the relative magnitudes of $p(\xi, \nu)$, $p(\eta, \nu)$, and $p(\zeta, \nu)$.

Neglecting the second-order term in (33) is equivalent to neglecting terms of order ν^2 and higher in (27). But what of terms of order ν^3 and higher in (27)? Surely these can be neglected, since they will become negligible compared with terms of order ν^2 when the limit $N \to$ is taken. Natural though this conclusion may seem, we shall see that it is mistaken. After solving (33) and taking the limits $N \to \infty$ and $\lambda \to 0$, we do not always come close to the stationary distribution of the musical chairs Markov process. It follows that there is a significant loss of accuracy involved in making the approximation that leads to the Fokker–Planck equation. This

---

[26] That is, $p(A, \nu)$ is an approximation to the ultralong-run probability of finding the system at one of the population states z in a subset A of $[0, 1]$.

[27] The constant of integration is zero because, for example, $\mu(\xi) = 0$.

loss of accuracy is presumably unimportant in the physics applications for which it is employed, but is not unimportant in our musical chairs model.

Returning to this section's program, we now integrate (33) before taking the limits $N \to \infty$ and $\lambda \to 0$. Integrating once between 0 and Y yields the equation

$$0 = -\mu(Y, v)p(Y, v) + \tfrac{1}{2}\nu \frac{d}{dy}\{\sigma^2(Y, v)p(Y, v)\}. \tag{34}$$

Integrating again, we obtain

$$p(y, \nu) = \frac{C_N}{\sigma^2(y, \nu)} \exp\left\{2N \int_0^y \frac{\mu(Y, \nu)}{\sigma^2(Y, \nu)}\, dY\right\}, \tag{35}$$

where C_N is a constant. The integral on the right side of (35) is referred to as a potential in the diffusion literature.

The shape of $p(y, \nu)$ is illustrated in Fig. 3. Pontryagin *et al.* and Vitt (1989, p. 338) discuss such distributions in some detail. Matters are simpler for us because we need only consider the relative heights of the peaks near ξ and ζ, as all of the probability of the stationary distribution is concentrated on ξ and ζ when N is large and λ small. For this purpose, the ratio

$$\frac{p(\zeta, \nu)}{p(\xi, \nu)} = \frac{\sigma^2(\xi, \nu)}{\sigma^2(\zeta, \nu)} \exp\left\{2N \int_\xi^\zeta \frac{\mu(y, \nu)}{\sigma^2(y, \nu)}\, dy\right\} \tag{36}$$

suffices, in which we take the limits $N \to \infty$ and $\lambda \to 0$.

THEOREM 6.

$$\lim_{\lambda \to 0} \lim_{N \to \infty} \frac{p(\zeta, \nu)}{p(\xi, \nu)} = 0.$$

Proof. First take the limit $\nu \to 0$ and then the limit $\lambda \to 0$ in the integral of (36) to obtain

$$\int_0^1 \left\{\frac{g_h(y) - g_d(y)}{g_h(y) + g_d(y)}\right\} dy = -\frac{3}{7} + \frac{2}{49}(10 - 3b) \log\left\{\frac{2b - 9}{2b - 2}\right\}. \tag{37}$$

Express the right side of (37) in the form $f(x)$, where $x = 7/(2b - 2)$. When $b = 5$, $x = \frac{7}{8}$. Expand $f(x)$ in powers of x. The coefficient of x^2 is zero, but higher powers have positive coefficients. It follows that f is strictly convex on $(0, 1)$. Since $f(0) = 0$ and $f(\frac{7}{8})$ is (just) negative, it follows that (37) is negative for all $b > 5$. For all $b > 5$, the limiting value of (36) is therefore zero. ■

Recalling that ξ corresponds to the risk-dominant equilibrium, Theorem (6) tells us that a system governed by the stochastic differential equation (32) would be close to the risk-dominant equilibrium nearly all the time in the ultralong run, regardless of the background fitness level b, provided that the mutation rate λ is sufficiently small and the population size N is sufficiently large. The approximation we are working with in this section is therefore insufficiently accurate for our purposes, since Section 5 showed that the payoff-dominant equilibrium is sometimes selected.

The replicator dynamics do not provide a good approximation of the asymptotic distribution of the process because they fail to capture certain transitions that play a crucial role in shaping the asymptotic distribution. The approximation of this section captures more transitions, and hence provides more hope of being a good tool for studying the asymptotic distribution, but still falls short.[28] The value of the integral (37) to which we are led in the current section is close to the integral (30) to which we were led in the preceding section, because the estimate $\log \theta \approx 2(\theta - 1)/(\theta + 1)$ is quite accurate. In consequence, the range of values of b over which the approximation of this section leads to a wrong prediction of the equilibrium selected in the ultralong run is small.

Is the approximation of this section adequate for such purposes as estimating expected waiting times, provided b is safely in the range where it provides a good approximation of the asymptotic distribution? We present an example. The following formula, adapted slightly from (24) of Pontryagin *et al.* (1989, p. 342), gives the expected waiting time for a process governed by the stochastic differential equation (32) to get from ζ to ξ:

$$\int_{\xi}^{\zeta} dy \int_{0}^{y} \frac{2N}{\sigma^2(w, \nu)} \exp\left(2N \int_{y}^{w} \frac{\mu(s, \nu)}{\sigma^2(s, \nu)}\, ds\right) dw. \tag{38}$$

---

[28] Could we not have drawn this conclusion simply from the observation that we are examining the case in which $N \to \infty$, which reduces the stochastic differential equation in (32) to the replicator dynamics? We might use the fact that (32) approaches the replicator dynamics as $N \to \infty$ to show that their sample-path properties and hence their long-run behaviors become similar, though we would expect (32) to provide a better approximation than does the replicator dynamics for finite N. But this does not tell us anything about how well (31) performs in approximating the asymptotic distribution, even though we know the replicator dynamics do not provide a good approximation. In particular, the asymptotic distributions derived from the replicator dynamics (8) and the Fokker–Planck equation (31) for large N are quite different. The latter concentrates all of its probability around the risk-dominant equilibrium, while the former attaches probability to either the risk-dominant or payoff-dominant equilibrium, depending upon initial conditions.

When $N = 100$, $\lambda = 0.001$, and $b = 10$, (38) is approximately 4,400. Recall that, using the exact model of Section 4, we estimated the corresponding waiting time as 5,000.

8. CONCLUSION

This paper has discussed a particular noisy evolutionary process. The process differs from that considered by Kandori *et al.* (1993) in that noise is intrinsic to the selection process rather than being derived entirely from mutations. We found that a version of the classical replicator dynamics suffices to describe long-run equilibrium selection. In discussing equilibrium selection in the ultralong run, we considered two approaches: the use of differential equations and the use of combinatorial methods. Both have their advantages, but only the second provided an exact result in our simple case. Finally, we confirmed that a model like ours is able to generate substantial reductions in the expected transit times between equilibria in the ultralong run, but waiting times are still sufficiently long that a long-run analysis will typically be more appropriate in applications.

Unlike Young (1993) and Kandori *et al.* (1993), we found that the payoff-dominant equilibrium is sometimes selected rather than the risk-dominant equilibrium in the ultralong run. The parameter that determines which is selected would normally be regarded as strategically irrelevant. Its relevance in this model is a reminder that in evolutionary analysis, the relevant payoffs are not the monetary rewards or even utilities that often appear in games, but rather fitnesses that are determined by the interaction of these rewards and the selection process.

REFERENCES

ANDERLINI, L. (1990). "Communication, Computability, and Common Interest Games," mimeo. St. John's College, Cambridge.

ANDREONI, J., AND MILLER, J. H. (1993). "Auctions with Artificial Adaptive Agents," mimeo. University of Wisconsin–Madison and Carnegie Mellon University, Pittsburgh.

BINMORE, K., AND SAMUELSON, L. (1993). "An Economist's Perspective on the Evolution of Norms," *J. Institutional Theoret. Econ.* **150,** 45–63.

BINMORE, K., AND SAMUELSON, L. (1994). "Muddling Through: Noisy Equilibrium Selection," mimeo. University of Wisconsin–Madison.

BOMZE, I., AND EICHBERGER, J. (1992). "Evolutive versus Naive Bayesian Learning," mimeo. Institute of Statistics and Computer Science, University of Vienna.

BÖRGERS, T., AND SARIN, R. (1993). "Learning through Reinforcement and Replicator Dynamics," mimeo. University College London.

BOYLAN, R. (1992). "Laws of Large Numbers for Dynamical Systems with Randomly Matched Individuals," *J. Econ. Theory* **57,** 473–504.

Boylan, R. (1995). "Continuous Approximation of Dynamical Systems with Randomly Matched Individuals," *J. Econ. Theory*, in press.

Cabrales, A. (1993). "Stochastic Replicator Dynamics," mimeo. University of California at San Diego.

Cabrales, A., and Sobel, J. (1992). "On the Limit Points of Discrete Selection Dynamics," *J. Econ. Theory* **57,** 407–419.

Carlsson, H., and van Damme, E. (1993). "Global Games and Equilibrium Selection," *Econometrica* **61,** 989–1018.

Crawford, V. (1991). "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games Econ. Behav.* **3,** 25–59.

Crawford, V. (1992). "Adaptive Dynamics in Coordination Games," mimeo. Department of Economics, University of California, San Diego.

Dekel, E., and Scotchmer, S. (1992). "On the Evolution of Optimizing Behavior," *J. Econ. Theory* **57,** 392–406.

Ellison, G. (1992). "Learning, Local Interaction, and Coordination," *Econometrica* **61,** 1047–1072.

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

Freidlin, M., and Wentzell, A. (1984). *Random Perturbations of Dynamical Systems.* New York: Springer-Verlag.

Fudenberg, D., and Harris, C. (1992). "Evolutionary Dynamics with Aggregate Shocks," *J. Econ. Theory* **57,** 420–441.

Gardiner, C. (1985). *Handbook of Stochastic Methods.* Berlin: Springer-Verlag.

Harsanyi, J., and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games.* Cambridge, MA: MIT Press.

Hofbauer, J., and Sigmund, K. (1988). *The Theory of Evolution and Dynamical Systems.* Cambridge, MA: Cambridge Univ. Press.

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

Kemeny, J., and Snell, J. (1960). *Finite Markov Chains.* Princeton, NJ: Van Nostrand.

Maynard Smith, J. (1982). *Evolution and the Theory of Games.* Cambridge, MA: Cambridge Univ. Press.

Miller, J. H., and Andreoni, J. (1991). "Can Evolutionary Dynamics Explain Free Riding in Experiments?," *Econ. Lett.* **36,** 9–15.

Norman, M. F. (1972). *Markov Processes and Learning Models.* New York: Academic Press.

Pontryagin, L., Andronov, A., and Vitt, A. (1989). "On the Statistical Treatment of Dynamical Systems," in *Noise in Nonlinear Dynamical Systems* (F. Moss and P. McClintock, Eds.). Cambridge, MA: Cambridge Univ. Press.

Revuz, D., and Yor, M. (1991). *Continuous Martingales and Brownian Motion.* Berlin: Springer-Verlag.

Risken, H. (1984). *The Fokker–Planck Equation.* Berlin: Springer-Verlag.

Roth, A., and Erev, I. (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games Econ. Behav.* **8,** 164–212.

Samuelson, L., and Zhang, J. (1992). "Evolutionary Stability in Asymmetric Games," *J. Econ. Theory* **57,** 363–391.

Seymour, R. (1993). "Continuous-Time Models of Evolutionary Games: Populations of Fixed Size," mimeo. University College London.

Taylor, P., and Jonker, L. (1978). "Evolutionarily Stable Strategies and Game Dynamics," *Math. Biosci.* **40,** 145–156.

van Damme, E. (1991). *Stability and Perfection of Nash Equilibria.* Berlin: Springer-Verlag.

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