Evolution and Equilibrium under Inexact Information[*]

William H. Sandholm
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706
whs@ssc.wisc.edu
www.ssc.wisc.edu/~whs

First Version: March 27, 1999
This Version: March 27, 2002

Abstract

We study a general model of stochastic evolution in games, assuming that players have inexact information about the game's payoffs or the population state. We show that when the population is large, its behavior over finite time spans follows an almost deterministic trajectory. While this result provides a useful description of disequilibrium behavior adjustment, it tells us little about equilibrium play.

We establish that the equilibrium behavior of a large population can be approximated by a diffusion. We then propose a new notion of stability called *local probabilistic stability* (LPS), which requires that a population which begins play in equilibrium settle into a fixed stochastic pattern around the equilibrium. We use the diffusion approximation to prove a simple characterization of LPS. While LPS accords closely with standard deterministic notions of stability at interior equilibria, it is significantly less demanding at boundary equilibria.

JEL Classification Numbers: C72, C73.

1. Introduction

The study of evolution in games rests on three basic principles. First, players are myopic, basing decisions about how to play the game on the current population state. Second, aggregate behavior exhibits inertia: the population state changes gradually over time. Third, population sizes are large, rendering individual players anonymous. These three assumptions hold in most contexts in which evolutionary models are applied. They are also mutually reinforcing: for example, behaving myopically is most sensible when opponents' behavior adjusts slowly and when one is sufficiently anonymous that possible repeated game effects can be ignored.

A fourth principle which seems natural to add to this list is that of inexact information. In most settings in which evolutionary models are appropriate, it is reasonable to expect players' knowledge about either the game or their opponents' behavior to be somewhat hazy. Indeed, imprecise knowledge seems most consistent with the other assumptions of the evolutionary model. If the population size is large, exact information about the population state may be difficult to obtain; if players make costly efforts to gather such information, it seems incongruous to then assume that they act upon it in a shortsighted fashion.

In this paper, we study a general model of stochastic evolution in games with large, finite populations, examining the evolution of behavior over finite time spans. Our main modeling restriction requires that players have inexact information about their strategic environment. We begin by showing that the population's behavior can approximated by a deterministic trajectory. While this result gives a precise description of disequilibrium behavior adjustment, it provides little information about equilibrium play.

This observation motivates the central results of the paper, which characterize equilibrium behavior under inexact information. We first prove that equilibrium behavior can be described by a diffusion. We then define new notion of evolutionary stability called *local probabilistic stability* (*LPS*), which requires that a large population which begins play in equilibrium settle into a fixed stochastic pattern around the equilibrium. We use the diffusion approximation to prove a simple characterization of local probabilistic stability. Local probabilistic stability of interior equilibria is closely related to stability under the deterministic dynamics. However, stability of boundary equilibria is less demanding, and can be characterized in terms of robustness of the deterministic dynamics to perturbations which do not leave the boundary.

Models of evolution in games can be split into two classes according to how the populations of players are described. Most research has focused on models with continuous populations of players. In such models, evolution is described directly in terms of a population-level law of motion; studying evolution means characterizing solutions to certain differential equations. Analysis of these models is relatively simple. However, the continuous populations are intended as approximations of finite populations, and the restrictions on aggregate behavior stand in for an explicit specification of individual behavior. It is therefore natural to ask how behavior in the continuous population models is related to behavior in the discrete population models for which they serve as a proxy.

Discrete population models are built up from descriptions of how individual agents behave. While such primitives are obviously desirable, proceeding from them carries a cost: discrete models of evolution can be considerably more difficult to analyze than their continuous population counterparts. Consequently, much of the work on these models has been restricted to very simple cases, most often the single population, two strategy case.

We consider evolution in large but finite populations and give approximate characterizations of behavior in terms of continuous state systems – ordinary differential equations and diffusions. We can therefore both specify our model in terms of individual behavior and characterize evolution in terms of relatively simple continuous state processes. In addition, our results suggest ways of interpreting the continuous population models used throughout the evolutionary literature.

In our model, a finite population of players repeatedly plays a game. Players occasionally receive opportunities to revise their behavior. A player who receives an opportunity decides how to act using a *decision procedure*, which for each population state specifies probabilities of switching between strategy pairs. We illustrate through examples that the decision procedure can embody optimizing, imitative, experimental, or other sorts of choice criteria.

The only restriction we place on the decision procedures is that they reflect an absence of exact information. Formally, we require that the probabilities with which the decision procedure offers its various recommendations change continuously in the population state. When players optimize, continuity can reflect uncertainty about opponent's behavior, or noise in the underlying payoffs. When players imitate, it can also reflect randomness in the choice of whom to mimic. Regardless

of its source, the continuity of the decision procedures captures the idea that small changes in aggregate behavior should not lead to large changes in players' responses.

A population size, a decision procedure and an initial population state define a Markov behavior process. While this evolutionary process is stochastic, we are able to show that when the population size is large, behavior adjusts in a nearly deterministic fashion. We associate with each decision procedure a *deterministic law of motion*, which is a vector field derived from the expected motion of the behavior processes. We establish that over any finite time span, the behavior of a large enough population is closely approximated by a solution to the differential equation defined by the deterministic law of motion.

Why should this be so? When the population size is large, any individual change in behavior has a small effect on the population state. Many revision opportunities pass without the transition probabilities changing significantly. Intuition based on the law of large numbers therefore suggests that the actual course of evolution should be largely determined by its expected direction of motion. Our result confirms the accuracy of the deterministic description of behavior over finite time spans.

This deterministic approximation provides a clear description of behavior away from equilibrium, where by an equilibrium we mean a rest point of the deterministic law of motion.[1] Unfortunately, this result does not enable us to determine which equilibria we should expect to persist. Away from rest points, the idiosyncratic noise from individual players' choices is inconsequential compared to the population's expected motion, which therefore governs its behavior. But at rest points, expected motion is zero; near rest points it is close to zero. Since solutions to differential equations are continuous in their initial conditions, it follows that if we fix the time span of interest in advance, a population that begins play close enough to any rest point will not stray far from the rest point during the span. Thus, the deterministic approximation tells us little about equilibrium behavior.

That expected motion is almost absent near rest points does not imply that a population near a rest point is in complete stasis. Rest points of the deterministic dynamics are points at which the *expected* flows of players between strategy are balanced. Since information is inexact, the *actual* flows between strategies are stochastic, and can occur at strictly positive levels. Because the deterministic approximation eliminates all but the expected changes in the use of each strategy, it

---

[1]   The connections between these rest points and the Nash equilibria of the underlying game depends on the players' decision rule. For examples, see Sections 2, 3.4, and 7.2.

renders these latter properties invisible. To understand equilibrium behavior, we must keep this behavioral flux in full view.

We accomplish this by defining the *local behavior process*, which magnifies deviations from the equilibrium by the square root of the population size. By viewing the population on this finer scale, we are able to perform a limit analysis which leaves the random variations in the population's behavior intact. We prove that over any finite time span, the local behavior process of a large population is closely approximated by a diffusion.

The magnification used to define the local behavior process is essential for proving convergence to a diffusion. However, this convergence result is only of interest if it provides us with information about the original behavior process, which describes the proportions of players choosing each strategy. Fortunately, if we rescale the limit diffusion in an appropriate way, we obtain a new diffusion which closely approximates the original behavior process near the equilibrium. Doing so enables us to make clear probability statements about this original process, and so reveals information about equilibrium behavior which is hidden when only the deterministic approximation is used.[2]

For this reason, the diffusion approximation provides a useful tool for the analysis of equilibrium stability. Since inexact information generally prevents the population's behavior from ever completely settling down, the right definition of stability must account for random variation of behavior around the equilibrium point. We call an equilibrium *locally probabilistically stable* (*LPS*) if a large population which begins play at the equilibrium settles into a fixed probability distribution around the equilibrium.

We use the diffusion approximation to establish a simple characterization of local probabilistic stability. We find that for generic interior equilibria, local probabilistic stability is equivalent to local stability under the deterministic dynamics.[3] Were we directly concerned with deterministic stability, we would test for it by examining the linearization of the dynamics around the equilibrium point; an equilibrium is stable if this linearized system is a contraction. To prove our characterization, we show that the drift coefficient of our diffusion is given by this same linearized system. When the equilibrium is in the interior of the state space, this observation is enough to connect the two forms of stability.

---

[2] For examples of this approach to describing equilibrium behavior, see Section 2.

[3] Of course, while the characterizations of these two notions of stability are nearly identical, the meanings of stability are quite different.

It is often desirable to specify decision procedures which reflect not only inexact information, but also *mutation*: occasional arbitrary behavior. In many contexts, mutation is prevalent enough that it appears in the limiting deterministic dynamics as a force pushing away from the boundaries. In such cases, all equilibria must be interior, and so the results described above completely characterize LPS.

However, if arbitrary behavior is quite rare, it may be more natural to suppose that the limiting dynamics can exhibit boundary equilibria. We show how mutations consistent with boundary equilibria can be introduced to our model, and then offer a characterization of local probabilistic stability for such equilibria. At boundary equilibria, LPS can be characterized in terms of the robustness of the deterministic dynamics to perturbations which remain on the boundary. That mutants playing an unused strategy would disrupt the deterministic system has no bearing on whether the equilibrium is locally probabilistically stable: a population playing such an equilibrium can still stay in a fixed distribution around the equilibrium for a long period of time.

Why don't movements into the interior of the state space matter? Consider an equilibrium at which strategy *i* is not used. When the population is at the equilibrium, the expected change in the number of players choosing strategy *i* is by definition zero. Since no one is using strategy *i* at the equilibrium, the number playing *i* cannot fall. The equilibrium condition then implies that it also cannot rise. These statements remain approximately correct in a neighborhood of the equilibrium. We can therefore show that the presence of strategy *i* always remains negligible. Consequently, even if the appearance of enough players choosing strategy *i* would cause the population to pull away from the equilibrium, random variations in behavior do not introduce enough players choosing strategy *i* to enable the population to leave.

To our knowledge, this paper is the first to characterize equilibrium behavior in an evolutionary model using a diffusion. However, a number of authors have proved special cases of our deterministic approximation result. Boylan (1995) shows how evolutionary processes based on random matching schemes converge to deterministic trajectories when the population size grows large. Binmore, Samuelson, and Vaughan (1995), Börgers and Sarin (1997), and Schlag (1998) consider particular models of evolution which converge to the replicator dynamics. Binmore and Samuelson (1999) prove a deterministic approximation result for discrete time models of evolution under a somewhat restrictive timing assumption. We apply an approximation result due to Kurtz (1970) to prove convergence to a

deterministic trajectory in a quite general model of evolution in games, and establish that inexact information is a sufficient condition for a deterministic approximation to be valid.[4]

Beggs (2002) considers a finite population evolutionary model in which selection pressures between strategies are weak. He shows that if the rate at which time passes is increased as larger populations are considered, a global diffusion approximation becomes possible. In contrast, we allow selection pressures to be weak or strong, and we fix the rate at which time passes independently of the population size. We show that when this rate is fixed, the appropriate global description of behavior is deterministic, but that diffusions can be used to study local behavior near equilibria.[5]

Our definition of local probabilistic stability depends on using a particular order of limits: the population size is taken to infinity first, followed by the time horizon.[6] Because the time horizon is held fixed while the population size limit is taken, LPS directly concerns behavior over finite time spans. As Binmore, Samuelson, and Vaughan (1995) have shown, one typically obtains quite different results if one considers the reverse order of limits, which focuses attention on infinite horizon behavior.[7] In undertaking a finite horizon analysis, we follow Binmore, Samuelson, and Vaughan (1995, p. 10-11), who argue that such analyses are more appropriate than infinite horizon analyses in many economic applications.[8]

Foster and Young (1990), Fudenberg and Harris (1992), and Cabrales (2000) model evolution in continuous populations using stochastic differential equations and

---

[4] Recently, Benaïm and Weibull (2003) have established an exponential bound on the probability of deviations from the deterministic trajectory.

[5] Corradi and Sarin (2000) offer a global diffusion approximation of behavior in a non-game-theoretic model of imitation which is quite different from the model studied here. In each period of their model, the entire population is replaced by newcomers, whose actions are determined through a process based on independent draws of pairs from the incumbent population. Similarly to Beggs (2002), they increase the rate at which these turnovers take place as they consider larger populations. In contrast, our model describes occasional strategy revisions by a fixed group of players in a game-theoretic setting, and the rate of revision is held fixed as larger population sizes are considered.

[6] However, our diffusion approximation only requires us to take the population size limit, and can be established over any finite time interval.

[7] In particular, taking the time horizon limit first leads one to consider equilibrium selection results, as studied by Foster and Young (1990), Kandori, Mailath, and Rob (1993), and Young (1993).

[8] In addition, Binmore, Samuelson, and Vaughan (1995) informally derive a global diffusion approximation with a vanishing noise term for their model of evolution, and they show that this approximation need not yield accurate descriptions of infinite horizon play. We formally establish a local diffusion approximation for our model, but this approximation too is only valid over finite time spans. We will return to the question of infinite horizon analysis and address the possibility of analyses over intermediate time horizons in the final sections of the paper.

characterize infinite horizon behavior. These authors assume directly that evolution is described by a diffusion rather than deriving this property from a more basic model. In addition, while in our model diffusions are only used to study behavior near equilibria, these authors use diffusions as a global description of behavior.

This difference arises because unlike us, these authors study evolution in the presence of aggregate shocks. For example, in Fudenberg and Harris (1992), payoffs contain a noise term which follows a Brownian motion. Hence, the influence of randomness on payoffs is correlated over time, generating aggregate disturbances in the evolutionary process. In contrast, we assume that conditional on the population state, payoffs and other random elements of the decision procedures are realized independently over time. Since in our model the noise influencing the evolutionary process is idiosyncratic, it becomes small when we consider how the proportions of players choosing each strategy evolve over time. We leave the study of finite population models with aggregate payoff noise for future research.

Section 2 introduces our results through two examples. Section 3 contains our formal model. Section 4 establishes the deterministic approximation, and Section 5 the diffusion approximation. Section 6 defines local probabilistic stability and characterizes stability of interior equilibria. Section 7 discusses mutations and characterizes stability of boundary equilibria. Section 8 provides further interpretations of local probabilistic stability. Concluding remarks, including further comments on related literature, are offered in Section 9. Proofs are relegated to the Appendix.

2. Examples

In the examples in this section, two populations of N players are repeatedly randomly matched to play a normal form game. Occasionally, players receive opportunities to change their behavior; each player's revision opportunities arrive via independent, rate 1 Poisson processes.

In the first example, players are repeatedly matched to play the game of matching pennies in Figure 1. In this example, we assume that when a player receives a revision opportunity, he learns the current strategy of a single opponent and plays a best response to that strategy. Let r represent the proportion of players in the column population who are playing strategy R, and let u represent the proportion of players in the row population playing strategy U.

We describe the players' decision procedure by the function d, where $d(s, s')$ is the probability that a player who receives a revision opportunity and is currently playing strategy s will switch to strategy s'. A column player currently choosing strategy R will switch to strategy L if the opponent he queries plays strategy U; hence, $d(R, L) = u$. Given the payoff matrix in Figure 1 and the verbal description of the decision procedure, it is easy to see that the function d is given by

$$d(U, D) = 1 - r \qquad\qquad d(L, R) = 1 - u;$$
$$d(D, U) = r; \qquad\qquad d(R, L) = u.$$

Since the probability that a column player is given the next revision opportunity is $\frac{1}{2}$, and since each player makes up fraction $\frac{1}{N}$ of his population, the expected change in the proportion of players choosing strategy R during a single revision opportunity is

$$\tfrac{1}{2N}((1 - r)d(L, R) - r\, d(R, L)) = \tfrac{1}{2N}((1 - r)(1 - u) - r\, u) = \tfrac{1}{2N}(1 - r - u).$$

Similarly, the expected change in the proportion of players choosing strategy U is

$$\tfrac{1}{2N}((1 - u)d(D, U) - u\, d(U, D)) = \tfrac{1}{2N}((1 - u)r - u\,(1 - r)) = \tfrac{1}{2N}(r - u).$$

Because the players' Poisson processes are independent, revision opportunities in the society as a whole arrive at rate $2N$. Thus, the expected increment per time unit is given by

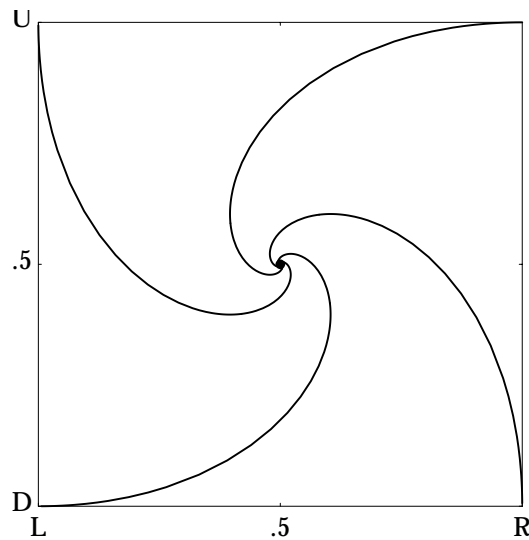|	L	R
U	-1,1	1,-1
D	1,-1	-1,1

Figure 1: Matching pennies

Figure 2: Evolution in matching pennies

$$f(r,\ u) = \begin{pmatrix} 1 - r - u \\ r - u \end{pmatrix}.$$

We call f the *deterministic law of motion* associated with decision procedure d.

In Theorem 4.1, we establish that for sufficiently large population sizes, the evolution of behavior over any finite time span is described by the deterministic law of motion. In this example, behavior over finite time spans is arbitrarily well approximated by solutions to the differential equation

$$\dot{x} = f(x),$$

where $x = (r,\ u)$. Some solutions to this equation are graphed in Figure 2. All solutions converge to the rest point $x^* = (\frac{1}{2},\ \frac{1}{2})$, which is also the unique Nash

	L	R
U	0,0	1,1
D	1,1	0,0

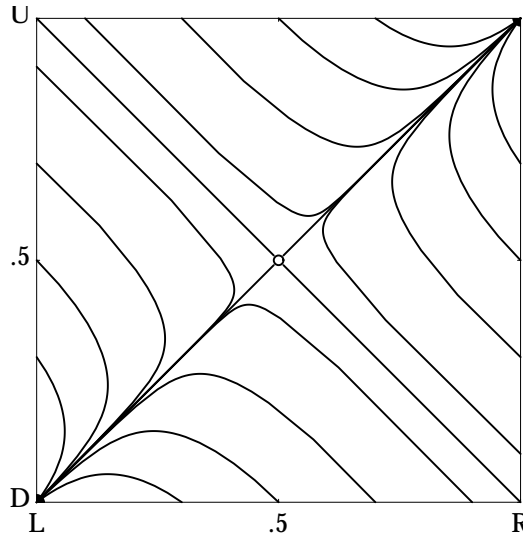Figure 3: A coordination game

Figure 4: Evolution in a coordination game

equilibrium of this game. Hence, regardless of its initial state, the behavior of a large population will quickly come to approximate x^*.[9]   Next, we consider the evolution of play in the coordination game in Figure 3. This time, we assume that when a player receives a revision opportunity, he learns the behavior of three players in the opposing population and plays a best response to this sample. If sampling is performed with replacement, the decision procedure is described by

$$d(U, D) = (1 - r)^3 + 3r(1 - r)^2; \qquad d(L, R) = u^3 + 3u^2(1 - u);$$
$$d(D, U) = r^3 + 3r^2(1 - r); \qquad d(R, L) = (1 - u)^3 + 3u(1 - u)^2.$$

---

[9]   It is worth noting that the decision procedure we have specified is not sensitive to the payoffs of the underlying game. Indeed, *any* payoffs with a counter-clockwise best response structure yield the same choice probabilities.   Hence, behavior converges to $x^* = (1/2, 1/2)$ regardless of the Nash equilibrium of the underlying game. However, as we increase the size of the samples drawn by the players, the limit point of the dynamics approaches the Nash equilibrium of the underlying game. Further discussion of this decision procedure can be found in Section 3.4.1.   For some surprising consequences of small sample sizes, see Sandholm (2001).

The law of motion associated with this decision procedure is therefore

$$g(r,\, u) = \begin{pmatrix} -r + 3u^2 - 2u^3 \\ -u + 3r^2 - 2r^3 \end{pmatrix}.$$

Trajectories from a variety of initial conditions are graphed in Figure 4. If the initial condition $(r_0,\, u_0)$ satisfies $r_0 + u_0 < 1$, the solution trajectory converges to the equilibrium $(D,\, L)$; if the initial condition satisfies $r_0 + u_0 > 1$, play converges to the equilibrium $(U,\, R)$. Thus, from most initial conditions, stochastic evolution leads the population to one of the pure equilibria of the game.

What happens in these games if play begins at the mixed equilibrium $x^* = (\tfrac{1}{2},\, \tfrac{1}{2})$? Since x^* is a rest point of both f and g, the solutions of both differential equations starting from x^* are degenerate. Moreover, if we fix a time T, the continuity of solutions to differential equations in their initial conditions implies that a population which begins play close enough to x^* will remain close to x^* through time T.

Thus, while the deterministic approximation tells us a great deal about the finite horizon behavior of populations which begin play out of equilibrium, it tells us little about populations which begin play in equilibrium. The deterministic approximation relies on the fact that when the population size is large, idiosyncratic sampling noise is averaged away, leaving only the expected motion of the system. Since rest points of the limiting system are points where expected motion is zero, the deterministic approximation suggests that very little happens at these points.

Of course, the rest point x^* is not a point at which the population's behavior ceases to evolve; it is simply a point where the expected flows of players between strategies cancel one another out. But since each player has very limited information about the population state, there is actually considerable idiosyncratic variation in the players' behavior. Near rest points, where expected motion is insignificant, this variation is the most prominent feature of play. To understand equilibrium behavior under inexact information, we must capture these stochastic aspects of play.

Let the *behavior process* $\{X_t^N\}_{t\geq 0} = \{(R_t^N,\, U_t^N)\}_{t\geq 0}$ describe the proportions of players choosing strategies R and U, and consider initial conditions $X_0^N = x_0^N$ which converge to x^* at rate $o(\tfrac{1}{\sqrt{N}})$.[10] We define the *local behavior process at x^**, $\{Z_t^N\}_{t\geq 0}$, by

---

[10] Since the state space for the behavior process $\{X_t^N\}_{t\geq 0}$ is a discrete grid, the initial conditions X_0^N generally cannot be identical to the limit rest point x^*.

$$Z_t^N \equiv \sqrt{N}\left(X_t^N - x^*\right).$$

The local behavior process magnifies the original behavior process by a factor of $\sqrt{N}$, enabling us to perform a finer analysis of behavior near the equilibrium.

Rescaling by $\sqrt{N}$ is helpful because it allows us to obtain a limiting characterization of equilibrium play. In Theorem 5.1, we show that if the population size is large enough, the local behavior process is closely approximated by a diffusion. The drift coefficient of this diffusion is described in terms of the derivative of the law of motion at the equilibrium x^*, which we denote $Df(x^*)$. In deterministic models, this derivative is used to characterize the behavior of trajectories starting near an equilibrium; in our stochastic model, we use this derivative to characterize equilibrium behavior itself.

In the matching pennies game, the local behavior process is approximated by the solution to

$$dZ_t = Df(x^*)Z_t\, dt + \tfrac{1}{\sqrt{2}}I\, dB_t = \begin{pmatrix} -1 & -1 \\ 1 & -1 \end{pmatrix} Z_t\, dt + \begin{pmatrix} \frac{1}{\sqrt{2}} & 0 \\ 0 & \frac{1}{\sqrt{2}} \end{pmatrix} dB_t$$

with initial condition $Z_0 \equiv 0$.[11] We call the solution to this stochastic differential equation the *local limit process at* x^*. The eigenvalues of $Df(x^*)$, $-1 \pm i$, both have negative real part, so the law of motion of Z_t is a contraction perturbed by a white noise process. This process usually moves towards the origin, but the noise term prevents it from ever settling down. By solving the stochastic differential equation, we can explicitly describe the local limit process: it is a zero-mean Gaussian process whose covariance matrix at time T is[12]

$$Cov(Z_T) = \begin{pmatrix} \tfrac{1}{4}(1 - e^{-2T}) & 0 \\ 0 & \tfrac{1}{4}(1 - e^{-2T}) \end{pmatrix} \rightarrow \begin{pmatrix} \tfrac{1}{4} & 0 \\ 0 & \tfrac{1}{4} \end{pmatrix}.$$

The local limit process is important because of what it tells us about the original behavior process X_t^N, which describes the actual proportions of players choosing each strategy. In particular, the random variable $X_T^N = x^* + \tfrac{1}{\sqrt{N}}Z_T^N \approx x^* + \tfrac{1}{\sqrt{N}}Z_T$ must be approximately normally distributed with mean $E(X_T^N) \approx x^*$ and covariance

---

[11] In this section, B_t represents a two-dimensional Brownian motion.
[12] We use lowercase time subscripts to refer to entire processes and uppercase time subscripts to refer to a process at a particular moment in time.

$$Cov(X_T^N) \approx \begin{pmatrix} \frac{1}{4N}(1-e^{-2T}) & 0 \\ 0 & \frac{1}{4N}(1-e^{-2T}) \end{pmatrix} \to \begin{pmatrix} \frac{1}{4N} & 0 \\ 0 & \frac{1}{4N} \end{pmatrix}.$$

Thus, thus, the behavior of a population which begins play at the mixed equilibrium x^* is immediately described by a normal distribution around x^*; the covariance of this distribution converges exponentially quickly to $\frac{1}{4N}I$ as time passes. When a large population which begins at an equilibrium is quickly described by some fixed distribution about the equilibrium, we call the equilibrium *locally probabilistically stable* (LPS).

Formally, a rest point is locally probabilistically stable if there is a zero-mean random variable Z_∞ such that

$$Z_\infty = \lim_{T \to \infty} \lim_{N \to \infty} Z_T^N,$$

where the limits are in distribution. As we noted in the Introduction, taking the time limit last focuses attention on behavior over a long but finite time span. When this limit exists, a population which begins play at an equilibrium settles into a fixed distribution around the equilibrium over this time span. When described on the scale of the original behavior process $X_t^N = x^* + \frac{1}{\sqrt{N}}Z_t^N$, the standard deviations of the limit distribution are of order $\frac{1}{\sqrt{N}}$.

The analysis above shows that in the matching pennies game, x^* is LPS. In contrast, the local behavior process for the mixed equilibrium of the coordination game is approximated by the solution to

$$dZ_t = Dg(x^*)Z_t \, dt + \frac{1}{\sqrt{2}}I \, dB_t = \begin{pmatrix} -1 & \frac{3}{2} \\ \frac{3}{2} & -1 \end{pmatrix} Z_t \, dt + \begin{pmatrix} \frac{1}{\sqrt{2}} & 0 \\ 0 & \frac{1}{\sqrt{2}} \end{pmatrix} dB_t$$

starting from $Z_0 \equiv 0$. The eigenvalues of $Dg(x^*)$ are $\frac{1}{2}$ and $-\frac{5}{2}$, so this stochastic differential equation has one expanding direction (along the 45° line) and one contracting direction (the orthogonal direction). The local limit process Z_t is again a zero-mean Gaussian process, this time with time T covariance matrix

$$Cov(Z_T) = \begin{pmatrix} \frac{1}{4}(e^T-1)+\frac{1}{20}(1-e^{-5T}) & \frac{1}{4}(e^T-1)-\frac{1}{20}(1-e^{-5T}) \\ \frac{1}{4}(e^T-1)-\frac{1}{20}(1-e^{-5T}) & \frac{1}{4}(e^T-1)+\frac{1}{20}(1-e^{-5T}) \end{pmatrix}.$$

The population's behavior is again described by normal distributions about the equilibrium. But as time passes, the correlation between the components of Z_T rapidly approaches 1, while the variances of the components of Z_T grow without bound. In other words, the distribution of behavior is stretched along the 45° line as the population heads towards one of the two pure equilibria. Consequently, the original behavior process X_t^N does not settle into a fixed distribution about the equilibrium x^*, and so the equilibrium x^* is not LPS.

In these examples, local probabilistic stability agreed with local stability under the deterministic dynamics. While this connection holds generically when we consider interior equilibria, at boundary equilibria the connection is broken: on the boundary, deterministic stability is more demanding than local probabilistic stability. We present a formal statement of this claim and an example in Section 7. We also have not addressed what local probabilistic stability tells us about whether the process X_t^N will escape the vicinity of an equilibrium. We take up this question in Section 8.

3. The Model

3.1 The Underlying Game

We consider the evolution of behavior in games played by $r \geq 1$ populations of players. For notational convenience, we assume that members of each population p can choose among n strategies. We let S_p denote the strategy set for population p, and let $S = \bigcup_p S_p$ denote the union of the strategy sets of each population. Let $\Delta_p \subset \mathbf{R}^n$ represent the simplex, so that elements of Δ_p are possible distributions of strategies in population p. The set $\Delta = \prod_p \Delta_p \subset \mathbf{R}^m$ contains all possible strategy distributions in the society as a whole.

We consider the evolution of behavior in large, finite populations. For notational convenience, we assume that each population has N members. If each player chooses a pure strategy, the set of possible strategy distributions is given by $\Delta^N = \{x \in \Delta: Nx_i \in \mathbf{Z} \text{ for all } i \in S\}$.

Each player's payoffs are represented by a random variable which depends on the player's strategy and the population state. We explicitly include payoff randomness to model settings in which players' decisions depend directly on payoff

realizations.[13]   Formally, for each $i \in S$ and $x \in \Delta^N$, the random variable $\pi_i(x)$ represents the payoffs to a player choosing strategy i when the population state is x. Payoffs are Markov, only depending on the past through the current state, and the payoffs received by different players during a single period are independent of one another.[14]   Finally, $\bar{\pi}_i(x)$ denotes expected payoffs.

3.2 Selection Mechanisms and Decision Procedures

Each evolutionary process can be characterized in terms of two components: a *selection mechanism*, which determines the times at which each player considers changing strategies, and a *decision procedure*, which specifies how players respond to such opportunities. We consider each in turn.

We find it convenient to model evolution in continuous time using *Poisson selection*. Under Poisson selection, all players' revision opportunities arrive via independent, rate 1 Poisson processes. Hence, a unit of time in our model is defined as the expected interval between a single player's revision opportunities; this unit does not change when we consider populations of different sizes. When there are r populations of size N, revision opportunities for the population as a whole follow a Poisson process with parameter rN, and each opportunity is equally likely to go to any player.[15]

Decision procedures provide the link between the game's payoffs and the players' behavior. Suppose that the population state is $x \in \Delta^N$ and that a player currently choosing strategy i receives a revision opportunity. Then $d^N(i, j, x)$ is the

---

[13]   For example, if players are randomly matched, it is often desirable to let their decisions depend on the payoffs they actually receive in their matches rather than their expected payoffs *ex ante*.

[14]   The latter assumption does not hold in all potential applications. For example, if we study players who are randomly matched, then for any finite population size, the payoffs received by different members of the same population are not quite independent, as conditioning on the matching of one player slightly alters the match probabilities of the others. Moreover, if there is only a single population and players are not matched against themselves, the payoff distribution will depend on the population size in a vanishing way. Fortunately, our model of evolution will permit transition probabilities which depend in a vanishing way on the population size, so explicitly including these finite population effects would not alter our results.

[15]   Versions of our results also hold in a discrete time version of our model. In this case, we assume that a new period begins every $(1/rN)$ time units. For our results to continue to hold, it is enough to assume that the number of players who receive revision opportunities during each period is constant. One can also assume that each period's revision opportunities are allocated via an i.i.d. process; in this case, we require that the probability p_N that any particular player receives an opportunity is such that Np_N converges as N approaches infinity. In either case, the expected number of revision opportunities that an individual receives during a single time unit is essentially fixed when N is large.

probability that this player switches to strategy j. Clearly, for all populations p, all $i \in S_p$, and all $x \in \Delta^N$, the *decision procedure* $d^N \colon S \times S \times \Delta^N \to [0, 1]$ must satisfy

$$d^N(i, j, x) = 0 \text{ whenever } j \notin S_p;$$

$$\sum_{j \in S_p} d^N(i, j, x) = 1.$$

That is, a player can only choose strategies available to members of his population, and for each strategy and population state, the probabilities of all possible switches sum to one.

A population size N, a decision procedure d^N and an initial condition $x_0^N \in \Delta^N$ define the *behavior process* $\{X_t^N\}_{t \geq 0}$. The behavior process is a pure jump Markov process taking values in Δ^N; the random variable X_T^N describes the society's aggregate behavior at time T. Our goal is to characterize behavior over finite time spans when the population size is large.

Finite population effects (due, for example, to sampling without replacement) can cause the decision procedure to depend on the population size. Fortunately, our results are not sensitive to such dependencies so long as they vanish sufficiently quickly. Formally, we assume that there exists a *limit decision procedure d*: $S \times S \times \Delta \to [0, 1]$ such that

(A1) $\displaystyle \lim_{N \to \infty} \sup_{x \in \Delta^N} \left| d^N(i, j, x) - d(i, j, x) \right| = 0$ for all $i, j \in S$.

That is, as the population size grows large, differences in the choice probabilities vanish uniformly over the set of strategy distributions. This assumption accommodates finite population effects. Moreover, this allowance for slight variations in the decision procedures implies that rare mutations need not affect our analysis; we discuss this point further in Section 7.

We find it reasonable to expect players' choice probabilities not to be unduly sensitive to the current population state. Whether because of payoff noise, imprecise information about opponents' behavior, or other reasons, very close states should not cause very different reactions by the players. To capture this idea formally, we assume

(A2) $d(i, j, x)$ is Lipschitz continuous in x for all $i, j \in S$.

3.3 Deterministic Laws of Motion

We begin our analysis by deriving the transition probabilities of the behavior process X_t^N. At each revision opportunity, a single player considers switching strategies. He either switches from his current strategy $i \in S_p$ to some new strategy $j \in S_p$ or decides to stay with strategy i. Hence, if the current population state is $x \in \Delta^N$ and the population size is N, the only states to which transitions are possible are of the form $x + \frac{1}{N}(\iota_j - \iota_i)$, where ι_i and ι_j are basis vectors in $\mathbf{R}^{rn}$. Since all players are equally likely to be granted the revision opportunity, the probability of a transition from state x to state $x + \frac{1}{N}(\iota_j - \iota_i)$ is given by

$$Q^N(x, x + \tfrac{1}{N}(\iota_j - \iota_i)) = \frac{x_i}{r} \, d^N(i, j, x).$$

We define $I^N : \Delta^N \to \mathbf{R}^{rn}$ to be the expected increment in X_t^N during the next revision opportunity conditional on the current population state:

$$I^N(x) = \sum_y (y - x) \, Q^N(x, y).$$

While this definition is quite compact, it will be more useful to express the expected increments directly in terms of the decision procedures d^N. Consider the expected change in the number of players choosing strategy $i \in S_p$. The probability that the player given the revision opportunity is playing strategy i is $\frac{x_i}{r}$; the probability that the player given the revision opportunity switches to strategy i is

$$\sum_{j \in S_p} \frac{x_j}{r} \, d^N(j, i, x).$$

To determine the expected change in the number of players choosing strategy i, we subtract the former expression from the latter, and then multiply by $\frac{1}{N}$ to account for each player's weight in his population. This yields

$$I_i^N(x) = \frac{1}{rN} \left(\sum_{j \in S_p} \left(x_j \, d^N(j, i, x) \right) - x_i \right).$$

When the population size is N, revision opportunities arrive in the society at rate rN. Hence, the expected increment of X_t^N per time unit, denoted $f^N : \Delta^N \to \mathbf{R}^m$,

is given by $f_i^N(x) = rNI_i^N(x)$. With this motivation, we define the *deterministic law of motion* associated with decision rules d^N, denoted $f: \Delta \to \mathbf{R}^{rn}$, by

$$f_i(x) = \left(\sum_{j \in S_p} x_j\, d(j, i, x) \right) - x_i = \sum_{j \in S_p} \Big(x_j\, d(j, i, x) - x_i\, d(i, j, x) \Big).$$

Assumption (A1) implies that the f^N converge uniformly to f, while assumption (A2) implies that f is Lipschitz continuous. The latter property implies that the differential equation

(D) $\dot{x} = f(x)$

admits a unique solution from every initial condition $x_0 \in \Delta$. We show in Section 4 that the stochastic behavior process X_t^N closely mirrors solutions to this deterministic dynamical system over finite time spans. Before stating this result, we offer two examples of decision procedures and their deterministic laws of motion.

3.4 Examples

For simplicity, our examples involve single population of players; both examples can be generalized to allow for multiple populations. Furthermore, we speak directly in terms of the limit decision procedures; the finite population decision procedures are the same up to a term which vanishes at rate $O(\frac{1}{N})$.

3.4.1 Sample Best Response

Suppose that when a player receives a revision opportunity, he samples the behavior of s members of the population. He then plays a best response to the distribution of players in his sample under the assumption that it is representative of the behavior of the population as a whole. We call this procedure, which was introduced in our examples in Section 2, the *sample best response procedure*.[16]

Let $B: \Delta \to \Delta$ denote the best response correspondence for the expected payoffs $\bar{\pi}$, and let S_x denote a multinomial random variable with parameters s and x. For simplicity, suppose that all possible realizations of the sample induce a unique best

---

[16]   Related procedures are considered by Young (1993) and Kaniovski and Young (1995). In these models, instead of choosing a best response to an incomplete sample of current behavior, players play a best response to an incomplete recollection of the history of play.

response. Then the choice probabilities for the sample best response decision procedure are given by

$$d^s(i, j, x) = P\left(B\left(\frac{S_x}{s}\right) = \iota_j\right).$$

The law of motion associated with this procedure is therefore

$$f_i^s(x) = \left(\sum_j x_j P\left(B\left(\frac{S_x}{s}\right) = \iota_i\right)\right) - x_i$$

$$= P\left(B\left(\frac{S_x}{s}\right) = \iota_i\right) - x_i.$$

Since the distribution of S_x is polynomial in x, so too are $d^s(i, j, x)$ and $f^s(x)$.

In contrast, suppose that the player receiving the revision opportunity is perfectly informed about the current population state. We call the resulting decision procedure the *best response decision procedure*. It is defined by

$$d(i, j, x) = 1_{\{B(x) = \iota_j\}}$$

whenever $x \in \Delta^u = \{x \in \Delta: B(x) \text{ is unique}\}$. This yields the law of motion

$$f(x) = B(x) - x.$$

on Δ^u. This last equation defines the best response dynamics (Gilboa and Matsui (1991)). Since the best response dynamics are discontinuous in the population state, they lie outside the scope of our analysis.

The law of large numbers implies that the sample best response dynamics converge to the best response dynamics as the sample size approaches infinity (although convergence is not uniform). Nevertheless, the two decision procedures lead to very different behavior near equilibria. Continuous dynamics move very slowly near rest points; for this reason our deterministic approximation will provide little information about behavior near equilibria. In contrast, the best response dynamics are discontinuous at equilibria: very small changes in behavior can lead to a sharp acceleration in the evolutionary process. Our assumption of inexact information will preclude this possibility.

3.4.2 Proportional Imitation

The sample best response procedure requires players to know the payoff structure of the game. Since in many settings it is unreasonable to expect players to have such knowledge, it is important to consider procedures which do not require it. Consider this procedure proposed by Schlag (1998). A player who receives a revision opportunity compares his current payoff realization to that of a randomly selected opponent. If his payoff is higher than hers, he continues to play the same strategy; otherwise, he switches to her strategy with a probability proportional to the difference in their payoffs.

This procedure, called *proportional imitation*, is described by

$$d(i, j, x) = x_j E\left[\beta\left(\pi_j(x) - \pi_i(x)\right)\right]^+,$$

where $\beta > 0$ is small enough that the choice probabilities are always between zero and one. The law of motion generated by proportional imitation is

$$\begin{aligned}
f_i(x) &= \sum_j \left(x_j x_i E[\beta(\pi_i(x) - \pi_j(x))]^+ - x_i x_j E[\beta(\pi_j(x) - \pi_i(x))]^+ \right) \\
&= \beta x_i \sum_j x_j E[\pi_i(x) - \pi_j(x)] \\
&= \beta x_i (\bar{\pi}_i(x) - \sum_j x_j \bar{\pi}_j(x)).
\end{aligned}$$

This law of motion is simply the replicator dynamics defined in terms of the expected payoffs of the game.[17]

4. Deterministic Approximation

Our first result establishes that the stochastic behavior process X_t^N can be arbitrarily well approximated by solutions to the differential equation (D). This approximation is valid over any finite time horizon so long as the population size is sufficiently large.

---

[17] The constant β only determines the speed of the evolutionary process.

Theorem 4.1 (*Deterministic Approximation*)

Fix $x_0 \in \Delta$, and let $\{x_t\}_{t \geq 0}$ be the solution to (D) *with initial condition x_0. Suppose that the initial conditions $X_0^N = x_0^N$ converge to x_0. Then for each $T < \infty$ and $\varepsilon > 0$,*

$$\lim_{N \to \infty} P\left(\sup_{t \in [0,T]} \left| X_t^N - x_t \right| < \varepsilon \right) = 1.$$

Theorem 4.1 follows immediately from an approximation result due to Kurtz (1970). The intuition behind the theorem can be explained as follows. At each revision opportunity, the increment in X_t^N is stochastic. However, during any time interval of length δ, the number of revision opportunities we should expect to occur is $\delta r N$, which grows without bound as the population size becomes large. On the other hand, the maximum change in any component of the population state during a single revision opportunity is $\frac{1}{N}$, so the total change in any component of X_t^N during the interval is bounded by δr. Thus, during any sufficiently brief interval, there are a very large number of revision opportunities, each of which generates nearly the same expected increment. Intuition from the law of large numbers suggests that the change in behavior during this interval should be almost completely determined by the expected motion of the system. This expected motion is captured by the differential equation (D).[18]

Theorem 4.1 offers a clear description of the finite horizon behavior of a population which begins play away from equilibrium, where by an equilibrium we mean a rest point of equation (D). However, the theorem tells us little about the behavior of a population which begins play in equilibrium. The theorem says that behavior can be closely approximated by solutions to a differential equation. Such solutions are continuous in their initial conditions: over any fixed time horizon T, a change in the initial conditions smaller than $\delta = \delta(\varepsilon, T)$ will not change behavior at time T by more than ε.[19] Since solutions starting from rest points are degenerate, it follows that solutions starting close enough to rest points move very little over finite time spans.

---

[18]   Binmore and Samuelson (1999) prove a deterministic approximation result in a discrete time framework. They assume that as larger population sizes N are considered, the number of periods which occur per time unit grows faster than order N^2. This guarantees that the occurrence of more than one revision opportunity in a single period becomes extremely unlikely. The discrete time model we describe in footnote 12 does not satisfy this restriction; Kurtz's (1970) results show that it is not needed for a deterministic approximation result to hold.

[19]   For a formal statement and proof, see, e.g., Robinson (1995, Theorem 5.3.3).

Corollary 4.2: *Let x^* be a rest point of* (D), *and fix $\varepsilon > 0$ and $T < \infty$. If the initial conditions $X_0^N = x_0^N$ converge to a point x_0 with $|x_0 - x^*| < \delta = \delta(\varepsilon, T)$, then*

$$\lim_{N \to \infty} P\left(\sup_{t \in [0,T]} |X_t^N - x^*| < \varepsilon \right) = 1.$$

Fix some finite time T. Corollary 4.2 tells us that if a large enough population begins play close enough to an equilibrium, it is quite unlikely to leave the vicinity of the equilibrium through time T.[20]

Theorem 4.1 shows that idiosyncratic noise tends to be drowned out by expected motion when the population size is large. Since at rest points there is no expected motion, the noise which is insignificant elsewhere takes on central importance. To understand equilibrium behavior, we require an analysis which captures this noise explicitly.

5. Diffusion Approximation

The deterministic approximation can be viewed as a law of large numbers for the behavior process. Unfortunately, Corollary 4.2 shows that this result provides limited information about equilibrium behavior. To obtain more information, we might seek a central limit theorem for the behavior process: by magnifying the behavior process by $\sqrt{N}$ about the equilibrium, we might hope to obtain a limit result which captures the random variations in the population's behavior.

We therefore define the *local behavior process at x^** by

$$Z_t^N \equiv \sqrt{N}\left(X_t^N - x^*\right).$$

Theorem 5.1 shows that when the population size is large, the local behavior process Z_t^N is nearly a diffusion. Therefore, if we express the original behavior process X_t^N in terms of Z_t^N, we can use this diffusion approximation to derive

---

[20]   Of course, if an equilibrium of a differential equation is unstable, even solutions which start extremely close to the equilibrium will *eventually* leave the vicinity of the equilibrium. However, we are concerned with behavior over some fixed, finite horizon. Corollary 4.2 says that if we fix the time span of interest in advance, solutions from points very close to the equilibrium will stay nearby during the span. For further discussion, see Section 8.

descriptions of the population's equilibrium behavior. In so doing, we obtain information about equilibrium behavior which is obscured when only the deterministic approximation is considered.

In order to establish the diffusion approximation, we need a somewhat stronger assumption concerning the convergence of the decision procedures: rather than requiring uniform convergence, we need a rate of convergence faster than $\frac{1}{\sqrt{N}}$.

(A3) $\sup\limits_{x\in\Delta^N} \left| d^N(i,j,x) - d(i,j,x) \right| \in o(\frac{1}{\sqrt{N}})$ for all $i, j \in S$.

Since most finite population effects vanish at rate $O(\frac{1}{N})$, this stronger assumption is not unduly restrictive. We also need the limit decision procedure to be continuously differentiable in the population state.

(A4) $d(i, j, \cdot)$ is continuously differentiable for all $i, j \in S$.

To characterize the random aspects of the evolutionary process, we need a measure of the dispersion of its increments. For this reason, we define the *incremental covariance* of the behavior process, $C^N \colon \Delta \to \mathbf{R}^{rn \times rn}$, by

$$C_{ij}^N(x) = \sum_{y} (y_i - x_i)(y_j - x_j) Q^N(x, y)$$

It will prove useful to express this function directly in terms of the decision rules. Note that if $i \neq j$, the only way that the number of players choosing both of these strategies can change in during a single revision opportunity is if the player with the opportunity switches from one strategy to the other. On the other hand, any revision which changes the number of players choosing strategy i contributes positively to the incremental variance of component i. Using this logic, one can show that

$$C_{ij}^N(x) = \begin{cases} -\frac{1}{rN^2}\left(x_i d^N(i,j,x) + x_j d^N(j,i,x) \right) & \text{if } i \neq j; \\ \frac{1}{rN^2}\sum\limits_{l \neq i}\left(x_i d^N(i,l,x) + x_l d^N(l,i,x) \right) & \text{if } i = j. \end{cases}$$

With this motivation, we define the function $a \colon \Delta \to \mathbf{R}^{rn \times rn}$ by

–23–

$$a_{ij}(x) = \begin{cases} -\left(x_i d(i, j, x) + x_j d(j, i, x)\right) & \text{if } i \neq j; \\ \sum_{l \neq i}\left(x_i d(i, l, x) + x_l d(l, i, x)\right) & \text{if } i = j. \end{cases}$$

When x^* is an equilibrium, we call the matrix $a(x^*)$ the *diffusion coefficient at x^**. The reason for this designation will become clear below.

A normalization will make our results easier to state. So far, we have expressed each population's behavior as a point in $\mathbf{R}^n$, where n is the number of strategies available to the population's members. However, since the population state must stay in the simplex Δ_p, it is only free to move in $n - 1$ dimensions. It is therefore convenient to change the coordinates we use to refer to population states. From this point forward, we view each set Δ_p as a subset of $\mathbf{R}^{n-1}$ rather than as a subset of $\mathbf{R}^n$. We accomplish this by identifying each element $(x_1, \ldots, x_{n-1}, x_n)$ in $\mathbf{R}^n$ with its projection $(x_1, \ldots, x_{n-1})$ in $\mathbf{R}^{n-1}$. Similarly, we consider the state space $\Delta = \prod_p \Delta_p$ a subset of $\mathbf{R}^k$, where $k = r(n - 1)$. Fortunately, all of our earlier definitions can still be used after minor modifications which account for this change in coordinates. In particular, $f: \Delta \to \mathbf{R}^k$ and $a: \Delta \to \mathbf{R}^{k \times k}$ are defined as before if we simply leave off all arguments and components corresponding to the nth strategy of each population.

Since we want to characterize behavior near equilibria, it will be useful to have a simple description of expected motion near equilibria. We therefore define the derivative of f, $Df: \Delta \to \mathbf{R}^{k \times k}$, which exists for all $x \in \Delta$ by assumption (A4).[21] We let $D^* = Df(x^*)$ denote the derivative of f at the equilibrium x^*.

The derivative D^* can be used to determine the stability of x^* under the differential equation

(D) $\dot{x} = f(x)$.

Taking a Taylor series of f about x^* reveals that

$$f(x) \approx f(x^*) + Df(x^*)(x - x^*) = D^*(x - x^*)$$

when x is close to x^*. It can therefore be shown that solutions of (D) near x^* are conjugate to solutions of the linear equation

(L) $\dot{y} = D^* y$

---

[21]   To define the derivative at points x on the boundary of Δ, it is enough to consider how f changes as we move from x in directions which stay within Δ.

near $y = 0$. Solutions of the linear equation (L) can be expressed compactly as

$$y_t = e^{D^* t} y_0,$$

where $e^{D^* t} = I + \sum_{k=1}^{\infty} \frac{(D^* t)^k}{k!}$ is a matrix exponential. The linear equation (L) will play a leading role in the analysis to come.

Before stating our result, we introduce a few additional definitions. First, let $a^* = a(x^*)$ denote the diffusion coefficient at x^*. Since a^* is symmetric and positive semidefinite,[22] it has a "square root": we can find a $\sigma^* \in \mathbf{R}^{k \times k}$ such that $\sigma^*(\sigma^*)' = a^*$. Next, let $\{B_t\}_{t \geq 0}$ denote a k-dimensional Brownian motion. Finally, our notion of convergence for the local behavior processes is weak convergence in $D([0, T], \mathbf{R}^k)$, the space of functions from $[0, T]$ to $\mathbf{R}^k$ which are right continuous and have left limits.

Theorem 5.1 (*Diffusion Approximation*)

Let x^* be a rest point of (D), and let the sequence of initial conditions $X_0^N = x_0^N$ converge to x^* at rate $o(\frac{1}{\sqrt{N}})$. Then for each $T < \infty$, Z_t^N converges weakly in $D([0, T], \mathbf{R}^k)$ to Z_t, the solution to

$$\text{(S)} \qquad dZ_t = D^* Z_t \, dt + \sigma^* dB_t$$

from initial condition $Z_0 \equiv 0$. This solution is given by

$$Z_t \equiv \int_0^t e^{D^*(t-s)} \sigma^* \, dB_s.$$

We call the process Z_t the local limit process at x^*.

Diffusions (i.e., solutions to stochastic differential equations) in $\mathbf{R}^k$ are characterized by two coefficients: a drift coefficient $\mu \colon \mathbf{R}^k \to \mathbf{R}^k$, which describes the expected direction of motion, and a diffusion coefficient $\sigma^2 \colon \mathbf{R}^k \to \mathbf{R}^{k \times k}$, which captures the dispersions of and correlations between the components of the process. Theorem 2 tells us that over finite time spans, the local behavior process Z_t^N is closely approximated by the local limit process Z_t, a diffusion with linear drift coefficient $\mu(z) = D^* z$ and constant diffusion coefficient $\sigma^2(z) = a^*$. That is, the distribution over paths through $\mathbf{R}^k$ induced by Z_t^N converges to the distribution

---

[22] We verify this in the Appendix (Lemma A.3).

over paths induced by the diffusion Z_t. Since the stochastic differential equation which describes Z_t is linear, we can solve it explicitly. Indeed, the integral representation of the local limit process Z_t immediately reveals that it is a zero-mean Gaussian process.

The local limit process is important because of what it tells us about the original behavior process X_t^N, which describes the proportions of players choosing each strategy. Fix some finite time T. Since $X_t^N = x^* + \frac{1}{\sqrt{N}} Z_t^N$, we can conclude from Theorem 5.1 that if N is large enough, the process X_t^N is approximately Gaussian through time T. In particular, at each moment $T' \in [0, T]$, $X_{T'}^N$ is approximately normally distributed about the equilibrium x^*, and has an approximate covariance of $\frac{1}{N} Cov(Z_T)$.

To prove Theorem 5.1, we appeal to a convergence theorem due to Stroock and Varadhan (1979). They consider sequences of Markov processes whose increments become vanishingly small. Roughly speaking, their result says that if the expected increments and incremental covariances of a sequence of Markov processes converge to some functions $\mu(\cdot)$ and $\sigma^2(\cdot)$, and if the probability of large increments vanishes, then the Markov processes themselves converge to the diffusion whose drift and diffusion coefficients are $\mu(\cdot)$ and $\sigma^2(\cdot)$. We use this observation to sketch the proof of Theorem 5.1; details can be found in the Appendix.

For simplicity, we assume that all decision rules are identical: $d^N \equiv d$ for all N. We first consider covariances. The deterministic approximation tells us that as the population size grows large, all variance in the original behavior processes X_t^N vanishes. To obtain a sequence of processes whose variances do not vanish, we must rescale the original processes by some exploding term. The analogy with the central limit theorem suggests that the correct term is $\sqrt{N}$.

A calculation verifies this intuition. The covariances of the components of X_t^N during a single revision opportunity are given by C^N. Multiplying this expression by the Poisson rate of rN yields the incremental covariance per time unit, rNC^N. To find the corresponding expression for $Z_t^N \equiv \sqrt{N}(X_t^N - x^*)$, we observe that increments of this process are $\sqrt{N}$ times larger the corresponding increments of X_t^N; this increases the covariance by a factor of $(\sqrt{N})^2 = N$. Thus, the incremental covariance per time unit of Z_t^N at state $z = \sqrt{N}(x - x^*)$ equals

$$rN^2 C^N(x^* + \tfrac{z}{\sqrt{N}}) = a(x^* + \tfrac{z}{\sqrt{N}}).$$

Since $a(\cdot)$ is continuous, this expression converges to $a(x^*)$ as N approaches infinity.

We next consider expected increments. The expected change in X_t^N during a single revision opportunity is $I^N(x) = \frac{1}{rN} f(x)$, so the expected increment per time unit is $f(x)$. Increments of $Z_t^N \equiv \sqrt{N}\left(X_t^N - x^*\right)$ are $\sqrt{N}$ times larger the corresponding increments of X_t^N, so the expected increment per time unit of Z_t^N at state $z = \sqrt{N}\,(x - x^*)$ is

$$\sqrt{N} f(x^* + \tfrac{z}{\sqrt{N}}).$$

The mean value theorem implies that

$$\sqrt{N} f(x^* + \tfrac{z}{\sqrt{N}}) = \sqrt{N}\ (f(x^*) + Df\,(x^* + \lambda \tfrac{z}{\sqrt{N}})\ \tfrac{z}{\sqrt{N}}).$$

for some $\lambda \in [0, 1]$. Were x^* were not a rest point (i.e., were $f(x^*) \neq 0$), the expected increment of Z_t^N would explode, and the diffusion approximation would fail. But since $f(x^*) = 0$, we find that

$$\sqrt{N} f(x^* + \tfrac{z}{\sqrt{N}}) = \sqrt{N}\ Df(x^* + \lambda \tfrac{z}{\sqrt{N}})\ \tfrac{z}{\sqrt{N}} = Df(x^* + \lambda \tfrac{z}{\sqrt{N}})z.$$

As N grows large, this expression converges to $Df(x^*)z = D^*z$, which is therefore the drift coefficient of the limit process.

We conclude that over any finite time span, Z_t^N converges weakly to the solution of the stochastic differential equation (S). The solution to this equation is obtained by introducing the integrating factor e^{-D^*t} and applying Ito's formula.

6. Local Probabilistic Stability

Because players' information is inexact, the flows of players between strategies are random. Even at rest points, where expected motion is absent, there is still considerable stochastic variation in players' behavior. The proper notion of equilibrium stability must account for this variation explicitly.

Theorem 5.1 shows that by examining the local behavior process, we can obtain a limiting description of equilibrium behavior which is independent of the population size. It is therefore natural to state our definition of stability in terms of this process. We call the equilibrium x^* *locally probabilistically stable* (*LPS*) if there is a zero-mean random variable Z_∞ such that

$$\lim_{T \to \infty} \lim_{N \to \infty} Z_T^N = Z_\infty,$$

where the limits are limits in distribution in $\mathbf{R}^k$.

Roughly speaking, LPS requires that when N is large, the random variable Z_T^N has nearly the same distribution as Z_∞. We can restate this requirement in terms of the original behavior process: when N is large, the mean and covariance of $X_T^N = x^*$ $+ \frac{1}{\sqrt{N}} Z_T^N$ are roughly x^* and $\frac{1}{N} Cov(Z_\infty)$. Thus, if an equilibrium is LPS, a population which begins play at that equilibrium settles into a fixed distribution around that equilibrium. The standard deviations of this distribution's components of are of order $\frac{1}{\sqrt{N}}$: the larger the population, the closer it will stay to the point x^*.

That we take the time limit last means that we are considering finite horizon behavior. That we take the time limit at all may seem to suggest that Z_∞ only describes behavior after a long time has passed. Fortunately, we shall see that whenever an equilibrium is LPS, the limit random variable Z_∞ describes behavior almost immediately.

We also offer a definition of instability of equilibrium. We say that the equilibrium x^* is *locally probabilistically unstable* (*LPU*) if

$$\lim_{T \to \infty} \lim_{N \to \infty} \left\| Cov(Z_T^N) \right\| = \infty,$$

where $\|A\| = \max_{i,j} |A_{ij}|$. If an equilibrium is LPU, then a population which starts at the equilibrium moves further and further away as time passes.

The diffusion approximation provides the basis for our characterization of local probabilistic stability. Since the local behavior process Z_t^N converges to the local limit process Z_t, it follows that

$$\lim_{T \to \infty} \lim_{N \to \infty} Z_T^N = \lim_{T \to \infty} Z_T.$$

Hence, stability can be characterized directly in terms of the local limit process. Since this process is a zero-mean Gaussian process, its limit behavior only depends on the limit behavior of its time T covariance matrix, $Cov(Z_T)$. If this matrix converges, then the equilibrium is LPS; if some component of $Cov(Z_T)$ explodes, the equilibrium is LPU.

Theorem 6.1 characterizes the local probabilistic stability of interior equilibria.

Theorem 6.1: *Let x^* be a rest point of* (D), *and suppose that a^* has full rank. Then if all eigenvalues of D^* have strictly negative real part, x^* is LPS; otherwise, it is LPU.*

The condition that the diffusion coefficient a^* has full rank is a requirement that random motions are possible in all directions from x^*. As long as x^* is in the interior of Δ, most decision procedures reflecting inexact information will generate such random variations. Of course, in any particular example it is easy to check the full rank condition directly.[23]

Theorem 6.1 provides a simple way of checking whether an interior equilibrium is LPS. Consider the equilibrium $x^* = (\frac{1}{2}, \frac{1}{2})$ of the examples from Section 2. In each case, the diffusion coefficient is the full rank matrix $a^* = \frac{1}{\sqrt{2}}I$, so we can apply the theorem. In the matching pennies game, the eigenvalues of the derivative D^* are $-1 \pm i$, so x^* is LPS. In the coordination game, the eigenvalues of D^* are $\frac{1}{2}$ and $-\frac{5}{2}$, so x^* is LPU.

The stochastic differential equation (S) which defines the local limit process is the linear equation (L) perturbed by random shocks. If equation (S) is a contraction (i.e., if all eigenvalues of D^* have negative real part), then whenever the process Z_t wanders far from the origin, the deterministic part of equation (S) forces the process back, guaranteeing stability. On the other hand, if D^* has a positive eigenvalue, then solutions to the linear equation (L) move away from the origin at an exponential rate from a most initial conditions; solutions to the stochastic equation (S) explode from *all* initial conditions, including the origin.

More formally, whether x^* is LPS turns on whether $\lim_{T\to\infty}Cov(Z_T)$ exists. Since Z_t is a linear diffusion, it can be shown that

$$Cov(Z_T) = \int_0^T e^{D^*t}a^* e^{(D^*)'t}dt.$$

Recall that e^{D^*t} is the matrix solution to equation (L). If all eigenvalues of D^* have negative real part, all solutions to (L) approach the origin exponentially quickly. Hence, the norm of e^{D^*t} falls at an exponential rate, and $\lim_{T\to\infty}Cov(Z_T)$ exists. Indeed, whenever x^* is LPS, $Cov(Z_T)$ converges to its limit at an exponential rate, so

---

[23] If a^* does not have full rank, then the condition that all eigenvalues of D^* be negative is sufficient but not necessary for x^* to be locally probabilistically stable. Essentially, the necessary condition requires negative eigenvalues for all eigenvectors corresponding to directions in which the population is able to move.

the limit distribution is not only describes behavior after a long time has passed, but in fact describes behavior almost immediately.

If some eigenvalue of D^* has positive real part, some component of e^{D^*t} increases exponentially; hence, $Cov(Z_T)$ diverges, and the equilibrium is LPU. The equilibrium is also LPU when the eigenvalue of D^* with the largest real part has real part zero. In contrast, in the latter case the dynamic stability of x^* as an equilibrium of the deterministic system (D) is indeterminate. At interior equilibria, this is the only way that the two predictions can differ.[24]

Why does this difference arise? When analyzing the dynamic stability of an equilibrium of the deterministic system (D), that D^* has an eigenvalue with real part zero tells us that the linear system (L) is not a good enough approximation to form the basis for stability analysis. This is the source of the indeterminacy. In contrast, when analyzing local probabilistic stability, eigenvalues with real part zero correspond to directions in which movements towards or away from the equilibrium are driven entirely by noise. For example, when $D^* = 0$, the local limit process is a Brownian motion: $Z_t = \sigma^* B_t$. Since $Cov(\sigma^* B_T) = a^* T$, the equilibrium must be locally probabilistically unstable.

7. Boundary Equilibria

7.1 Mutations and Boundary Equilibria

It is often of interest to study decision rules which allow for occasional arbitrary behavior, commonly called *mutation*. Mutation ensures that a strategy which is currently unused can always be chosen at the next revision opportunity. It follows that if rates of mutation are bounded away from zero, mutation will appear in the law of motion *f* as a force leading away from the boundary. This guarantees that all equilibria will lie in the interior of the state space, and hence that Theorem 6.1 is enough to characterize LPS.

On the other hand, in settings where arbitrary behavior is very rare, the right deterministic approximation may be one which admits boundary equilibria. Mutations and boundary equilibria can coexist if we treat the mutation rate the same

---

[24] At first glance, cases in which the real part of some eigenvalue is zero may seem quite rare. But as Binmore and Samuelson (1999) note, this must be true of any rest point which lies in a non-trivial component of rest points; such components are common features of dynamics for extensive form games.

way we treat the population size: as a parameter whose allowable values depend on the precision we demand in our approximations.

As an example, suppose that players who receive revision opportunities usually follow the proportional imitation rule from Section 3.4.2, but occasionally select a strategy at random. The resulting decision procedure is

$$d(i, j, x) = (1 - \eta) \, x_j \, E\Big[\beta\big(\pi_j(x) - \pi_i(x)\big)\Big]^+ + \frac{\eta}{n}.$$

When there are no mutations $(\eta = 0)$, then as we saw earlier, the proportional imitation rule generates the replicator dynamics as its law of motion. This is still true if mutations are possible but sufficiently rare: we can still approximate the stochastic evolutionary process X_t^N by solutions to the replicator dynamics over any finite time span and with any degree of precision, so long as the population size is large enough and the mutation rate is small enough.

To establish this formally, we let the mutation rates η^N associated with the processes X_t^N approach zero as N grows large. The resulting decision procedures d^N converge uniformly to the unperturbed proportional imitation rule, and so satisfy Assumption (A1). Hence, we may apply Theorem 4.1: for pairs (N, η^N) far enough along the sequence, the deterministic approximation holds. The same logic holds for the diffusion approximation, but to ensure that Assumption (A3) is satisfied we must impose a tighter restrictions on the set of parameter pairs (N, η^N) we may consider. These points will be illustrated in an example below.[25]

7.2 Local Probabilistic Stability of Boundary Equilibria

Since many commonly studied dynamics admit boundary equilibria, it is important to be able to check which boundary equilibria are LPS. We therefore offer a simple characterization. Since it is based on the diffusion approximation from Theorem 5.1, this characterization of LPS is valid in the presence of mutations, so long as we consider parameter pairs (N, η^N) which satisfy assumption (A3).

The following lemma is basic to understanding boundary behavior.

---

[25]   Our formal technique for allowing arbitrarily small mutation rates is to allow them to vary with the population size. In doing so, we are not suggesting that players in larger populations are more likely to behave arbitrarily than those in smaller ones. Rather, this is simply a way of specifying the combinations of population sizes and mutation rates for which our approximation results hold.

Lemma 7.1: *If x^* is a rest point of* (D) *with $x_i^* = 0$, then $a_{ij}^* = a_{ji}^* = 0$ for all j.*

Proof: Since x^* is a rest point,

$$f_i(x^*) = \left(\sum_{j \in S_p} x_j^* d(j, i, x^*) \right) - x_i^* = 0.$$

We have assumed that $x_i^* = 0$; therefore, $x_j^* d(j, i, x^*) = 0$ for all j. Since $a_{ij}^* = a_{ji}^* = -\left(x_i^* d(i, j, x^*) + x_j^* d(j, i, x^*) \right)$ when $i \neq j$ and $a_{ii}^* = \sum_{l \neq i} \left(x_i^* d(i, l, x^*) + x_l^* d(l, i, x^*) \right)$, the lemma follows. ■

If x^* is a rest point at which strategy i is not used, then when the population is at state x^*, the expected increment in the number of players choosing strategy i must be zero. But the number choosing this strategy cannot fall, and so cannot rise either. Hence, near x^*, the probability of *any* change in the use of strategy i must be close to zero.

Lemma 7.1 implies that the diffusion coefficient of a boundary equilibrium cannot have full rank. Therefore, Theorem 6.1 cannot be applied to test for local probabilistic stability. The lemma also helps us establish an important property of the local limit process.

Proposition 7.2: *If x^* is a rest point with $x_i^* = 0$, then its local limit process Z_t satisfies $(Z_T)_i \equiv 0$ for all $T \geq 0$.*

Suppose that a large population begins play near a rest point x^* at which strategy i is unused. Lemma 7.1 tells us that the probability that a player switches to or from strategy i during the next revision opportunity is very small. Proposition 7.2 extends this observation over time: even after a long interval has passed, the probability that strategy i is adopted by a significant fraction of the population remains negligible.

If x^* corresponds to a pure strategy profile, Proposition 7.2 implies that very little variation in the use of *any* strategy is observed. Consequently, x^* must be locally probabilistically stable.

Corollary 7.3: *If x^* is rest point which corresponds to a pure strategy profile, the local limit process is the null process*: $Z_T \equiv 0$ *for all $T \geq 0$. Therefore, x^* is LPS.*

Even if the deterministic dynamics lead away from x^*, these dynamics move very slowly at points very close to x^*. Since there is little random variation in behavior, the population never wanders far enough from the equilibrium for the deterministic dynamics to draw the population away.

Corollary 7.3 implies that $Cov(X_T^N) \approx \frac{1}{N} Cov(Z_T) = 0$. Thus, the random variations in behavior observed at interior rest points are absent at rest points on the vertices. Even when players' information is inexact, their behavior near monomorphic rest points is almost completely noise free.

More generally, Proposition 7.2 says that when a population begins play near a boundary equilibrium, strategies outside the support of the equilibrium are never adopted to any significant extent. This suggests that a complete characterization of local probabilistic stability may be possible if we ignore directions of motion corresponding to unused strategies: that is, directions heading away from the boundary.

To prepare for such a result, we recall our convention of only explicitly representing $n - 1$ out of n strategies in each population, so that population states are elements of $\mathbf{R}^k = \mathbf{R}^{r(n-1)}$ rather than $\mathbf{R}^{rn}$. To this we add a new convention: that each strategy which we omit is in the support of the equilibrium. Let $\hat{S} \subset S$ be the set of strategies which are represented explicitly, and let $\hat{S}_u \subset \hat{S}$ be the set of strategies which are unused at x^*: $\hat{S}_u = \{i \in \hat{S}: x_i^* = 0\}$. Denote the cardinality of $\hat{S} - \hat{S}_u$ by κ. By our convention, all unused strategies in S are represented explicitly in $\hat{S}$, and are hence in $\hat{S}_u$.

We include the unused strategies at first so that we can be certain to ignore them in our analysis. We let $r^* \in \mathbf{R}^{\kappa \times \kappa}$ be the *reduced diffusion coefficient*, which is obtained by eliminating all rows and columns of a^* corresponding to unused strategies. By Lemma 7.1, these rows and columns consist entirely of zeros. Similarly, let $R^* \in \mathbf{R}^{\kappa \times \kappa}$ be the *reduced derivative of f at x^**: this is the derivative matrix $D^* = Df(x^*)$ with the rows and columns for all $i \in \hat{S}_u$ deleted.

Our new stability result generalizes Theorem 6.1, showing that the diffusion coefficient a^* and derivative D^* in the statement of that theorem can be replaced by their reduced counterparts.

Theorem 7.4: *Let x^* be a rest point, and suppose that r^* has full rank. Then if all eigenvalues of R^* have strictly negative real part, x^* is LPS; otherwise, it is LPU.*

Theorem 7.4 provides a simple, general method for determining local probabilistic stability. We illustrate this through an example. Consider a single population of players who are repeatedly randomly matched to play the symmetric game in Figure 5. In this game, strategy U is dominant; if this strategy is eliminated, a Hawk-Dove game remains. When a player receives a revision opportunity, he usually follows a simple imitative decision procedure, but occasionally chooses a strategy at random. With probability $(1 - \eta^N)$, the player compares the payoff he received in his previous match to the payoff a randomly chosen opponent received in her previous match. If the opponent received a higher payoff, the player switches to her strategy; otherwise, he stays with his original strategy. With the remaining probability of η^N, the player chooses a strategy arbitrarily.

	U	M	D
U	1,1	1,0	1,0
M	0,1	0,0	1,1
D	0,1	1,1	0,0

Figure 5: A symmetric game

Let $x = (u, m, d)$ represent the population state. The decision procedure above is described by[26]

$$d^N(U, M) = \tfrac{\eta^N}{3}, \qquad\qquad d^N(U, D) = \tfrac{\eta^N}{3};$$
$$d^N(M, U) = (1 - \eta^N)(u + m)u + \tfrac{\eta^N}{3}, \qquad d^N(M, D) = (1 - \eta^N)(u + m)dm + \tfrac{\eta^N}{3};$$
$$d^N(D, U) = (1 - \eta^N)(u + d)u + \tfrac{\eta^N}{3}, \qquad d^N(D, M) = (1 - \eta^N)(u + d)md + \tfrac{\eta^N}{3}.$$

---

[26] For simplicity, we assume that players can be randomly matched against themselves. Assuming that this does not occur would not alter our results.

In the absence of mutations, this decision procedure is equivalent to the proportional imitation rule with $\beta = 1$. Therefore, if $\{\eta^N\}$ is a vanishing sequence, the law of motion induced by the decision procedures d^N is the replicator dynamics:

$$f(x) = \begin{pmatrix} u(1 - \overline{\Pi}(x)) \\ m(d - \overline{\Pi}(x)) \\ d(m - \overline{\Pi}(x)) \end{pmatrix}.$$

(Here, $\overline{\Pi}(x) = u + 2md$ represents the average payoffs in the population when the current state is x.)

Theorem 4.1 links these dynamics to the evolutionary process with mutations. If we fix any time span T and degree of precision ε, then for any parameter pair (N, η^N) far enough along the sequence, the corresponding process X_t^N stays within ε of the solution to the replicator dynamics through time T with probability at least $1 - \varepsilon$.

Some solution trajectories of the replicator dynamics are sketched in Figure 6. The unique symmetric Nash equilibrium $(u, m, d) = (1, 0, 0)$ is an equilibrium under
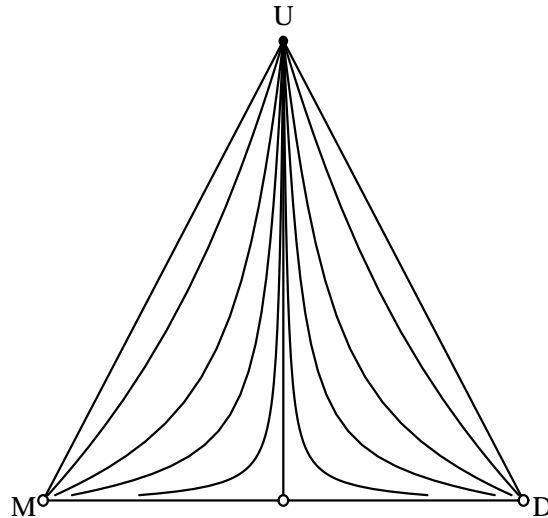

Figure 6: Evolution in a symmetric game

these dynamics, and solutions from all interior initial conditions lead here. The states $(0, \frac{1}{2}, \frac{1}{2})$, $(0, 1, 0)$, and $(0, 0, 1)$ are also equilibria. Now suppose we choose the η^N to vanish at rate $o(\frac{1}{\sqrt{N}})$, thereby focusing on parameter pairs (N, η^N) in which mutation rates are relatively small. If we do this, Assumption (A3) is satisfied, so we can use the diffusion approximation to determine which equilibria are LPS.

Leaving off the redundant component representing strategy D, we let $X_t^N = (U_t^N, M_t^N)$ represent the proportions of players choosing strategies U and M, and consider the local behavior process $Z_t^N \equiv \sqrt{N}\left(X_t^N - x^*\right)$ for the equilibrium $x^* = (0, \frac{1}{2}, \frac{1}{2})$. Applying Theorem 5.1, we find that this process is approximated by the local limit process Z_t, which is the solution to

$$dZ_t = \begin{pmatrix} \frac{1}{2} & 0 \\ -\frac{1}{2} & -\frac{1}{2} \end{pmatrix} Z_t\, dt + \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{\sqrt{8}} \end{pmatrix} dB_t$$

starting from $Z_0 = 0$. The derivative matrix D^* has eigenvalues $-\frac{1}{2}$ and $\frac{1}{2}$, suggesting that x^* might not be stable. However, since strategy U is unused, we remove the row and column of D^* corresponding to this strategy, obtaining the reduced derivative $R^* = -\frac{1}{2}$. Since R^* is negative, x^* is LPS.

Solving the stochastic differential equation, we find that the local limit process Z_t is a zero-mean Gaussian process whose covariance matrix at time T is

$$Cov(Z_T) = \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{8}(1 - e^{-T}) \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{8} \end{pmatrix}.$$

The strategy U component of this process is degenerate at zero, while the strategy M component converges in distribution to a normal distribution with mean zero and variance $\frac{1}{8}$.

We can use this description of the local limit process to make explicit statements about behavior in finite populations under positive mutation rates. Fix a time $T \geq 0$, a constant $\varepsilon > 0$ and a disk $A \subset \mathbf{R}^2$. Then for all parameter pairs (N, η^N) far enough along the sequence, $\left|P(Z_T^N \in A) - P(Z_T \in A)\right| < \varepsilon.$[27] Therefore, even in the presence of mutations, the local behavior process Z_t^N and the original behavior process X_t^N are described by fixed distributions about the equilibrium.

If the population instead begins near any of the pure equilibria, Corollary 7.3 tells us that the local behavior process is approximated by the null process: $Z_t \equiv 0$. Thus, despite what Figure 6 seems to suggest, all four rest points are LPS.[28]

---

[27]   This follows from Theorem 5.1 and the Portmanteau Theorem (Durrett (1996, Theorem 8.1.1)). The result holds for any set $A \subset \mathbf{R}^2$ such that the set $\{m \in \mathbf{R}: (0, m) \in \delta A\}$ has Lebesgue measure zero. Analogous probability statements can also be made for the original behavior process X_t^N.

[28]   It is important to remember that local probabilistic stability is used to analyze game/decision procedure pairs. Decision procedures based on imitation, like the one considered in this example, often generate boundary rest points which are not near Nash equilibria; our analysis suggests that such rest points may be more stable than expected. If our intended application led us to consider a different

8. Interpreting Local Probabilistic Stability

To this point, we have provided a definition of local probabilistic stability in terms of the local behavior process Z_t^N, as well as a characterization of this concept in terms of the derivative matrix of the deterministic law of motion (D). We now consider the following question: what does this notion of stability tell us about whether the original behavior process X_t^N will leave the vicinity of an equilibrium? We first address this question intuitively and then discuss one possible formalization.

To begin, recall the deterministic approximation results from Section 4. Theorem 4.1 tells us that if we fix the time horizon T in advance, then when the population size N is large enough, the behavior process X_t^N will remain close to an appropriate solution to (D) through time T. When the initial condition X_0^N is an equilibrium x^*, Corollary 4.2 shows that X_t^N does not move away from x^* over this time horizon, regardless of the stability of x^* under (D).

Now suppose the process X_t^N begins at an equilibrium x^* which is unstable under (D). If X_t^N were able to take a "first step" away from x^*, leaving the small neighborhood of x^* where expected motion is extremely slow, then stronger deterministic forces would take over, enabling X_t^N to leave the vicinity of x^*. Since x^* is a rest point of (D), this first step from x^* must occur through the combined action of random fluctuations and of the weak deterministic forces existing very close to x^*. It is precisely this combination of forces which is captured by the local behavior process Z_t^N and by local probabilistic stability.

To put this differently (though still somewhat loosely), the deterministic approximation (D) captures stability of equilibrium once X_t^N is outside a small neighborhood of x^*, while the local behavior process Z_t^N captures stability within this small neighborhood. The only possible discrepancy between these analyses occurs when an equilibrium is unstable in the former sense but stable in the latter. In this case, we would expect the process X_t^N to move away from x^* if a "first step" away from the equilibrium took place, but we do not expect this first step to occur.

To formalize these ideas, one would need to show that local probabilistic stability was relevant to answering the following question: when the population size N and

---

decision procedure, this would likely yield different deterministic dynamics, and hence different LPS rest points.

the time horizon T are large, will X_T^N lie outside the vicinity of x^*? If we perform this analysis by fixing T and taking N to infinity (a "finite horizon" analysis), Corollary 4.2 tells us that escape will never occur. On the other hand, if we fix N and take T to infinity (an "infinite horizon" analysis), escape is guaranteed to occur: indeed, X_t^N will visit all states in some recurrent class $R^N \subset \Delta^N$ infinitely often.[29]

These two approaches can be viewed as polar attempts to characterize X_T^N when both N and T are large. Neither provides a discriminating local stability analysis when the process X_t^N is initially very close to equilibrium. To probe the middle ground, one can look at the behavior of X_T^N while taking N and T to infinity simultaneously. More precisely, one can choose a sequence of pairs (N, T); for each such pair, one can examine the behavior of the population size N process through time T. A local stability analysis would address whether X_T^N is outside the vicinity of x^* for pairs (N, T) far enough along the chosen sequence.[30] The informal arguments above suggest the possibility of four growth rate regimes: as the growth rate of N relative to T is reduced (i.e., as one moves from the finite horizon analysis towards the infinite horizon analysis), one might pass through regimes (*i*) where escape occurs for no equilibria; (*ii*) where escape only occurs for locally probabilistically unstable equilibria; (*iii*) where escape occurs for these equilibria as well as for locally probabilistically stable equilibria which are deterministically unstable; (*iv*) where escape occurs for all equilibria. The analysis of cases with simultaneous limits is an interesting topic for future research.

9. Conclusion

We study stochastic evolution under inexact information, focusing special attention on equilibrium play. We establish that finite horizon behavior away from equilibria can be described by a differential equation, and that finite horizon behavior near equilibria can be described by a diffusion. We define a new notion of evolutionary stability which explicitly accounts for the random variations in behavior created by inexact information, and use the diffusion approximation to establish a simple characterization of stability. While at interior equilibria, local

---

[29] It is well known that this recurrence property makes it possible to prove equilibrium selection results in stochastic evolutionary models – see the following section for further discussion.

[30] Sandholm and Pauzner (1998) perform an analysis in which N and T simultaneously diverge in the context of the Kandori, Mailath, and Rob (1993) model. However, in their analysis (unlike in the analysis proposed here), only a single stochastic process is analyzed, and the population size in this process grows as time elapses.

probabilistic stability accords closely with standard deterministic notions of stability, it is significantly less demanding at equilibria which lie on the boundary.

It is useful to contrast the model studied here with models of stochastic fictitious play (Fudenberg and Kreps (1993), Kaniovski and Young (1995), Benaïm and Hirsch (1999), Hofbauer and Sandholm (2002)). These models consider small groups of players who repeatedly play a normal form game. Players choose best responses to their recollections of the history of play. Suppose that payoffs are noisy, or that players' recollections of the history of play are incomplete. Then in certain classes of games, both the time average of past play and the players' choice probabilities converge with probability one over the infinite horizon; limit values approximate Nash equilibria of the underlying game.

In stochastic fictitious play models, the state variable represents the time average of past play. Consequently, the increments in the state variable become vanishingly small as time passes, making infinite horizon convergence results possible. Since in our model the state variable represents current behavior, increments are of fixed size. For this reason, the behavior process in our model need not converge over the infinite horizon. Instead, the possible limiting time averages of play can be described by stationary distributions, one for each recurrent class of the behavior process.

Most work on stochastic evolution in games has focused on infinite horizon behavior, with results stated in terms of stationary distributions. For example, Foster and Young (1990), Kandori, Mailath, and Rob (1993), and Young (1993) study stochastically stable equilibria, which are equilibria which receive all weight in the stationary distribution as the rate of mutation vanishes. These models offer unique predictions of infinite horizon behavior, even in games which exhibit multiple strict equilibria.

In the current context, Benaïm and Weibull (2003) have recently shown that when the population size is large, nearly all mass in any stationary distribution must lie near the minimal center of attraction of the deterministic law of motion. They also use techniques from large deviation theory to show that under certain general conditions, the limiting stationary distribution is unique and places all mass on a single component of this set. Applying this approach to obtain equilibrium selection results is an important topic for future research.

Appendix

We begin by stating the convergence results of Kurtz (1970) and Stroock and Varadhan (1979) used to prove Theorems 4.1 and 5.1; along the way we prove Theorem 4.1. Since both convergence results also have discrete time formulations, Theorems 4.1 and 5.1 can also be proved for discrete time versions of our model; see footnote 10.

For each $N \in \mathbf{N}$, let $\{Y_t^N\}_{t \geq 0}$ be a pure jump Markov process which takes values in some countable set $V^N \subset \mathbf{R}^m$. Each process has a Poisson rate λ^N which is independent of the current state. Let $Q^N(y, y')$ be the probability that when a jump occurs from state y it lands at state y'; we allow $Q^N(y, y)$ to be strictly positive. For each $y \in V^N$, define

$$b^N(y) = \sum_{y' \in V^N} (y' - y) \lambda^N Q^N(y, y');$$

$$d_\varepsilon^N(y) = \sum_{y: |y'-y| > \varepsilon} \lambda^N Q^N(y, y')$$

We first present a version of a result on the convergence of pure jump Markov processes to deterministic flows due to Kurtz (1970, Theorem 2.11)

Theorem A.1: *Suppose that each V^N is a subset of some bounded set $V \subset \mathbf{R}^m$, and that there is a Lipschitz continuous function b: $V \to \mathbf{R}^m$ and a constant $k \in \mathbf{R}$ such that*

$$(1.1) \quad \lim_{N \to \infty} \sup_{y \in V^N} \left| b^N(y) - b(y) \right| = 0;$$

$$(1.2) \quad \lim_{N \to \infty} \sup_{y \in V^N} d_{k/N}^N(y) = 0.$$

Let the initial conditions $Y_0^N = y_0^N$ converge to $y_0 \in S$. Then for all $T < \infty$ and $\varepsilon > 0$,

$$\lim_{N \to \infty} P\left(\sup_{t \in [0,T]} \left| Y_t^N - y_t \right| < \varepsilon \right) = 1,$$

where $\{y_t\}_{t \geq 0}$ is the solution to $\dot{y} = b(y)$ from initial condition y_0.

Condition (1.1) requires that the expected increments per time unit converge

uniformly to some limit function. Condition (1.2) demands that the probability of jumps bigger than order $\frac{1}{N}$ vanishes as N approaches infinity. Under these conditions, the finite horizon behavior of the pure jump Markov processes can be arbitrarily well approximated by the solution to a differential equation.

Proof of Theorem 4.1:

We apply Theorem A.1 to the process X_t^N. In this case, $b^N(x) \equiv f^N(x)$ and $b(x) \equiv f(x)$, so condition (1.1) follows from Assumption (A1). The Lipschitz continuity of b follows from Assumption (A2). Finally, condition (1.2) follows immediately from the fact that at most one player changes strategies during each revision opportunity. ∎

We now present a result on the convergence of Markov processes to diffusions due to Stroock and Varadhan (1979). Our presentation follows Durrett (1996). In particular, the following result follows from Theorems 8.7.1, 5.2.2, 5.4.1, and 5.4.5 and Lemma 8.8.2 of Durrett (1996). For all i and j we define

$$c_{ij}^N(y) = \sum_{y' \in V^N} (y_i' - y_i)(y_j' - y_j)\lambda^N Q^N(y,y');$$
$$\phi^N(y) = \sum_{y' \in V^N} |y' - y|^4 \lambda^N Q^N(y,y').$$

Theorem A.2: *Suppose there exist Lipschitz continuous functions b: $\mathbf{R}^m \to \mathbf{R}^m$, c: $\mathbf{R}^m \to \mathbf{R}^{m \times m}$, and σ: $\mathbf{R}^m \to \mathbf{R}^{m \times m}$ such that $c \equiv \sigma\sigma'$ and such that for all $R < \infty$,*

(2.1) $\quad \lim_{N \to \infty} \sup_{|y| < R} |b^N(y) - b(y)| = 0;$

(2.2) $\quad \lim_{N \to \infty} \sup_{|y| < R} |c^N(y) - c(y)| = 0;$

(2.3) $\quad \lim_{N \to \infty} \sup_{|y| < R} \phi^N(y) = 0;$

where the suprema are taken only over $y \in V^N$. Then if the initial conditions $Y_0^N = y_0^N$ approach $y_0 \in \mathbf{R}^m$, then for each $T < \infty$, Y_t^N converges weakly in $D([0, T], \mathbf{R}^m)$ to Y_t, the solution to

$$dY_t = b(Y_t)dt + \sigma(Y_t)dB_t$$

from initial condition $Y_0 = y_0$.

Conditions (2.1) and (2.2) require that the infinitesimal means and covariances of the pure jump Markov processes converge uniformly to some Lipschitz continuous functions. The fourth moment condition (2.3) bounds the probability of "large" jumps. If these conditions hold, then for large N, the finite horizon behavior of the pure jump Markov processes can be approximated arbitrarily well by a diffusion.

Proof of Theorem 5.1:

We apply Theorem A.2 to the local behavior process $Z_t^N \equiv \sqrt{N}(X_t^N - x^*)$, establishing that the functions $b^N(z)$ and $c^N(z)$ converge to $b(z) = D^*z$ and $a(z) = a^*$, respectively. To start, observe that since $(x_0^N - x^*) \in o(\frac{1}{\sqrt{N}})$, the initial conditions $Z_0^N = z_0^N = \sqrt{N}(x_0^N - x^*)$ converge to the origin. To establish the result, we need to show that conditions (2.1), (2.2), and (2.3) hold. Since at most one player can change strategies at any switching opportunity, condition (2.3) is easily verified. To establish condition (2.1), recall that the function $b^N(x)$ associated with the process X_t^N is just $f^N(x)$. Thus, since $Z_t^N \equiv \sqrt{N}(X_t^N - x^*)$, we find that

$$b^N(z) = \sqrt{N} f^N\left(x^* + \tfrac{z}{\sqrt{N}}\right).$$

Since $f(x^*) = 0$, the mean value theorem tells us that for all N and all $z \in \mathbf{R}^k$ there exists a $\lambda(z) \in [0, 1]$ such that

$$f\left(x^* + \tfrac{z}{\sqrt{N}}\right) = Df\left(x^* + \lambda(z)\tfrac{z}{\sqrt{N}}\right)\tfrac{z}{\sqrt{N}}.$$

Fix an $R < \infty$. Since assumption (A2) implies that $f^N: \Delta \to \mathbf{R}^k$ converges to f at rate $o(\frac{1}{\sqrt{N}})$, we find that

$$
\begin{aligned}
\limsup_{\substack{N\to\infty \\ |z|<R}} & \left|b^N(z) - D^*z\right| \\
&= \limsup_{\substack{N\to\infty \\ |z|<R}} \left|\sqrt{N}f^N\left(x^* + \tfrac{z}{\sqrt{N}}\right) - Df(x^*)z\right| \\
&\leq \limsup_{\substack{N\to\infty \\ |z|<R}} \left(\left|\sqrt{N}\left(f^N\left(x^* + \tfrac{z}{\sqrt{N}}\right) - f\left(x^* + \tfrac{z}{\sqrt{N}}\right)\right)\right| + \left|\sqrt{N}f\left(x^* + \tfrac{z}{\sqrt{N}}\right) - Df(x^*)z\right|\right) \\
&= \limsup_{\substack{N\to\infty \\ |z|<R}} \left(\left|\sqrt{N}\left(f^N\left(x^* + \tfrac{z}{\sqrt{N}}\right) - f\left(x^* + \tfrac{z}{\sqrt{N}}\right)\right)\right| + \left|Df\left(x^* + \lambda(z)\tfrac{z}{\sqrt{N}}\right)z - Df(x^*)z\right|\right)
\end{aligned}
$$

$$= 0.$$

This establishes condition (2.1).

To verify condition (2.2), observe that the function $c^N(x)$ associated with the process X_t^N is $rNC^N(x)$, where the function $C^N(\cdot)$ is defined in the text. Thus, the function $c^N(z)$ corresponding to $Z_t^N \equiv \sqrt{N}\left(X_t^N - x^*\right)$ is given by

$$c^N(z) = rN^2 C^N\left(x^* + \tfrac{z}{\sqrt{N}}\right).$$

Define $a^N: \Delta \to \mathbf{R}^{k \times k}$ by $a^N(x) = rN^2 C(x)$, so that $c^N(z) = a^N(x^* + \tfrac{z}{\sqrt{N}})$. By assumption (A3), the functions a^N converge uniformly to the function a defined in the text. Furthermore, a is continuous by assumption (A4), and $a(x^*) = a^*$ by definition. Hence,

$$\limsup_{N \to \infty \ |z| < R} \left|c^N(z) - a^*\right|$$
$$\leq \limsup_{N \to \infty \ |z| < R} \left(\left| a^N\left(x^* + \tfrac{z}{\sqrt{N}}\right) - a\left(x^* + \tfrac{z}{\sqrt{N}}\right) \right| + \left| a\left(x^* + \tfrac{z}{\sqrt{N}}\right) - a(x^*) \right| \right)$$
$$= 0.$$

This establishes condition (2.2).

It is clear from its definition that the matrix $a^* \in \mathbf{R}^{k \times k}$ is symmetric; moreover,

Lemma A.3: $a^* \in \mathbf{R}^{k \times k}$ is positive semidefinite.

Proof: It is enough to establish this for the representation of $a^* \in \mathbf{R}^{rn \times rn}$ in the original coordinates. Let $m_{ij} = m_{ji} = x_i^* \, d(i, j, x^*) + x_j^* \, d(j, i, x^*)$. Then for $z \in \mathbf{R}^{rn}$,

$$z'a^*z = \sum_i \sum_j z_i z_j a_{ij}^* = \sum_i z_i^2 \sum_{j \neq i} m_{ij} - \sum_i \sum_{j \neq i} z_i z_j m_{ij} = \sum_i \sum_{j \neq i} (z_i^2 - z_i z_j) m_{ij}$$
$$= \sum_{i>1} \sum_{j>i} (z_i^2 + z_j^2 - 2 z_i z_j) m_{ij} = \sum_{i>1} \sum_{j>i} (z_i - z_j)^2 m_{ij} \geq 0. \ \square$$

By Lemma A.3, there exists a matrix $\sigma^* \in \mathbf{R}^{k \times k}$ such that $\sigma^*(\sigma^*)' = a^*$. Therefore, Theorem A.2 implies that Z_t^N converges weakly to the solution to

(S) $dZ_t = D^* Z_t \, dt + \sigma^* dB_t$

with initial condition $Z_0 \equiv 0$. ∎

Proof of Theorem 6.1:

By Theorem 5.1, the process Z_t^N converge weakly to the process Z_t in $D([0, T], \mathbf{R}^k)$ for all $T < \infty$. Hence, the random variable Z_T^N converges in distribution to Z_T for each $T < \infty$. To determine the stability of x^* it is enough to consider whether Z_T converges in distribution as T approaches infinity.

Since Z_t is a mean zero Gaussian process, each random variable Z_T is multivariate normal with mean zero. Thus, if the covariance matrix $Cov(Z_T)$ converges as T approaches infinity, x^* is stable; if some diagonal component of $Cov(Z_T)$ heads to infinity, then x^* is unstable. Equation (5.6.14)' of Karatzas and Shreve (1991) tells us that

$$Cov(Z_T) = e^{D^*T} \int_0^T e^{-D^*t} a^* e^{-(D^*)'t} dt \, e^{(D^*)'T}.$$

Bringing all terms under the integral sign and performing the substitution $t = T - s$ yields

$$Cov(Z_T) = \int_0^T e^{D^*t} a^* e^{(D^*)'t} dt.$$

Since a^* is symmetric, positive semidefinite, and has full rank, it is positive definite, so there exists a full rank matrix $A \in \mathbf{R}^{k \times k}$ and a diagonal matrix $\Lambda \in \mathbf{R}^{k \times k}$ with strictly positive diagonal entries such that $a^* = A \Lambda A'$. Thus, letting $M^t = e^{D^*t} A$, we see that

$$Cov(Z_T) = \int_0^T M^t \Lambda (M^t)' \, dt$$

We can therefore compute each element of $Cov(Z_T)$ as follows:

$$(Cov(Z_T))_{jk} = \int_0^T \sum_l \sum_m M_{jl}^t \Lambda_{lm} M_{km}^t dt$$
$$= \int_0^T \sum_l M_{jl}^t M_{kl}^t \Lambda_{ll} \, dt.$$

Since e^{D^*t} is the matrix solution to the linear system (L), it follows from

Proposition 4.3.4 and Theorem 4.3.5 of Robinson (1995) that each component of e^{D^*t} is a linear combination of terms of the form $t^k e^{\alpha t} \cos \beta t$ and of the form $t^k e^{\alpha t} \sin \beta t$, where $\alpha + i\beta$ is an eigenvalue of D^* with multiplicity greater than k. Conversely, if $\alpha + i\beta$ is an eigenvalue of D^*, terms of both forms above with $k = 0$ must appear in some component of e^{D^*t}. Since A has full rank, these statements are also true of the matrix M^t.

If $\alpha < 0$ for all eigenvalues $\alpha + i\beta$ of D^*, then it is clear from the foregoing that the absolute value of each integrand $\sum_l M_{jl}^t M_{kl}^t \Lambda_{ll}$ decreases exponentially in t. Hence, each integrand is integrable, so $\lim_{T \to \infty} Cov(Z_T)$ exists and x^* is stable.

On the other hand, suppose that there is an eigenvalue $\alpha + i\beta$ of D^* with $\alpha \geq 0$. Then by the discussion above, there must be a pair (j, k) such that M_{jk}^t contains a term of the form $ce^{\alpha t} \cos \beta t$ with $c \neq 0$. Therefore, the integrand of the jth variance term, $(Cov(Z_T))_{jj}$, is

$$\sum_l (M_{jl}^t)^2 \Lambda_{ll} \geq c^2 e^{2\alpha t} (\cos \beta t)^2 \Lambda_{kk}.$$

Since $\alpha \geq 0$, this expression is not integrable. Thus, $(Cov(Z_T))_{jj}$ diverges, and x^* is unstable. This completes the proof of the theorem. ∎

Proof of Proposition 7.2:

Recall that $\hat{S}$ is the set of the k strategies which are represented explicitly as coordinates in $\mathbf{R}^k$, which includes all strategies which are unused at x^*. The set of unused strategies is denoted $\hat{S}_u$: that is, $\hat{S}_u = \{i \in \hat{S}: x_i^* = 0\}$. We proceed with three lemmas.

Lemma A.4: *If $i \in \hat{S}_u$, $\sigma_{ij}^* = 0$.*

Proof: Lemma 7.1 tells us that $a_{ii}^* = 0$. Since $a_{ii}^* = \sum_j (\sigma_{ij}^*)^2$ by definition, the result follows. ❏

Lemma A.5: *If $i \in \hat{S}_u$ and $j \in \hat{S} - \hat{S}_u$, then $D_{ij}^* = \frac{\partial f_i}{\partial x_j}(x^*) = 0$.*

Proof: Suppose that $\frac{\partial f_i}{\partial x_j}(x^*) \neq 0$. Then since $f_i(x^*) = 0$, and since $x_j^* \in (0, 1)$ by our naming convention, there is a real number k such that $x^* + k\iota_j \in \Delta$ and $f_i(x^* + k\iota_j) < 0$. But since $(x^* + k\iota_j)_i = 0$, this contradicts that Δ is forward invariant under the flow

defined by f. ❏

Lemma A.6: *If $i \in \hat{S}_u$ and $j \in \hat{S} - \hat{S}_u$, then $(e^{D^*t})_{ij} = 0$ for all $t \geq 0$.*

Proof: Follows from Lemma A.5 and the fact that $e^{D^*t} = I + \sum_{k=1}^{\infty} \frac{(D^*t)^k}{k!}$. ❏

Fix $i \in \hat{S}_u$. Lemmas A.4 and A.6 imply that $(e^{D^*t})_{ik}\sigma^*_{kj} = 0$ for all $j, k \in \hat{S}$. Thus, $(e^{D^*t}\sigma^*)_{ij} = \sum_k (e^{D^*t})_{ik}\sigma^*_{kj} = 0$ for all $j \in \hat{S}$. That is, the ith row of $e^{D^*t}\sigma^*$ is the zero vector for all $t \geq 0$. Since $Z_t \equiv \int_0^t e^{D^*(t-s)}\sigma^* \, dB_s$, we conclude that for all $t \geq 0$,

$$(Z_t)_i = \sum_j \int_0^t (e^{D^*(t-s)}\sigma^*)_{ij} \, d(B_s)_j = 0. \ \blacksquare$$

Proof of Theorem 7.4:

We begin with two definitions. We say that $A \in \mathbf{R}^{k \times k}$ is $\hat{S}_u$-*null* if $A_{ij} = A_{ji} = 0$ whenever $i \in \hat{S}_u$; we say that A *reduces to* $\tilde{A} \in \mathbf{R}^{\kappa \times \kappa}$ if $\tilde{A}$ is obtained from A by eliminating all rows and columns corresponding to i in $\hat{S}_u$.

Since r^* is symmetric, positive semidefinite, and has full rank, it is positive definite. Let ρ^* be a square root of r^*: $\rho^*(\rho^*)' = r^*$. Lemma 7.1 says that a^* is $\hat{S}_u$-null, and it reduces to r^* by definition; we can therefore choose σ^* to be the $\hat{S}_u$-null matrix which reduces to ρ^*. Moreover, Lemma A.5 and the definition of the exponential matrix imply that e^{D^*t} reduces to e^{R^*t} for all $t \geq 0$. Therefore, Lemma A.6 implies that $e^{D^*t}\sigma^*$ is $\hat{S}_u$-null and reduces to $e^{R^*t}\rho^*$.

The covariance matrix of the random variable Z_T is

$$Cov(Z_T) = \int_0^T e^{D^*t} a^* e^{(D^*)'t} dt = \int_0^T (e^{D^*t}\sigma^*)(e^{D^*t}\sigma^*)' \, dt.$$

Thus, since $e^{D^*t}\sigma^*$ is $\hat{S}_u$-null and reduces to $e^{R^*t}\rho^*$, $Cov(Z_T)$ is $\hat{S}_u$-null and reduces to

$$\tilde{\Sigma}_T = \int_0^T e^{R^*t} r^* e^{(R^*)'t} dt = \int_0^T (e^{R^*t}\rho^*)(e^{R^*t}\rho^*)' \, dt.$$

Therefore, since r^* is symmetric and positive definite, the proof is completed by applying the proof of Theorem 6.1 to $\tilde{\Sigma}_T$. $\blacksquare$

References

Beggs, A. (2002). "Stochastic Evolution with Slow Learning," *Econ. Theory* **19**, 379-405.

Benaïm, M., and M. W. Hirsch (1999). "Mixed Equilibria and Dynamical Systems Arising from Repeated Games," *Games Econ. Behav.* **29**, 36-72.

Benaïm, M., and J. W. Weibull (2003). "Deterministic Approximation of Stochastic Evolution in Games," forthcoming, *Econometrica*.

Binmore, K. J., and L. Samuelson (1999). "Evolutionary Drift and Equilibrium Selection," *Rev. Econ. Stud.* **66**, 363-393.

Binmore, K. J., L. Samuelson, and R. Vaughan (1995). "Musical Chairs: Modeling Noisy Evolution," *Games Econ. Behav.* **11**, 1-35.

Börgers, T. and R. Sarin (1997). "Learning through Reinforcement and the Replicator Dynamics," *J. Econ. Theory* **77**, 1-14.

Boylan, R. T. (1995). "Continuous Approximation of Dynamical Systems with Randomly Matched Individuals," *J. Econ. Theory* **66**, 615-625.

Cabrales, A. (2000). "Stochastic Replicator Dynamics," *Int. Econ. Rev.* **41**, 451-481.

Corradi, V., and R. Sarin (2000). "Continuous Approximations of Stochastic Evolutionary Game Dynamics," *J. Econ. Theory* **94**, 163-191.

Durrett, R. (1996). *Stochastic Calculus: A Practical Introduction*. Boca Raton: CRC Press.

Foster, D., and H. P. Young (1990). "Stochastic Evolutionary Game Dynamics," *Theor. Population Biol.* **38**, 219-232.

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

Fudenberg, D., and D. M. Kreps (1993). "Learning Mixed Equilibria," *Games Econ. Behav.* **5**, 320-367.

Gilboa, I., and A. Matsui (1991). "Social Stability and Equilibrium," *Econometrica* **59**, 859-867.

Hofbauer, J., and W. H. Sandholm (2002). "On the Global Convergence of Stochastic Fictitious Play," *Econometrica* **70**, 2265-2294.

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

Kaniovski, Y. M., and H. P. Young (1995). "Learning Dynamics in Games with Stochastic Perturbations," *Games Econ. Behav.* **11**, 330-363.

Karatzas, I., and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus*, 2nd ed. New York: Springer.

Kurtz, T. G. (1970). "Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes," *J. Appl. Prob.* **7**, 49-58.

Robinson, C. (1995). *Dynamical Systems*: *Stability*, *Symbolic Dynamics*, *and Chaos*. Boca Raton: CRC Press.

Sandholm, W. H. (2001). "Almost Global Convergence to p-Dominant Equilibrium," *Int. J. Game Theory* **30**, 107-116.

Sandholm, W. H., and A. Pauzner (1998). "Evolution, Population Growth, and History Dependence," *Games Econ. Behav.* **22**, 84-120.

Schlag, K. H. (1998). "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits," *J. Econ. Theory* **78**, 130-156.

Stroock, D. W., and S. R. S. Varadhan (1979). *Multidimensional Diffusion Processes*. New York: Springer.

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