Econometrica, Vol. 71, No. 3 (May, 2003), 873–903

DETERMINISTIC APPROXIMATION OF STOCHASTIC EVOLUTION IN GAMES

BY MICHEL BENAÏM AND JÖRGEN W. WEIBULL[1]

This paper provides deterministic approximation results for stochastic processes that arise when finite populations recurrently play finite games. The processes are Markov chains, and the approximation is defined in continuous time as a system of ordinary differential equations of the type studied in evolutionary game theory. We establish precise connections between the long-run behavior of the discrete stochastic process, for large populations, and its deterministic flow approximation. In particular, we provide probabilistic bounds on exit times from and visitation rates to neighborhoods of attractors to the deterministic flow. We sharpen these results in the special case of ergodic processes.

KEYWORDS: Game theory, evolution, approximation, large deviations, Markov chains.

1. INTRODUCTION

MANY MODELS IN EVOLUTIONARY game theory hypothesize an infinitely large population of interacting agents, usually represented as a continuum, and describe the evolutionary process as deterministic, defined in terms of a system of ordinary differential or difference equations. These equations concern changes in population shares, one for each pure strategy in the game, and the changes are viewed as averages over a large number of individual strategy switches. Prime examples are different versions of the replicator dynamics (Taylor and Jonker (1978), Taylor (1979), Maynard Smith (1982)). For a wide class of such dynamics it has been established that dynamic (Lyapunov) stability implies Nash equilibrium, and that the limit point to any convergent trajectory through any initial population state with all pure strategies present is a Nash equilibrium.[2] For a certain subclass it has been shown that, even if the solution trajectory diverges, all iteratively strictly dominated pure strategies nevertheless vanish asymptotically.[3] Moreover, attractors that contain essential components and strategically stable

[1] The authors thank the former editor and four anonymous referees for helpful comments. We are also grateful for comments from Carlos Alos-Ferrer, Stefano Demichelis, Georg Pflug, William Sandholm, and participants at presentations at Alicante University, the First World Congress of Game Theory Society (Bilbao), the Bolzano 2000 European Economic Association Annual Congress, Copenhagen University, MIT, the Henri Poincaré Institute (Paris), Odense University, and Vienna University. We are grateful for travel funding from the Hedelius and Wallander Foundation. We thank Ariane Lambert-Mogilianski for getting us to meet. This manuscript is a revised version of Benaïm and Weibull (2000).

[2] See, e.g., Bomze (1986), Nachbar (1990), and Weibull (1995).

[3] See, e.g., Samuelson and Zhang (1992) and Hofbauer and Weibull (1996).

sets of Nash equilibria have been identified for certain classes of such deterministic dynamics.[4] An important question for the relevance of these results, and for deterministic population models more generally, is whether these dynamics are good approximations of stochastic population processes that arise from individual strategy adaptation in finite but large populations. The present study addresses three versions of this question, each version corresponding to a precise meaning of "good approximation." In the special case of ergodic processes, we also provide asymptotic results for the stochastic process that go beyond those for its deterministic approximation dynamics.

Technically, this is achieved by applying and extending mathematical results in the theory for large stochastic deviations to a class of Markov chains that live in compact polyhedra. We interpret these Markov chains as population processes in finite games, where individuals are recurrently drawn from finite populations to review their strategy choice in the game. There is one population for each player position, and all populations are of the same finite size N. Each individual has a pure strategy that he or she uses if called upon to play the game. At discrete times $t = 0, \delta, 2\delta, \ldots$ exactly one individual is given the opportunity to change his or her strategy in the light of some information about the current payoff to one or more pure strategies in her player position. All individuals have the same probability of being drawn for such a strategy review, and we assume statistical independence across populations and over time.

The state of the population process is defined as the vector of population shares associated with the pure strategies in the game. Hence, a population state is mathematically equivalent to a mixed-strategy profile, with all probabilities being multiples of the factor $1/N$. The state space is thus a finite grid in the polyhedron of mixed-strategy profiles, and a population share can only change by $\pm 1/N$ units at each review opportunity. We set the time interval δ between successive strategy reviews equal to $1/N$, and study the limit as $N \to \infty$. Hence, the expected time span between two successive strategy revision opportunities for any given individual is kept constant as N is increased. The approximation results apply when the population size N is large, and the approximation is defined in terms of a system of ordinary differential equations derived from the transition probabilities of the Markov chain—the so-called mean-field equations.[5]

We provide examples in which these differential equations are of the type studied in evolutionary game theory. For example, different versions of the replicator dynamic can arise if strategy choices are based on imitation or smooth best replies. However, the stochastic processes here differ qualitatively from those arising in stochastic fictitious play.[6] The latter processes have "decreasing gain"

[4] See, e.g., Swinkels (1993), Ritzberger and Weibull (1995), and Demichelis and Ritzberger (2001).

[5] In this respect, our approach is close to that of Börgers and Sarin (1997). They study a model of stochastic reinforcement learning where, like here, the size of jumps are reduced at the same rate as the time rate of jumps is increased. However, their stochastic processes have continuum state spaces and take jumps of unequal size, while our stochastic processes have discrete state spaces and take jumps of equal size $(1/N)$.

[6] See Fudenberg and Kreps (1993), Kaniovski and Young (1995), Fudenberg and Levine (1998), and Benaïm and Hirsch (1999b), or, for a more mathematical treatment, Benaïm (1999).

in the sense that the magnitude of state transitions is decreasing over time. For in fictitious play the state is the vector of accumulated empirical frequencies of strategies used in all previous rounds, and the change in these frequencies tends to zero over time. These processes thus slow down over time. By contrast, the processes studied here have "constant gain:" the effect on the state at each strategy revision is of equal magnitude ($1/N$) all the time.

We establish precise connections between the long-run behavior of the stochastic process and its deterministic approximation. In particular, we show that if the deterministic solution through the initial state of the stochastic process at some point in time enters a basin of attraction, then the stochastic process will enter any given neighborhood of the attractor inside its basin of attraction in a finite and deterministic time, with a probability that exponentially approaches one as the population size goes to infinity. The process will remain in this neighborhood for a random time that exceeds an exponential function of the population size. During the random time interval spent in the neighborhood, the process spends almost all time near a certain subset of the attractor, the Birkhoff center of the flow restricted to the attractor. If the process is ergodic, then it will eventually end up near the Birkhoff center of a generically unique attractor.

Our analysis proceeds in four steps, the first of which is a sharpening of fairly well-known results that leads up to more novel results. The first approximation result concerns the deviation of the stochastic process from the solution trajectory of its mean-field approximation dynamics during a given time interval. More exactly, for any bounded time interval and finite population size N, we provide an upper bound on the probability that the stochastic population process will depart more than a prescribed distance from the deterministic solution trajectory during that time interval. This upper bound goes exponentially to zero as the population size goes to infinity. In this part of the analysis, the time horizon is thus fixed and finite, while the population size is taken to infinity. Such "averaging theorems" are part of the folklore in the literatures on random perturbations of dynamical systems.[7] However, the present result provides a sharp exponential estimate of the deviation probability, based on martingale inequalities. It generalizes and sharpens results for games in Boylan (1995), Binmore, Samuelson, and Vaughn (1995), Börgers and Sarin (1997), Binmore and Samuelson (1997), Corradi and Sarin (1999), and Sandholm (1999).

The second approximation result concerns the first exit time from sets. We show that, for large population sizes N, the exit time from any neighborhood of the deterministic solution is probabilistically very large. More exactly, for any given time $t > 0$, the exit time from any neighborhood of the closure of the forward orbit (the orbit from time zero on) through the initial state of the stochastic process exceeds t for all but finitely many population sizes N, with probability one. Consequently, the exit time from the basin of attraction of an attractor

[7] See, e.g., Freidlin and Wentzell (1984), and, on stochastic approximation see, e.g., Duflo (1996, 1997), Kushner and Yin (1997), and Benaïm (1998, 1999).

in the deterministic dynamics is large when the population is large. Moreover, we provide bounds for the probability distribution of the exit time from certain neighborhoods of such attractors for large but finite populations.

The third approximation result concerns empirical visitation rates to sets, i.e., the long-run time fraction that the stochastic process spends in a given set. Using this random variable, we identify a certain set that has the property that the Markov chain spends almost all time, in the long run, at the set. This set is called the minimal center of attraction of the deterministic flow. It contains all stationary states and periodic orbits, and it is contained in the (more easily identified) Birkhoff center of the dynamics. We also provide results for conditional visitation rates, i.e., empirical visitation rates until the first exit time from a given neighborhood of the set in question. Conditional visitation rates are actually highly relevant because of the possibility of meta-stability, i.e., the possibility (in the limit as $N \to \infty$) of an infinite expected exit time from a subset—for example, a neighborhood of an attractor—that the stochastic process with probability one eventually will leave.

In the first result, we fix the time horizon and let the population go to infinity. In results two and three, we first take the time horizon to infinity for a fixed population size, and only thereafter do we take the population size to infinity, thereby studying the asymptotic behavior of the stochastic process when the population is fixed but large. Our results on exit times and visitation rates build on ideas in Freidlin and Wentzell (1984) and Benaïm (1998, 1999), and are complementary to the asymptotic results in Ellison (1993), Binmore, Samuelson, and Vaughn (1995), and Binmore and Samuelson (1997).[8]

We finally study large deviations in the special case when the population process is ergodic—and thus admits a unique invariant probability distribution over the state space. Following Ellison (2000), we define the "radius" and "co-radius" of attractors of the deterministic approximation flow, and establish an asymptotic result for the support of the invariant distribution for large finite populations. The main difference is that while Ellison keeps the population size fixed and changes the stochastic micro model by taking the mutation rate to zero, we keep the stochastic micro model fixed and take the population size to infinity. In keeping the mutation rate fixed in the limit, we follow Binmore and Samuelson (1997) and Young (1998b, Section 4.5).

The rest of the text is organized as follows. The studied class of Markov chains is defined in Section 2 and the general approximation results are presented in Section 3. Section 4 briefly discusses applications to a few micro models of boundedly rational strategy choice. Section 5 provides asymptotic results for ergodic population processes. Section 6 shows how the analysis can be adapted to continuous-time Poisson processes. Mathematical proofs are given in an Appendix at the end of the paper.

[8] For a discussion of the relation between the two iterated limits in question, see also Gale, Binmore, and Samuelson (1995) and Börgers and Sarin (1997).

2. A CLASS OF STOCHASTIC PROCESSES

Consider a finite n-player game with player roles (or "positions") $i \in I = \{1, \ldots, n\}$, finite pure strategy sets $S_i = \{1, \ldots, m_i\}$, a set of pure-strategy profiles $S = \times_i S_i$, mixed-strategy simplices

$$(1) \qquad \Delta(S_i) = \left\{ x_i \in \mathbb{R}_+^{m_i} : \sum_{h \in S_i} x_{ih} = 1 \right\},$$

and polyhedron $\square(S) = \times_i \Delta(S_i)$ of mixed-strategy profiles $x = (x_1, \ldots, x_n)$. The polyhedron $\square(S)$ is thus a subset of $\mathbb{R}^M$, for $M = \sum_i m_i$. For any player role i and pure strategy $h \in S_i$, let $e_i^h \in \Delta(S_i)$ denote the corresponding unit vector—the mixed strategy for player i that assigns unit probability to pure strategy h.

For each player role i there is a population of N individuals. Each individual is at every moment in time associated with a pure strategy in her strategy set. An individual in population i who is associated with pure strategy $h \in S_i$ is called an *h-strategist*. At times $t \in \mathbb{T} = \{0, \delta, 2\delta, \ldots\}$, where $\delta = 1/N$, and only then, exactly one individual has the opportunity to change his or her pure strategy. This individual is randomly drawn, with equal probability for all nN individuals, and with statistical independence between successive draws.[9] With this fixed relationship between population size and period length, the expected time interval between two successive draws of one and the same individual is n, independently of the population size N.[10] We will call the times $t \in \mathbb{T}$ *transition times*—the only times when a transition *can* take place.

The specific models to be studied each define a Markov chain $X^N = (X^N(t))_{t \in \mathbb{T}}$ with finite state space $\square^N(S)$ in the polyhedron $\square(S)$ of mixed-strategy profiles.[11] The *state* $X^N(t)$ at any time $t \in \mathbb{T}$ specifies, for every player role $i \in I$ and pure strategy $h \in S_i$, the share $X_{ih}^N(t)$ of h-strategists in population i. The only state transitions that can occur are that one individual in one population changes pure strategy. For every player role i and pair (h, k) of pure strategies for that role, we assume that there exists a continuous function $p_{ik}^h : \square(S) \to [0, 1]$ such that $p_{ik}^h(x) = 0$ if $x_{ik} = 0$, and

$$(2) \qquad p_{ik}^h(x) = \Pr\left[X_i^N\left(t + \frac{1}{N}\right) = x_i + \frac{1}{N}(e_i^h - e_i^k) \,\middle|\, X^N(t) = x \right]$$

for all $i \in I, h, k \in S_i$, $N \in \mathbb{N}$, and $x \in \square^N(S)$. In other words: the conditional probability that a k-strategist in population i will become a h strategist is continuous in the current state. In particular, it is independent of calendar time t

[9] The subsequent analysis is valid also if n individuals were simultaneously (and statistically independently) drawn, one from each player population. The only difference is that the stochastic process would be a factor n faster, and thus the vector field F a factor n stronger.

[10] There are N draws per time unit, and each time the probability that a particular individual is drawn is one over the total population size, nN.

[11] The finite set $\square^N(S)$ is the subset ("grid") of points x in the polyhedron $\square(S)$ such that Nx_{ih} is a nonnegative integer for all i and h.

and population size N.[12] The corresponding transition probabilities are, for any $v \in \mathbb{R}^M$,

$$\Pr\left[X^N\left(t+\frac{1}{N}\right) = x + \frac{1}{N}v \,\middle|\, X^N(t) = x\right] = \begin{cases} p_{ik}^h(x) & \text{if } v_i = e_i^h - e_i^k \text{ and } v_j = 0 \ \forall j \neq i, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

2.1. *The Induced Vector Field*

For any player role $i \in I$ and pure strategy $h \in S_i$, the *expected net increase* in the subpopulation of h-strategists from one transition time to the next, conditional upon the current state x, is

$$F_{ih}(x) = \sum_{k \neq h} p_{ik}^h(x) - \sum_{k \neq h} p_{ih}^k(x). \tag{4}$$

It follows from the probability specification above that $F_{ih} : \Box(S) \to \mathbb{R}$ is bounded and continuous, with $\sum_h F_{ih}(x) \equiv 0$ and $F_{ih}(x) \geq 0$ if $x_{ih} = 0$.

Recall that the polyhedron $\Box(S)$ is a subset of $\mathbb{R}^M$, where $M = \sum_i m_i$. Let $m = M - n$, let E^1 denote the m-dimensional hyperplane of $\mathbb{R}^M$ that contains $\Box(S)$, and let E^0 be the parallel m-dimensional subspace, the tangent space of E^1:

$$E^1 = \left\{x \in \mathbb{R}^M : \sum_h x_{ih} = 1 \quad \forall i\right\} \quad \text{and} \quad E^0 = \left\{x \in \mathbb{R}^M : \sum_h x_{ih} = 0 \quad \forall i\right\}. \tag{5}$$

We identify E^1 and E^0 with $\mathbb{R}^m$, call this space E, and view the function F as a mapping from E to E. This function is assumed to be bounded and locally Lipschitz continuous.[13] We will call F the *vector field* associated with the Markov chain X^N. By "open sets" in $\Box(S)$ we mean "open in the Euclidean topology induced on $\Box(S) \subset E$."

REMARK 1: In the special case of symmetric n-player games, an alternative setting is that of a single population consisting of N individuals. The present machinery then applies by letting the Markov chain X^N have the finite state space Δ^N in the common unit simplex Δ of mixed strategies. The *state* $X^N(t) \in \Delta^N$ at any time t now specifies, for each pure strategy h in the common pure-strategy set, the share of h-strategists in the population.

[12] This excludes, in particular, the possibility that individual behaviors change as the population size changes, from, say, some form of best-reply behavior in small populations to some form of imitation behavior in large populations.

[13] The function F is bounded and continuous on the polyhedron, by virtue of these properties of the transition probabilities, and can hence be extended to the whole space E while preserving these properties. We strengthen the continuity assumption by requiring local Lipschitz continuity. A function F: $E \to E$ has this property if for every compact subset $C \subset E$ there exists a scalar λ_C such that $\|F(x) - F(y)\| < \lambda_C \|x - y\|$ for all $x, y \in C$, where $\|\cdot\|$ is a norm on $E = \mathbb{R}^m$.

3. DETERMINISTIC APPROXIMATION

We are interested in deterministic approximation of Markov chains X^N in the class defined above, when the population size N is large, and thus the time interval $\delta = 1/N$ between successive transition times is short. The key element for such approximation is the vector field $F : E \to E$ defined above, which, for large populations and short time intervals, gives the expected net increase in each population *share* during the time interval, *per time unit*. (There are N transition times per time unit and N individuals in each player population.) The associated *mean-field equations*

$$\dot{x}_{ih} = F_{ih}(x) \quad \forall\, i \in I, h \in S_i, x \in E \tag{6}$$

together specify this limiting deterministic dynamic (a dot over a state variable denotes its time derivative).

In force of the Picard-Lindelöf Theorem, the system (6) of first-order ordinary differential equations has a unique solution through every point x in E. Moreover, as noted above, the sum of all population shares in each population remains constant over time, and no population share can turn negative. Hence, the system of equations (6) defines a *solution mapping* $\xi : \mathbb{R} \times \square(S) \to E$ that leaves each mixed-strategy simplex $\Delta(S_i)$, and hence also the polyhedron $\square(S)$ of mixed-strategy profiles, *forward invariant*. In other words, the system of differential equations determines a solution for all times $t \in \mathbb{R}$, and if the initial state is in $\square(S)$, then also all future states are in $\square(S)$.[14] We will frequently call ξ the *induced flow*.

We are now in a position to address the three questions in what precise sense, if any, the induced flow ξ approximates the Markov chain $X^N = (X^N(t))_{t \in \mathbb{T}}$ when N is large. As N changes, in general also the initial state $X^N(0) \in \square^N(S)$ has to change, since the finite state space $\square^N(S) \subset \square(S)$ changes with N (all population shares are multiples of $1/N$). We will frequently write "$X^N(0) \to x$" as a shorthand notation for "for every positive integer N, $X^N(0) \in \square^N(S)$, and $X^N(0)$ converges to $x \in \square(S)$ as N goes to plus infinity."

3.1. *Trajectories Over Bounded Time Intervals*

Our first result gives an exact form to the heuristic "law of large numbers" that says that the stochastic population process with high probability moves close to the associated deterministic population flow during any given bounded time interval, granted the population is large enough. We measure the fit of the deterministic approximation over bounded time intervals in terms of the

[14] More exactly: $\xi(0, x) = x$ for all x, $\partial \xi_{ih}(t, x)/\partial t = F_{ih}[\xi(t, x)]$ for all i, h, x, and t, and $\xi_i(t, x) \in \Delta(S_i)$ for all $i \in I$, $x \in \square(S)$, and $t > 0$. The time domain of the solution mapping ξ can be taken to be the whole real line in force of the compactness of $\square(S)$.

interpolated continuous-time process $\widehat{X}^N$ defined by piecewise affine interpolation of the Markov chain X^N.[15]

In order to state this result more precisely, let $\|\cdot\|_\infty$ denote the L^∞-norm on $E = \mathbb{R}^m$. Then $\|\widehat{X}^N(t) - \xi(t,x)\|_\infty$ represents the deviation of the interpolated Markov chain from the deterministic approximation solution ξ at time t, measured as the largest deviation in any population share at time t:

$$\|\widehat{X}^N(t) - \xi(t,x)\|_\infty = \max_{i\in I,\, h\in S_i} \left|\widehat{X}^N_{ih}(t) - \xi_{ih}(t,x)\right|. \tag{7}$$

The random variable

$$D^N(T,x) = \max_{0\le t\le T} \|\widehat{X}^N(t) - \xi(t,x)\|_\infty \tag{8}$$

is thus the *maximal deviation* in any population share during a bounded time interval $[0,T]$.[16] The proof of the following result is based on exponential martingale inequalities, enabling an exponential upper bound:

LEMMA 1: *There exists a scalar $c > 0$ such that, for any $\varepsilon > 0$, $T > 0$, and any N large enough*:

$$\Pr[D^N(T,x) \ge \varepsilon \mid X^N(0) = x] \le 2me^{-\varepsilon^2 cN},$$

for all $x \in \square^N(S)$.

In other words: for a fixed game, vector field F, deviation ε, and finite time horizon T, the probability of a larger deviation in any of the population shares is, for large enough populations, bounded from above by an exponentially decreasing function of the population size N. Consequently, for any given finite time horizon, the deterministic population flow ξ, induced by the vector field F, uniformly approximates the stochastic process over the time interval arbitrarily well, provided the population is sufficiently large: the probability in Lemma 1 goes to zero as N goes to plus infinity. This last claim is not a new result, however. Binmore, Samuelson, and Vaughn (1995) establish a version of this claim for a particular process in symmetric 2×2 coordination games; see also Boylan (1995), Börgers and Sarin (1997), Binmore and Samuelson (1997), Corradi and Sarin (1999), and Sandholm (1999). The value added in Lemma 1 is the exponential and hence summable ($\sum_N e^{-aN} < +\infty$) bound, which allows us to go beyond earlier results by way of the Borel-Cantelli Lemma; see Propositions 1–3.[17]

[15] The values of the interpolated process, at any time $t \in [n\delta, (n+1)\delta]$, are defined by

$$\widehat{X}^N(t) = X^N(n\delta) + \frac{t-n\delta}{\delta}[X^N((n+1)\delta) - X^N(n\delta)].$$

[16] It is immaterial whether one writes "max" or "sup" in this equation.

[17] The *Borel-Cantelli Lemma* states that if a sequence $\langle A_n\rangle_{n\in\mathbb{N}}$ of events A_n is such that the sum of their probabilities is finite, $\sum_n \Pr(A_n) < +\infty$, then the probability is zero that infinitely many of them occur, $\Pr(\limsup_{n\to\infty} A_n) = 0$.

3.2. *Exit Times from Sets*

The results in this subsection concern exit times from sets in the state space. They give alternative exact forms to the heuristic "law of large numbers" that if the population is large, and the deterministic population flow remains forever in some subset of the state space, then also the stochastic process will remain there for a very long time with a probability arbitrarily close to one, granted the population is large enough.[18]

For any Borel set $U \subset E$, and given X^N with $X^N(0) \in U$, the *exit time* from U is the random variable

(9) $$\tau^N(U) = \inf\{t \geq 0 : \widehat{X}^N(t) \notin U\}.$$

For any state $x \in \square(S)$, let $\gamma(x)$ be the *orbit* of the deterministic flow ξ through x—the set of states $y \in \square(S)$ such that $y = \xi(t, x)$ for some $t \in \mathbb{R}$—and let $\gamma^+(x)$ be the *forward orbit*—the set of states $y \in \square(S)$ such that $y = \xi(t, x)$ for some $t \geq 0$. In other words, the orbit through a state x is the set of all states that have been or will be reached, granted the state at time zero is x. The forward orbit is the subset of these states that are reached from time zero on. Neither the orbit nor the forward orbit need be closed sets. This is, for instance, the case if the solution trajectory converges to a stationary state from some nonstationary state: it does not reach the stationary state in finite time. Hence, without moving far away from the deterministic flow the stochastic process may anyhow leave a neighborhood of the forward orbit, since such a neighborhood need not contain a neighborhood of the limit state in question.[19] Therefore, we instead consider neighborhoods of the closure $\overline{\gamma^+(x)}$ of the forward orbit through an initial state x. The stationary state in the just mentioned example clearly belongs to the closure of the forward orbit, and any neighborhood of this closure is also a neighborhood of the limit state.

Combining Lemma 1 with the Borel-Cantelli Lemma, we obtain that, for any open set U containing the closure of the forward orbit, and for any $t > 0$, the exit time $\tau^N(U)$ from U exceeds t for all but finitely many $N \in \mathbb{N}$, with probability one.[20] Hence, we have the following result.[21]

PROPOSITION 1: *Suppose $U \subset E$ is open, $\overline{\gamma^+(x)} \subset U$, and $X^N(0) \to x$. Then*

$$\Pr\left[\lim_{N\to\infty} \tau^N(U) = +\infty\right] = 1.$$

Proposition 1 can be applied to attractors of the deterministic flow. The following result establishes that every neighborhood of the attractor, within its basin

[18] Ellison (1993) establish such a result, for an ergodic process, for exit times from neighborhoods of strict Nash equilibria in recurrently played 2×2-coordination games.

[19] A *neighborhood* of a set A is a set that contains an open set B that contains A.

[20] Here the summability of the exponential bound in Lemma 1 is important.

[21] This and many subsequent results implicitly refer to a common probability space that can easily be constructed; see Appendix for details.

of attraction, contains some neighborhood from which the exit time is with high probability very large when the populations are large. In this sense, attractors are good predictors also for the stochastic process, granted it starts near them. We obtain this result by way of a more powerful and precise result (see Appendix, Lemma 3) which gives exponential bounds on the rate at which the exit time goes to infinity as the population size goes to infinity.

Formally, a nonempty compact set $A \subset \Box(S)$ is an *attractor* of the deterministic flow ξ if it is *invariant*, $\gamma(x) \subset A$ for all $x \in A$, and has a neighborhood U with the property that $\lim_{t\to\infty} d[\xi(t,x), A] = 0$ uniformly in $x \in U$, where $d(x, C)$ denotes the (Hausdorff) distance between a point x and a closed set C. The *basin of attraction* of an attractor A is the set

$$B(A) = \left\{x \in E : \lim_{t\to\infty} d[\xi(t,x), A] = 0\right\}.$$

PROPOSITION 2: *Let $A \subset \Box(S)$ be an attractor of the flow ξ; let $V \subset B(A)$ be any neighborhood of A. Then there exists a neighborhood $U \subset V$ of A with closure $\overline{U}$ in $B(A)$, such that* $\Pr[\liminf_{N\to\infty} \tau^N(U) = +\infty] = 1$ *if* $X^N(0) \in U$.

Together with Lemma 1, this implies that if the deterministic flow through the initial state of the stochastic process at some point in time enters the basin of attraction of some attractor, then the stochastic process will enter any given neighborhood of that attractor within a finite and deterministic time with a probability that exponentially approaches one as the population size goes to infinity, and the process will remain in this neighborhood for a random time which with probability one exceeds any upper bound as the population size goes to infinity.

By contrast, if an invariant set A is not an attractor, then there exists some outgoing deterministic solution orbit, from initial states near A. One can then show that if the stochastic process starts near such an initial state, then it will depart in finite time from any neighborhood of A, with probability one as the population size is taken to infinity. More exactly, suppose $A \subset \Box(S)$ is a compact invariant set under ξ, and let U be a neighborhood of A such that the complement to A in U contains no invariant set. In other words, U is an *isolating* neighborhood of A. For any state $x \in \Box(S)$, let $\alpha(x)$ be the alpha-limit set of x, i.e., the set of states $y \in \Box(S)$ such that $\lim_{k\to+\infty} \xi(t_k, x) = y$ for some unbounded decreasing sequence of times $t_k < 0$. Let ∂U denote the boundary of the neighborhood U, and let U' be the subset of U that consists of states $x \in U$ such that $\alpha(x) \in A$ and $\xi(t,x) \in \partial U$ for some $t > 0$. In other words, U' consists of those states x in the neighborhood U that belong to solution orbits that originate arbitrarily close to A in the distant past and that reach the boundary of the neighborhood in finite time. It is well-known that if A is not an attractor, then U' is nonempty; see, e.g., Conley (1978).

PROPOSITION 3: *Suppose $A \subset \Box(S)$ is a compact invariant set with isolating neighborhood U, and let U' be as defined above. If A is not an attractor of ξ, then $U' \neq \varnothing$, and* $\Pr[\limsup_{N\to\infty} \tau^N(U) < +\infty] = 1$ *if* $X^N(0) \to x' \in U'$.

Note the contrast with Proposition 1, which states that the stochastic process, in the limit as N goes to infinity, remains forever in a neighborhood of every forward orbit—even if its limit is an unstable state. How can these two results be reconciled? The answer lies in the hypothesis concerning the initial state. Suppose x is a stationary but unstable state in the deterministic approximation dynamics. The forward orbit is thus $\{x\}$, and Proposition 1 says that if the initial states of the stochastic processes X^N converge to x as $N \to \infty$, then the exit time from any neighborhood U of x is probabilistically very large when the population is large, in the sense that it exceeds any given time t for all but finitely many N, with probability one. By contrast, Proposition 3 says that the stochastic process almost surely will leave any such neighborhood U in finite time, granted the initial states of the stochastic processes X^N converge to some state x' in the nonempty subset $U' \subset U$. Clearly $x \notin U'$, and thus x' lies at some (possibly small but) finite distance from the stationary state x. By definition, $x \in \alpha(x')$, but it takes the deterministic flow an infinite amount of time to reach x' from x, and, by continuity, an arbitrarily long time to reach x' from an initial state that is arbitrarily close to x. However, starting at $x' \neq x$, the deterministic flow leaves U after a finite time. The two results essentially say that the same holds for the stochastic processes. The contrast is particularly stark in the special case when x is a repellor. Then the above set U' is the whole set U except for one point, namely x, so in this special case the conclusion in the above proposition can be strengthened to the claim that if $X^N(0) \to x' \in U$, where $x' \neq x$, then the limit superior of the sequence $\tau^N(U)$, as $N \to \infty$, is almost surely finite. By contrast, if x instead were an asymptotically stable state, then Proposition 1 would imply that if $X^N(0) \to x' \in U$, then $\tau^N(U)$ almost surely goes to infinity as $N \to \infty$.

3.3. *Visitation Rates to Sets*

We next study how often, in the long-run, the stochastic process visits a given set. Such time fractions are called empirical visitation rates, and we here study both unconditional and conditional such rates.

First, for any Borel set $U \subset E$ and time T, let $V^N(U, T)$ be the fraction of transition times that the Markov chain X^N visits U in the time interval $[0, T]$, its (empirical) *visitation rate* in U during that time interval:

$$V^N(U, T) = \frac{1}{|\mathbb{T}(T)|} \sum_{t \in \mathbb{T}(T)} 1_{\{X^N(t) \in U\}}, \tag{10}$$

where $\mathbb{T}(T) = \mathbb{T} \cap [0, T]$ is the subset of the transition times that fall in the time interval $[0, T]$, and where $|\mathbb{T}(T)|$ is the number of elements in $\mathbb{T}(T)$. The stochastic process X^N defines the random variable $V^N(U, T)$, for any given time horizon T, "target" set U, and population size N.

The first result is stated in terms of the so-called *minimal center of attraction* of the deterministic flow ξ. This set, which we will denote $M(\xi)$, contains all stationary states and all periodic orbits of ξ. Formally, the *minimal center of attraction*

$M(\xi) \subset \square(S)$ of ξ is the closure of the union of the supports of Borel probability measures that are invariant under ξ.[22] The minimal center of attraction is a subset of the so-called Birkhoff center of the flow. The Birkhoff center essentially consists of those states that, as initial states, are passed nearby infinitely many times in the future. Formally, for any state $x \in \square(S)$, let $\omega(x) \subset \square(S)$ be its *omega limit set*, i.e., the set of states $y \in \square(S)$ such that $\lim_{k \to +\infty} \xi(t_k, x) = y$ for some unbounded increasing sequence of times $t_k > 0$. The *Birkhoff center* $B(\xi) \subset \square(S)$ of ξ is the closure of the set of states $x \in \square(S)$ such that $x \in \omega(x)$. By the *Poincaré Recurrence Theorem* (see, for example, Mañé (1987)), $M(\xi) \subset B(\xi)$. It is usually easier to find the Birkhoff center than the minimal center of attraction, and these two sets do not differ much, if at all, in many applications.

The deterministic flow will eventually be close to the minimal center of attraction. Standard results on Markov chains imply that any realization of the stochastic process will eventually settle into some invariant probability measure. For large populations, such a measure has to be close to some invariant measure under the deterministic flow, by Lemma 1:

PROPOSITION 4: *Suppose $X^N(0) \to x$. For any open set $U \subset E$ containing $M(\xi)$:*

$$\lim_{N \to \infty} \left[\liminf_{T \to \infty} V^N(U, T) \right] = 1 \quad a.s.$$

In other words: for large populations the Markov chain almost surely spends almost all time, in the long run, at the minimal center of attraction of the deterministic flow, and hence, a fortiori, at its Birkhoff center.

In many applications, the minimal center of attraction has several disjoint components. In order to get more predictive power, we accordingly focus on *conditional* visitation rates, i.e., visitation rates that are conditional on the event that the stochastic process remains in some pre-specified neighborhood of the component in question. Suppose U and C' are Borel sets, $U \subset C \subset \square(S)$, and $X^N(0) \in C$. The *conditional visitation rate* in U with respect to the superset C is defined as $V^N[U, \tau^N(C)]$.[23] In other words, this is the visitation rate in U until the first exit from C. For any invariant set A in the flow ξ, let $\xi_{|A}$ denote the restriction of ξ to A, and let $M(\xi_{|A})$ denote the minimal center of attraction of the flow $\xi_{|A}$.[24] Note that $M(\xi_{|A}) \subset M(\xi) \cap A$. Likewise, let $B(\xi_{|A}) \supset M(\xi_{|A})$ denote the Birkhoff center of the restricted flow $\xi_{|A}$.

We are now in a position to state our approximation result for conditional visitation rates. The result concerns any open set $U \subset C$ that contains the minimal center of attraction of this restricted flow. The claim is that if the process starts

[22] A Borel probability measure μ is *invariant under* ξ if $\mu(A(t)) = \mu(A)$ for any Borel set A and time t, where $A(t) = \{y \in \square(S) : y = \xi(t, x) \text{ for some } x \in A\}$ is a Borel set by continuity of ξ.

[23] If the first exit time $\tau^N(C)$ is infinite, we define $V^N[U, \tau^N(C)]$ as $\liminf_{T \to \infty} V^N(U, T)$.

[24] For any invariant set $A \subset \square(S)$ under the flow ξ, one may unambiguously define the mapping $\xi_{|A} : \mathbb{R} \times A \to A$ by $\xi_{|A}(t, x) = \xi(t, x)$ for all $t \in \mathbb{R}$ and $x \in A$.

in U, then its conditional visitation rate in U, during its stay in C, is almost surely one as the population size N goes to plus infinity. In other words, as long as the process remains in C, which is typically a very long time (see Proposition 2, and Lemma 3 in the Appendix), it spends almost all time near the minimal center of attraction of the flow restricted to the attractor in question. A fortiori, it spends almost all time near the Birkhoff center of the flow restricted to the attractor. This result can be established along the same lines as the preceding result. However, the conditioning on the exit time from the superset C requires some additional probabilistic machinery. More specifically, we obtain the result by first establishing that every limit point of the sequence $V^N[\cdot, \tau^N(C)]$ is an invariant probability measure with support in A (Proposition 8; see Appendix).

PROPOSITION 5: *Let $A \subset \square(S)$ be an attractor of ξ and let $C \subset B(A)$ be a compact neighborhood of A. If $U \subset C$ is an open neighborhood of $M(\xi_{|A})$ and $X^N(0) \in U$ for all N, then*

$$\lim_{N\to\infty} V^N[U, \tau^N(C)] = 1 \quad a.s.$$

Combined with Lemma 1 and Proposition 2, this result implies that, if the deterministic solution through the initial state of the stochastic process at some point in time enters a basin of attraction, then the stochastic process will not only enter any given neighborhood of that attractor and remain there for a long time if the population is large, but during this time interval, the process will actually spend almost all time at a the minimal center of attraction of the flow restricted to the attractor—a subset of the attractor.

3.4. *Metastability*

Despite the above "positive" approximation results, it is not excluded that the stochastic process eventually stays far away from its deterministic approximation. In particular, even if the process starts in some basin of attraction of the deterministic flow it may with probability one eventually leave this basin and remain outside forever—a phenomenon sometimes called metastability. As the above results show, the time until such an event occurs may be probabilistically so long, for large populations, that this phenomenon is of little practical relevance. However, this remains at least a logical possibility under certain circumstances, which we here identify.

First, suppose the stochastic process has positive "switching" probabilities in the following sense:

$$[\mathrm{C1}]: \qquad 0 < x_{ik} < 1 \;\Rightarrow\; p_{ik}^h(x) > 0 \quad \text{for some} \quad h \neq k.$$

In other words, if some, but not all, individuals in a player population i use pure strategy k, then it is possible that a k-strategist abandons his or her current strategy k. This property is held by many micro choice models of evolution and learning in games.

Secondly, suppose the stochastic process has zero probability of switching to a currently unused strategy:

$$[\text{C2}]: \qquad x_{ih} = 0 \;\Rightarrow\; p^h_{ik}(x) = 0 \quad \text{for all} \quad k \neq h.$$

Stochastic processes with this property arise from micro models based on strategy choice by way of imitation—then no individual ever switches to a currently unused strategy. Such a process evidently never leaves the boundary of the polyhedron $\square(S)$ once this has been reached. Consequently, if both conditions [C1] and [C2] are met, then the stochastic population process sooner or later hits one of the vertices of the polyhedron of mixed-strategy profiles—a pure-strategy profile—and remains there forever. Since the state space is finite, this happens with probability one, irrespective of the initial state, and for any finite population size:

REMARK 2: If condition [C2] is met, then the Markov chain is not ergodic. If both conditions [C1] and [C2] are met, then the chain reaches a vertex of the polyhedron $\square(S)$ in finite time and remains there forever, with probability one.

4. EXAMPLES

We here briefly consider applications to some micro models of boundedly rational strategy choice in recurrently played games, models that have been discussed in the literature on evolution and learning in games. For this purpose, consider a finite game in normal form, with strategy sets as specified in Section 2, and with payoff functions $u_i : \square(S) \to \mathbb{R}$ derived as usual from some pure-strategy payoff functions $\pi_i : S \to \mathbb{R}$, for all player roles $i \in I$. We assume that all individuals in the same player population have the same preferences, given by these payoff functions.[25] We first briefly consider two imitation behaviors and thereafter two best-reply behaviors.

4.1. *Aspiration and Random Imitation*

This first choice model formalizes the decision rule "If you are dissatisfied with the current performance of your strategy, imitate a randomly drawn individual in your own population." More exactly, if the current strategy performs below some aspiration level, then, and only then, does the reviewing individual switch to the strategy of a randomly drawn individual in her own population (this strategy may happen to be the same as her own). Such behaviors are discussed in Gale, Binmore, and Samuelson (1995), Björnerstedt and Weibull (1996), and Binmore

[25] This assumption is standard in evolutionary game theory and facilitates connections with standard solution concepts in noncooperative game theory. However, the present analytical machinery does not require this. An interesting avenue for future research is to consider heterogeneous populations.

and Samuelson (1997). This choice model induces a population process in the class defined in Section 2, and it meets conditions [C1] and [C2].

It is easily shown that if the aspiration levels within each player population i are statistically independent and uniformly distributed on an interval $[a_i(x), b_i(x)]$ that contains the range of the payoff function π_i, then the resulting mean vector field is

$$(11) \qquad F_{ih}(x) = \frac{1}{n}\frac{u_i(e_i^h, x_{-i}) - u_i(x)}{b_i(x) - a_i(x)} x_{ih}.$$

In particular, if the aspiration distributions are state independent, $a_i(x) \equiv \alpha_i$ and $b_i(x) \equiv \beta_i$ for some $\alpha_i < \beta_i$ and for all i, then (11) is but a player-specific time-rescaling of the Taylor (1979) version of the replicator dynamics. If instead each aspiration level follows its population's average payoff in such a way that $a_i(x) \equiv \alpha_i u_i(x)$ and $b_i(x) \equiv \beta_i u_i(x)$ for some $\alpha_i < \beta_i$ (and all payoffs are positive), then we obtain a player-specific time-rescaling of the Maynard Smith (1982) version of the replicator dynamics. In this case it is as if individuals aspire to a multiple of the current average payoff in their population, where the multiplicator is a random variable uniformly distributed on some interval $[\alpha_i, \beta_i]$ (such that $[a_i(x), b_i(x)]$ contains the range of payoffs).

Note the small amount of information needed for this choice model. In particular, a strategy reviewing individual need not even know her own strategy set or payoff function, nor any other aspect of the game, and no information is needed about the population state. What is needed is knowledge of a randomly drawn individual's pure strategy—which in an extensive-form game may mean that the randomly drawn individual actually has to "tell" what his strategy is (what he would have done at unreached information sets).

4.2. *Aspiration and Imitation of Success*

In this choice model, the strategy reviewing individual compares the performance of her strategy with the performance of the strategy of a randomly drawn individual in her own population. The individual switches to the other individual's strategy if the realized payoff difference exceeds a random threshold value. This threshold may be a switching cost or an observational error, or may emanate from idiosyncratic preference differences between individuals in the same player population. Clearly also this choice model generates a population process in the class defined in Section 2, and also in this case conditions [C1] and [C2] are met.

If the threshold distributions are uniform, with a support $[a_i(x), b_i(x)]$ that covers the range of possible payoff differences between any two pure strategies in player role i, then it is easily shown that the mean vector-field is twice that in equation (11), and hence one obtains the Taylor (1979) and Maynard Smith (1982) replicator dynamics as special cases. See Kandori (1996) and Schlag (1998) for similar choice models and observations.

4.3. *Pure and Perturbed Best and Better Replies*

In Kandori, Mailath, and Rob (1993) individuals in the aggregate switch to better or best replies to last period's population state. In order to discuss such behaviors, let the set of pure *best replies* in player role i to any mixed-strategy profile x be denoted $\beta_i(x)$, and let $\gamma_i(x)$ be the set of pure *better replies*.[26] Suppose that the individual drawn for strategy review with probability $1-\varepsilon$ switches to a pure best (better) reply to the current population state, with equal (conditional) probability for all best (better) replies, and otherwise switches to some *other* pure strategy with the remaining probability $\varepsilon \geq 0$. The induced vector field (4) corresponding to the case of best replies becomes

$$F_{ih}(x) = \begin{cases} \dfrac{1-\varepsilon}{|\beta_i(x)|} - x_{ih} & \text{if } h \in \beta_i(x), \\ \dfrac{\varepsilon}{m_i - |\beta_i(x)|} - x_{ih} & \text{otherwise,} \end{cases} \tag{12}$$

and likewise for the case of better replies. Clearly these vector fields are Lipschitz continuous almost everywhere, but have discontinuities on the boundaries of the sets where the pure best-reply (better-reply) correspondences are constant. Consequently, both pure ($\varepsilon = 0$) and perturbed ($0 < \varepsilon < 1$) best-reply (better-reply) behaviors generate population processes that fall slightly outside the class analyzed here. Approximation results for the stochastic processes generated by these behaviors thus call for a generalization of the present analysis.

4.4. *Smooth Best Replies*

Suppose instead that the reviewing individual makes a noisy observation of the current average payoff to each pure strategy in her strategy set, and chooses a pure strategy that has the highest observed value. More precisely, an individual drawn from population i for strategy review observes, for each pure strategy $h \in S_i$, the sum $u_i(e_i^h, x_{-i}) + \varepsilon_{ih}$, where $\{\varepsilon_{ih}\}_h$ are independent and identically distributed according to the extreme-value distribution function $G(\varepsilon_{ih} \leq z) = \exp[-\exp(\sigma z)]$, for $\sigma > 0$. As is well-known in the random-utility discrete choice literature (see, e.g., Anderson, de Palma, and Thisse (1992)), this leads to conditional choice probabilities of the logit form. In the present context this boils down to

$$F_{ih}(x) = \frac{\exp[\sigma u_i(e_i^h, x_{-i})]}{\sum_{k\in S_i} \exp[\sigma u_i(e_i^k, x_{-i})]} - x_{ih}. \tag{13}$$

[26] The pure *better replies* to a mixed-strategy profile x are those pure strategies that earn at least the same payoff as when x is played (see Ritzberger and Weibull (1995) for an analysis of the better-reply correspondence). In order to identify the better or best replies, expected payoffs to pure strategies need to be known. In a finite population, the expected payoff to a strategy is the same as its average payoff when played against all individuals in all other player populations. Such information is assumed in Kandori, Mailath, and Rob (1993).

The induced Markov chain X^N is ergodic. Moreover, as $\sigma \to +\infty$, this vector field converges pointwise to the pure best-reply vector field (wherever this is well defined).

5. LARGE DEVIATIONS

For any player role $i \in I$, pure strategies $h, k \in S_i$, let the "$(i, k \to h)$-*switch*" be the vector $v \in E$ defined by $v_i = e_i^h - e_i^k$ and $v_j = 0$ for all $j \neq i$. Let $V \subset E$ be the (finite) set consisting of all such vectors v. For each vector $v \in V$ and $x \in \Box(S)$, let $\mu_x(v) = p_{ik}^h(x)$. Hence, for each population state x, μ_x is a discrete probability measure on E with support in V. Moreover, in view of equation (3), we have

$$\mu_x(v) = \Pr\left[X^N\left(t + \frac{1}{N}\right) = x + \frac{1}{N}v \,\middle|\, X^N(t) = x\right] \tag{14}$$

for all $v \in V$ and $x \in \Box(S)$. Observe that the mean value of μ_x is the vector field F at x:

$$\sum_{v \in V} v\mu_x(v) = F(x). \tag{15}$$

We will say that the family $\{\mu_x : x \in \Box(S)\}$ is *nondegenerate* if the probability that a strategy reviewing individual switches to any other strategy is always positive. Formally,

$$\text{[C3]} \qquad x_{ik} > 0 \;\Rightarrow\; p_{ik}^h(x) > 0 \quad \text{for all } h.$$

This condition is more stringent than condition [C1], and it is incompatible with condition [C2]. Condition [C3] clearly implies that X^N is an irreducible and aperiodic Markov chain. In particular, for each $N \in \mathbb{N}$ there exists a unique invariant probability measure μ^N such that

$$\lim_{t \to \infty} \Pr[X^N(t) \in B] = \mu^N(B) = \lim_{T \to \infty} V^N(B, T) \tag{16}$$

for every Borel set B in $\Box(S)$.

Moreover, for large N, the process spends almost all time near a particular attractor of the deterministic flow ξ. This attractor can be characterized in terms of the game data and the micro model underlying the stochastic population process. We achieve this characterization by relying on certain results for large deviations. In the present model, the relevant large deviations are finite sequences of many small "jumps" (each of size $1/N$) that take the stochastic process out of one of the deterministic flow's basins of attraction and into another. In order to state this result, some more terminology and notation is needed.

First, for any $T > 0$, let $C[0, T]$ denote the set of continuous functions $\psi : [0, T] \to \Box(S)$, functions that map times t in the bounded interval $[0, T]$ to mixed-strategy profiles $x = \psi(t)$. We will call such functions *paths* and endow this function space with the topology of uniform convergence. For each mixed-strategy profile x, let $C_x[0,T]$ denote the subset of functions $\psi \in C[0,T]$ with $\psi(0)=x$, i.e., paths "starting" at state x. The basic approach in the theory of large deviations is to estimate the probability that the interpolated stochastic process $\widehat{X}^N$, with initial state $X^N(0)\to x$, follows any given path $\psi \in C_x[0,T]$. It turns out that this probability is of the order of magnitude of $\exp[-Nc_{x,T}(\psi)]$, where $c_{x,T}(\psi)$ is a nonnegative number, possibly infinite, to be called the *cost* of the path ψ. In the present context, the *cost* of ψ is defined as follows:

$$c_{x,T}(\psi) = \begin{cases} \int_0^T L[\psi(t), \dot{\psi}(t)]\,dt & \text{if } \psi \text{ is absolutely continuous,} \\ +\infty & \text{otherwise,} \end{cases} \tag{17}$$

where, for each state $x \in \square(S)$ and "direction" $v \in E$, $L(x,v)$ is the so-called *Cramér transform* of μ_x, defined as follows:

$$L(x,v) = \sup_{u\in E}\left[\langle u, v\rangle - \ln\left(\sum_{w\in V} e^{\langle u,w\rangle}\mu_x(w)\right)\right]. \tag{18}$$

While the numerical evaluation of the Cramér transform is complicated in general, the function L is known to have the following qualitative features (see, e.g., Benaïm (1998)):

[P1] For each $x \in \square(S)$, $v \to L(x,v)$ defines a convex and nonnegative function.

[P2] $L(x,v)=0$ iff $v=F(x)$.

[P3] $L(x,v)<\infty$ iff $v\in V(\mu_x)$, where $V(\mu_x)$ is the convex hull of the support of μ_x.[27]

In other words: given any initial state x, $L(x,v)$ defines an "instantaneous cost" at x, associated with every direction v from x, a cost that is a convex function of the direction, with minimum value zero in the direction of the vector field F of the deterministic flow. The "instantaneous cost" is infinite in those directions that have zero probability in the population process X^N (c.f. equation (3)). Equation (17) simply defines the cost $c_{x,T}(\psi)$ of a path ψ as the integral of the instantaneous costs along the path. It follows from the mentioned properties that a path ψ from any given state x has zero cost, $c_{x,T}(\psi)=0$, if and only if ψ is the solution to the mean-field equation (6) with initial condition $\psi(0)=x$. By contrast, if during some subinterval of times $t\in[0,T]$ the tangent vector $\dot{\psi}(t)$ to the path ψ at t falls outside $V(\mu_{\psi(t)})$, i.e., points in zero-probability directions for the stochastic process, then the path ψ has infinite cost.

[27] Formally,

$$V(\mu_x) = \left\{ \sum_{v\in V,\,\mu_x(v)>0} p(v)v : p(v)\geq 0,\ \sum_{v\in V,\,\mu_x(v)>0} p(v)=1 \right\}.$$

For any two sets $A, B \subset \Box(S)$ we define the cost $c(A, B)$ of going from A to B as the least costly path from a point in A to a point in B:

$$c(A, B) = \inf\{c_{x,T}(\psi) : x \in A,\ T > 0,\ \psi \in C_x[0, T],\ \psi(T) \in B\}. \tag{19}$$

Let A be an attractor of the flow ξ, with basin of attraction $B(A)$. Following Ellison (2000), we define the *radius* of A as the lowest cost of going from A to anywhere outside its basin of attraction:

$$R(A) = c[A, \Box(S) \backslash B(A)]. \tag{20}$$

Hence, the radius is a measure of the "cost of escaping" from the attractor. Likewise, the *co-radius* of an attractor is defined as the highest cost to the attractor from anywhere outside its basin of attraction:

$$CR(A) = \sup_{x \notin B(A)} c(\{x\}, A).$$

Evidently, $CR(A) = \sup_{x \notin B(A)} c[\{x\}, B(A)]$. The following proposition establishes that if the deterministic flow has an attractor A that is costlier to leave than to reach, then, for sufficiently large populations, the stochastic population process will spend virtually all time near that attractor, in fact near its minimal center of attraction. This result is similar to Theorem 1 in Ellison (2000). The main difference is that while Ellison keeps the population size fixed and changes the underlying micro model by taking the mutation rate to zero, we keep the underlying micro model fixed and take the population to infinity.

PROPOSITION 6: *Let A be an attractor of ξ, and suppose $U \subset \Box(S)$ is an open neighborhood of $M(\xi_{|A}) \subset A$. If $R(A) > CR(A)$, then*:
(a) $\lim_{N\to\infty} \lim_{T\to\infty} V^N(U, T) = 1$ *a.s.*
(b) $\lim_{N\to\infty} \lim_{T\to\infty} \Pr[X^N(T) \in U] = 1$.

In other words, if the radius exceeds the co-radius of an attractor A of the deterministic flow, then the asymptotic visitation rate to any neighborhood of the minimal center of attraction in A is arbitrarily close to 1, granted the population size is large enough; and likewise for the asymptotic probability that the process will be in such a neighborhood. Since this result cannot hold simultaneously for two disjoint attractors, an implication of the result is that the inequality $R(A) > CR(A)$ can hold for at most one minimal attractor, that is, an attractor that does not properly contain another attractor.

The numerical evaluation of the radius and co-radius of a given attractor is in general difficult, including computationally demanding applications of variational calculus. A challenging line of future research is thus to provide numerical estimates of these quantities. For work in this direction, the following proposition, which follows directly from Theorem 4.3 in Freidlin and Wentzell (1984), may be useful. Let $H : E^2 \to \mathbb{R}$ be defined by the second term in the above definition of the Cramér transform,

$$H(x, u) = \ln\left[\sum_{v \in V} e^{\langle u, v\rangle} \mu_x(v)\right]. \tag{21}$$

PROPOSITION 7: *Let A be an attractor of ξ. Suppose g is a real-valued and continuous function defined on $\overline{B(A)}$, such that g vanishes on A but is positive and continuously differentiable on $B(A)\backslash A$, with $\nabla g(x) \neq 0$ and $H[x, \nabla g(x)] = 0$ for all such x. Then $g(x) = c(A, \{x\})$ for all $x \in \overline{B(A)}$, and thus*

$$R(A) = \inf_{x \in \partial B(A)} g(x).$$

Applied to symmetric 2×2 coordination games, this approach can be used to generate results similar to those in Binmore and Samuelson (1997); see Benaïm and Weibull (2000, revised 2001) for details.

6. POISSON CLOCKS

The stochastic processes studied in this paper have exactly one individual drawn for strategy review at distinct deterministic times, separated by a fixed time interval of length $\delta = 1/N$. It would be more natural to assume that these review times instead are random. Is the analysis robust in this respect? The canonical continuous-time model of random "arrival times" is that of a Poisson process. It is not difficult to verify that all the qualitative results obtained in this paper remain valid if we replace the discrete-time process X^N by a continuous-time process Y^N whose transition times are generated by a Poisson process with constant intensity. More exactly, we may replace the Markov chain X^N by any Markov process Y^N such that, at any time $t \in \mathbb{R}$ and for all $\tau > 0$,

$$\Pr\left[Y^N(t+\tau) = x + \frac{1}{N} v \,\middle|\, Y^N(t)\right] = x = N\mu_x(v)\tau + o(\tau), \tag{22}$$

where $\mu_x(v)$ is the discrete probability measure defined in equation (14) (see also equation (3)).

Assume, for instance, that each individual has a "Poisson clock," with constant intensity $1/n$, all clocks being statistically independent. Each time an individual's clock "rings," the individual reviews her strategy choice. To see that this results in an equation of the above form, suppose that each individual η, where $\eta \in \mathcal{M} = \{1, 2, \ldots, nN\}$, reviews her strategy choice at random times $0 = T_0(\eta) < T_1(\eta) < T_2(\eta) < \cdots$, where the random variables $T_k(\eta) - T_{k-1}(\eta)$, for $\eta \in \mathcal{M}$ and $k \in \mathbb{N}$, are i.i.d. exponentially distributed with mean value n. When an individual η in any population $i \in I$ is given the opportunity to revise her strategy choice, she switches from her current strategy $k \in S_i$ to strategy $h \in S_h$ with some conditional probability $q_{ik}^h(x)$. Thus,

$$\Pr\left[Y_i^N(t+\tau) = x + \frac{1}{N}\left(e_i^h - e_i^k\right) \,\middle|\, Y^N(t) = x\right] = N\frac{x_{ik}}{n} q_{ik}^h(x)\tau + o(\tau). \tag{23}$$

We explain in the Appendix how Lemma 1 can be adapted to such a continuous-time process, while proofs of the other results are left to the reader.

Dept. of Mathematics, Université de Cergy-Pontoise, 2 Avenue Adolphe Chauvin, Cedex, F-95 302 Cergy-Pontoise, France,

and

Dept. of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, U.S.A.

Manuscript received June, 2000; final revision received June, 2002.

APPENDIX

Let λ be the Lipschitz constant of F on the compact set $\square(S) \subset E$, with respect to the L^∞-norm, let $\|\cdot\|_2$ denote the L^2-norm, and let $\|F\|_2$ be the maximum of $\|F(x)\|_2$ on $\square(S)$.[28]

Let U_n, for $n \in \mathbb{N}$, be the difference between the step taken by the Markov chain X^N between periods n and $n+1$, per time unit, and the vector field at the state:

$$U_n = \frac{1}{\delta}[X^N((n+1)\delta) - X^N(n\delta)] - F(X^N(n\delta)), \tag{24}$$

where $\delta = 1/N$ is the length of a period. Let $\mathcal{F}_k$, $k \in \mathbb{N}$ denote the sigma-field generated by $\{X^N(0), \dots, X^N(k\delta)\}$. The following result, giving an upper bound on the difference U_k, turns out to be useful for the proof of Lemma 1:

LEMMA 2: *Let $\Gamma = (\sqrt{2} + \|F\|_2)^2$. For any $\theta \in \mathbb{R}^m$:*

$$E\left(e^{\langle \theta, U_k \rangle} | \mathcal{F}_k\right) \le e^{\Gamma \|\theta\|_2^2/2}.$$

PROOF: By definition of U_k it is easy to verify that

$$\|U_k\|_2 \le \max_{i,h,k} \|e_i^h - e_i^k\|_2 + \|F\|_2 = \sqrt{\Gamma}. \tag{25}$$

Let $g(t) = \log E(e^{t\langle \theta, U_k \rangle} | \mathcal{F}_k)$. The function $g : \mathbb{R} \to \mathbb{R}$ is convex and satisfies $g(0) = g'(0) = 0$, $g''(t) \le \|\theta\|_2^2 \Gamma$. Therefore $g(1) \le \|\theta\|_2^2 \Gamma/2$. *Q.E.D.*

A. *Proof of Lemma* 1

Let $U : \mathbb{R}_+ \to E$ be the map defined by $U(t) = U_k$ for $k\delta \le t < (k+1)\delta$. Likewise, let $\overline{X}^N$ be the continuous-time (right-continuous) *step process* generated by the Markov chain X^N : $\overline{X}^N(t)$ is defined for all $t \in \mathbb{R}_+$ by $\overline{X}^N(t) = X^N(k\delta)$ for $k\delta \le t < (k+1)\delta$. Suppose that $X^N(0) = x \in \square(S)$. Then

$$\begin{aligned} \widehat{X}^N(t) - x &= \int_0^t [F(\overline{X}^N(s)) + U(s)]\, ds \\ &= \int_0^t [F(\widehat{X}^N(s)) + F(\overline{X}^N(s)) - F(\widehat{X}^N(s)) + U(s)]\, ds. \end{aligned} \tag{26}$$

Since $\xi(t,x) - x = \int_0^t F(\xi(s,x))\, ds$, we obtain

$$\|\widehat{X}^N(t) - \xi(t,x)\|_\infty \le \lambda\left[\int_0^t (\|\widehat{X}^N(s) - \xi(s,x)\|_\infty)\, ds + 2\delta T\right] + \Psi(T), \tag{27}$$

[28] The Lipschitz constant of F on the compact set $\square(S)$ is $\lambda = \lambda_C$, for $C = \square(S)$; see footnote 13.

where

$$\Psi(T) = \max_{0\le t\le T} \left\| \int_0^t U(s)\,ds \right\|_\infty . \tag{28}$$

Grönwall's inequality implies

$$D^N(T,x) = \max_{0\le t\le T} \|\widehat{X}^N(t) - \xi(t,x)\|_\infty \le [\Psi(T) + 2\delta\lambda T]e^{\lambda T}. \tag{29}$$

Thus, for $\delta \le (\varepsilon/4\lambda T)e^{-\lambda T}$,

$$\Pr[D^N(T,x) \ge \varepsilon] \le \Pr\left[\Psi(T) \ge \frac{\varepsilon}{2} e^{-\lambda T}\right]. \tag{30}$$

Our next goal is to estimate the probability on the right-hand side. For $k \in \mathbb{N}$, let

$$Z_k(\theta) = \exp\left(\sum_{i=0}^{k-1} \langle \theta,\ \delta U_i\rangle - \frac{\Gamma}{2} k\delta^2 \|\theta\|_2^2\right). \tag{31}$$

According to Lemma 3, $(Z_k(\theta))_{k\in\mathbb{N}}$ is a supermartingale. Thus, for any $\beta > 0$,

$$\begin{aligned} \Pr\left[\max_{0\le k\le n}\left\langle \theta, \sum_{i=0}^{k-1} \delta U_i\right\rangle \ge \beta\right] &\le \Pr\left[\max_{0\le k\le n} Z_k(\theta) \ge \exp\left(\beta - \frac{\Gamma}{2}\|\theta\|_2^2 n\delta^2\right)\right] \\ &\le \exp\left(\frac{\Gamma}{2}\|\theta\|_2^2 n\delta^2 - \beta\right). \end{aligned} \tag{32}$$

Let $u_1, \dots, u_m$ be the canonical basis of $E = \mathbb{R}^m$, $\varepsilon > 0$, and $u = \pm u_i$ for some i. Set $\beta = \varepsilon^2/(\Gamma n\delta^2)$ and $\theta = (\beta/\varepsilon)u$. Then

$$\begin{aligned} \Pr\left[\max_{0\le k\le n}\left\langle u, \sum_{i=0}^{k-1} \delta U_i\right\rangle \ge \varepsilon\right] &= \Pr\left[\max_{0\le k\le n}\left\langle \theta, \sum_{i=0}^{k-1} \delta U_i\right\rangle \ge \beta\right] \\ &\le \exp\left(\frac{-\varepsilon^2}{2\Gamma n\delta^2}\right). \end{aligned} \tag{33}$$

It follows that

$$\Pr[\Psi(T) \ge \varepsilon] \le 2m \exp\left(\frac{-\varepsilon^2}{2\delta\Gamma T}\right). \tag{34}$$

Therefore,

$$\Pr\left[\Psi(T) \ge \frac{\varepsilon}{2} e^{-\lambda T}\right] \le 2m \exp\left(-\varepsilon^2 \frac{e^{-2\lambda T}}{8\delta\Gamma T}\right), \tag{35}$$

and the claimed inequality follows, with

$$c = \frac{e^{-2\lambda T}}{8T\sqrt{\sqrt{2} + \|F\|_2^2}}. \tag{36}$$

B. *Construction of a Common Probability Space*

Let $I = (0, 1)$ and $\Omega = I^{\mathbb{N}}$. A point in Ω is denoted $\omega = (\omega_i)_{i\in\mathbb{N}}$. Let λ be the uniform probability on $(0, 1)$, and let $\mathbb{P} = \lambda^{\mathbb{N}}$, that is, for any finite collection of Borel sets $A_1, \ldots, A_n \subset I : \mathbb{P}[\omega \in \Omega : \omega_1 \in A_1, \ldots, \omega_n \in A_n] = \prod_{i=1}^{n} \lambda(A_i)$. Let $B(\Omega)$ be the associated product σ-field. We proceed to show that any finite Markov chain can be constructed on the probability space $(\Omega, B(\Omega), \mathbb{P})$.

Let $M = \{1, \ldots, d\}$ be any finite set and $P = (p_{ij})$ a Markov transition matrix over M (that is, all p_{ij} are nonnegative, and $\sum_j p_{ij} = 1$ for all $i \in M$). For each $i \in M$, partition I according to the probabilities $p_{i1}, \ldots, p_{id}$: let $q_{i0} = 0$, and for all $j \in M$ let $q_{ij} = q_{ij-1} + p_{ij}$ (thus $q_{id} = 1$). For each $i \in M$, let $f_i : I \to M$ be the associated "state indicator" function, that is, $f_i(\theta) = j$ if $q_{ij-1} < \theta \leq q_{ij}$. Using this function, one obtains a Markov chain X on $(\Omega, B(\Omega), \mathbb{P})$ with the given transition matrix P, and with any initial state $x \in M$, by letting $X : \Omega \to M^{\mathbb{N}}$ be defined recursively by $X_0(\omega) = x$ and

$$X_{t+1}(\omega) = f_i(\omega_{t+1}) \quad \text{where} \quad i = X_t(\omega).$$

C. *Proof of Proposition* 1

Let $C = \square(S)\backslash U$, a compact subset of $\square(S)$, disjoint from $\overline{\gamma^+(x)}$. Set $\varepsilon = d(C, \overline{\gamma^+(x)})$, where $d(\cdot,\cdot)$ is the Hausdorff metric, and let $t > 0$. By continuity of the flow ξ there exists a positive integer N^o such that for all $N > N^o$

$$\sup_{0\leq s\leq t} \|\xi[s, X^N(0)] - \xi(s, x)\| < \varepsilon/2. \tag{37}$$

Then $\{\tau^N(U) \leq t\} \subset \{D^N[t, X^N(0)] \geq \varepsilon/2\}$, for all $N > N^o$. Thus, by Lemma 1:

$$\sum_{N=N^o}^{\infty} \Pr[\tau^N(U) \leq t] \leq 2m \sum_{N=N^o}^{\infty} e^{-\varepsilon^2 cN/4} < \infty. \tag{38}$$

Hence by the Borel-Cantelli Lemma, the event $\{\tau^N(U) \leq t$ for infinitely many $N\}$ has zero probability.

D. *Proof of Proposition* 2

Proposition 2 follows from claim (c) of the following result; see Remark 3 below. Claim (a) is used in the proof of claims (b) and (c), as well as in the proof of Proposition 5.

LEMMA 3: *Let $A \subset \square(S)$ be an attractor for the flow ξ, let $C \subset B(A)$ be compact, and suppose $X^N(0) \in C$ for all N. Then there exists an open neighborhood U of $A \cup C$ with closure $\overline{U}$ in $B(A)$, and a scalar $\alpha > 0$, such that (with $m = M - n$):*
(a) $\Pr[\tau^N(U) \leq t] \leq 2m(t+1)e^{-\alpha N} \quad \forall t \geq 0, N \in \mathbb{N}$.
(b) $E[\tau^N(U)] \geq (1/4m)e^{\alpha N} - 1 \quad \forall N \in \mathbb{N}$.
(c) $\liminf_{N\to\infty}[(1/N)\ln \tau^N(U)] \geq \alpha \quad$ *a.s.*

PROOF: Since A is an attractor, it is possible to find an open neighborhood U of $A \cup C$, having compact closure $\overline{U} \subset B(A)$, such that

$$\xi(t, \overline{U}) \subset U \subset \overline{U} \subset B(A) \tag{39}$$

for all $t > 0$. Now fix $\varepsilon > 0$ small enough so that

$$N_\varepsilon(\xi(1, \overline{U})) \subset U \subset N_\varepsilon(U) \subset B(A) \tag{40}$$

where $N_\varepsilon(U)$ denotes the ε-neighborhood of the set U. Let t be a positive integer, and let

$$D_t^N = \sup_{0\leq k\leq t-1} D^N(1, \widehat{X}^N(k)). \tag{41}$$

For N large enough, $\widehat{X}^N(0) \in U$. Therefore, $D_t^N < \varepsilon$ implies $\tau^N(U) > t$. Hence

$$\begin{aligned}(42)\qquad \Pr[\tau^N(U) \le t] \le \Pr\big(D_t^N \ge \varepsilon\big) &\le \sum_{k=0}^{t-1} \Pr[D^N(1, \widehat{X}^N(k)) \ge \varepsilon] \\ &= \sum_{k=0}^{t-1} E(\Pr[D^N(1, \widehat{X}^N(k)) \ge \varepsilon | \widehat{X}^N(k)]) \\ &\le 2mt \exp[-\varepsilon^2 cN],\end{aligned}$$

where the last inequality follows from Lemma 1. If $t \in \mathbb{R}_+$, then

$$(43)\qquad \Pr[\tau^N(U) \le t] \le \Pr[\tau^N(U) \le \lfloor t \rfloor + 1] \le 2m(t+1)\exp[-\varepsilon^2 cN],$$

where $\lfloor t \rfloor$ is the largest integer not exceeding t. This proves assertion (a), for $\alpha = \varepsilon^2 c$.

To prove assertion (b) from (a), we use the fact that

$$(44)\qquad E[\tau^N(U)] = \int_0^\infty \Pr(\tau^N(U) > t)\, dt \ge \int_0^\infty \max\{0, 1 - a(t+1)\}\, dt,$$

where $a = 2m\exp[-\alpha N]$. Therefore,

$$(45)\qquad E[\tau^N(U)] \ge \int_0^{\frac{1-a}{a}} [1 - a(t+1)]\, dt = \frac{(1-a)^2}{2a} \ge \frac{1}{2a} - 1,$$

which gives (b).

Turning to assertion (c), finally, let $\beta > 0$. From assertion (a) we obtain, for N sufficiently large,

$$(46)\qquad \Pr(\tau^N(U) \le \exp[(\alpha - \beta)N]) \le 4m\exp(-\beta N)$$

or, equivalently,

$$(47)\qquad \Pr\left[\frac{1}{N}\ln \tau^N(U) \le \alpha - \beta\right] \le 4m\exp(-\beta N).$$

Hence, by the Borel-Cantelli Lemma,

$$(48)\qquad \liminf_{N\to\infty}\left[\frac{1}{N}\log \tau^N(U)\right] \ge \alpha - \beta \quad \text{a.s.}$$

Since this inequality holds for all $\beta > 0$, it also holds for $\beta = 0$.

REMARK 3: Let $V \subset B(A)$ be any neighborhood of A. Then there exists a neighborhood $U \subset V$ of A with compact closure in $B(A)$ such that the conclusions of Lemma 3 hold provided $X^N(0) \in \overline{U}$. The proof is exactly the same since it is always possible to find such a U satisfying equation (39).

E. *Proof of Proposition* 3

Let U_ε be an ε-neighborhood of U. For $\varepsilon > 0$ small enough, every point in U' leaves U_ε. More exactly, for all $x' \in U'$ there exists a time $T > 0$ such that $\xi(T, x') \in \partial U_\varepsilon$. Hence, for $X^N(0) = x'$, $\tau^N(U) = +\infty$ implies $D^N(T, x') \ge \varepsilon$. By Lemma 1, this gives

$$(49)\qquad \Pr[\tau^N(U) = +\infty] \le 2me^{-\varepsilon^2 cN},$$

and the claim follows by the Borel-Cantelli Lemma.

F. *Proof of Proposition* 4

By continuity of p_{ih}^k the Markov chain X^N is Feller.[29] Let $C \subset E$ be a compact set disjoint from $M(\xi)$. Since $\square(S)$ is compact, the sequence of probability measures $\langle V^N(\cdot, k\delta)\rangle_{k\in\mathbb{N}}$ is relatively compact, in the topology of weak convergence. By a standard result for Markov chains, every limit point of $\langle V^N(\cdot, k\delta)\rangle_{k\in\mathbb{N}}$ is almost surely an invariant measure of the chain X^N (see, e.g., Lemma 1.IV.21 in Duflo (1996)). In other words, some subsequence of the sequence $\langle V^N(\cdot, k\delta)\rangle_{k\in\mathbb{N}}$ converges weakly to some invariant measure μ_N.[30] Let μ_N be such a measure. Our averaging result, Lemma 1, implies also that the limit points of the sequence $(\mu_N)_{N\in\mathbb{N}}$ are invariant probability measures under the flow ξ (see, e.g., Benaïm (1998, Corollary 3.2)). For any $x \in \square(S)$, let $d(x, C)$ be the (Hausdorff) distance from x to the set C. For any $\varepsilon > 0$, let $f_\varepsilon : \square(S) \to \mathbb{R}$ be defined by $f_\varepsilon(x) = \max\{0, 1 - d(x, C)/\varepsilon\}$. Clearly f_ε is continuous with $f_\varepsilon(x) = 1$ if $x \in C$ and $f_\varepsilon(x) = 0$ if $d(x, C) \geq \varepsilon$. Let $\langle N_k\rangle_{k\in\mathbb{N}}$ be an unbounded increasing sequence such that the associated subsequence $\langle \mu_{N_k}\rangle_{k\in\mathbb{N}}$ converges weakly, and

$$\lim_{k\to\infty} \mu_{N_k}(C) = \limsup_{N\to\infty} \mu_N(C). \tag{50}$$

Then

$$\lim_{k\to\infty} \mu_{N_k}(C) \leq \lim_{k\to\infty} \int f_\varepsilon(x)\mu_{N_k}(dx) = \int f_\varepsilon(x)\mu(dx) \tag{51}$$

where μ is a probability measure that is invariant under ξ. Since $C \cap M(\xi) = \varnothing$, $\int f_\varepsilon(x)\mu(dx) = 0$ for $\varepsilon > 0$ small enough. Therefore, $\limsup_{N\to\infty} \mu_N(C) = 0$, and thus $\lim_{N\to\infty} \mu_N(C) = 0$. Since this holds for any invariant measure μ_N to which some subsequence of $\langle V^N(\cdot, k\delta)\rangle_{k\in\mathbb{N}}$ converges weakly, we have

$$\lim_{N\to\infty}\Big[\limsup_{t\to\infty} V^N(C, t)\Big] = 0 \quad \text{a.s.} \tag{52}$$

Hence, if $U \subset E$ is an open set containing $M(\xi)$, then $C = \square(S)\backslash U$ is a compact set disjoint from $M(\xi)$, and the claimed result holds.

G. *Proof of Proposition* 5

For all $N \in \mathbb{N}$, let $T^N = \tau^N(C)$ and $a_N = N$ in Proposition 8 below. Then condition (53) is met, and, by Lemma 3 (a), so is condition (54).

PROPOSITION 8: *Let $\langle T^N\rangle$ be a sequence of nonnegative finite random variables. Assume that there exists a* (*deterministic*) *sequence $\langle a_N\rangle$ of positive real numbers such that* (53) *and* (54) *below hold. Then the limit points of $V^N(\cdot, T^N)$, in the weak* topology, are almost surely invariant probability measures under ξ.*

$$\lim_{N\to\infty} \frac{a_N}{\log N} = +\infty; \tag{53}$$

$$\sum_N \Pr(T^N \leq t a_N) < +\infty \quad \textit{for all } t \geq 0. \tag{54}$$

[29] The process X^N is said to be *Feller* if for all continuous functions $f : \square(S) \to \mathbb{R}$, the following function $P(f) : \square(S) \to \mathbb{R}$ is continuous:

$$[P(f)](x) = E[f(X^N(\delta)) \mid X^N(0) = x].$$

[30] Let $\mathcal{P}$ denote the space of (Borel) probability measures on $\square(S)$. For $\mu \in \mathcal{P}$ and $f : \square(S) \to \mathbb{R}$ in $L^1(\mu)$ write $\mu(f) = \int f\, d\mu$. A sequence $\langle \mu_n\rangle$ of such measures is said to *converge weakly* to μ if $\lim_{n\to\infty} \mu_n(f) = \mu(f)$ for every continuous function $f : \square(S) \to \mathbb{R}$. The space $\mathcal{P}$ endowed with this topology is a compact metric space.

Hence, every limit point of $\langle V^N[\cdot,\tau^N(C)]\rangle_N$ is a measure μ that is invariant under ξ and whose support is contained in the set C in the statement of Proposition 5. The support of an ξ-invariant measure being an invariant set, μ is supported by A. Let U be an open neighborhood of $M(\xi_{|A})$. It follows that every limit point μ of $\langle V^N[\cdot,\tau^N(C)]\rangle_N$ satisfies $\mu(U)=1$. This implies the claim in Proposition 5 (for details, see the last argument in the above proof of Proposition 4). It thus remains to prove Proposition 8. The key step for this is the following lemma.

Let $f:\square(S)\to\mathbb{R}$ be a Lipschitz continuous function with max-norm not exceeding one; $\|f\|=\sup_{x\in\square(S)}\|f(x)\|\leq 1$. Let $\mathrm{Lip}(f)$ be its Lipschitz constant. It is convenient here to write the flow ξ in the form $\xi_t(x)=\xi(t,x)$.

LEMMA 4: *Let $\langle T^N\rangle$ be a sequence of nonnegative finite random variables with the properties assumed in Proposition* 8. *For all $t\geq 0$:*

$$\lim_{N\to\infty}\frac{1}{T^N}\int_0^{T^N}\big[f(\widehat{X}^N(s))-f\circ\xi_t(\widehat{X}^N(s))\big]\,ds=0\quad a.s.$$

PROOF: Fix $t>0$. For every positive integer k set

$$W_k=\int_{kt}^{(k+1)t}\big[f(\widehat{X}^N(s+t))-f\circ\xi_t(\widehat{X}^N(s))\big]\,ds \tag{55}$$

and set $K^N=\lceil T^N/t\rceil-1$ where $\lceil x\rceil$ is the integer part of a real number x. With this notation,

$$\begin{aligned}\frac{1}{T^N}\int_0^{T^N}\big[f(\widehat{X}^N(s))-f\circ\xi_t(\widehat{X}^N(s))\big]\,ds&=\frac{1}{T^N}\int_0^{T^N}\big[f(\widehat{X}^N(s))-f(\widehat{X}^N(t+s))\big]\,ds+\sum_{k=0}^{K^N}W_k\\&\quad+\int_{\lceil T^N/t\rceil t}^{T^N}\big[f(\widehat{X}^N(s+t))-f\circ\xi_t(\widehat{X}^N(s))\big]\,ds\\&\leq\frac{1}{T^N}\left(2t+\sum_{k=0}^{K^N}W_k+2t\right).\end{aligned}\tag{56}$$

By the Borel-Cantelli Lemma, $T^N\to+\infty$ almost surely. Therefore, it suffices to prove that

$$\lim_{N\to\infty}\frac{1}{T^N}\sum_{k=0}^{K^N}W_k=0\quad\text{a.s.}\tag{57}$$

To save on notation, we write $\mathcal{G}_s=\mathcal{F}_{\lceil Ns\rceil}$ where we recall that $\mathcal{F}_n$ is the sigma-field generated by $\{X^N(0),\dots,X^N(n/N)\}$. Let $V_k=E(W_k|\mathcal{G}_{kt})$ and $U_k=W_k-V_k$. We then have

$$V_k=\int_{kt}^{(k+1)t}E\big[E\big(f[\widehat{X}^N(s+t)]-f\circ\xi_t[\widehat{X}^N(s)]|\mathcal{G}_s\big)|\mathcal{G}_{kt}\big]\,ds\tag{58}$$

and, according to Lemma 1,

$$\begin{aligned}E\big(f[\widehat{X}^N(s+t)]-f\circ\xi_t[\widehat{X}^N(s)]|\mathcal{G}_s\big)&\leq\mathrm{Lip}(f)\cdot E\big(D^N[t,\widehat{X}^N(s)]|\mathcal{G}_s\big)\\&\leq\mathrm{Lip}(f)\big[E\big(D^N[t,\widehat{X}^N(s)]\cdot\mathbf{1}_{\{D^N[t,\widehat{X}^N(s)]\geq\varepsilon\}}|\mathcal{G}_s\big)+\varepsilon\big]\\&\leq\mathrm{Lip}(f)[2m|\square(S)|\exp(-\varepsilon^2cN)+\varepsilon]\end{aligned}\tag{59}$$

where $|\square(S)|$ is the diameter of $\square(S)$ (the largest distance between any two points in $\square(S)$). It then follows that

$$\frac{1}{T^N}\left|\sum_{k=0}^{K_N}V_k\right|\leq\frac{1}{t}\mathrm{Lip}(f)(2m|\square(S)|\exp[-\varepsilon^2cN]+\varepsilon).\tag{60}$$

Since $\varepsilon > 0$ is arbitrary, this implies that

$$\lim_{N\to\infty} \frac{1}{T^N} \sum_{k=0}^{K_N} V_k = 0 \tag{61}$$

with probability one. It remains to prove that

$$\lim_{N\to\infty} \frac{1}{T^N} \sum_{k=0}^{K^N} U_k = 0 \quad \text{a.s.} \tag{62}$$

By definition, U_k is measurable with respect to $\mathcal{G}_{(k+1)t}$ and satisfies $E(U_k|\mathcal{G}_{kt}) = 0$. Observe also that

$$|U_{k+1}| \le |W_k| + |V_k| \le 4t. \tag{63}$$

Therefore, exactly as in the proof of inequality (34) in the proof of Lemma 1, we obtain

$$\Pr\left(\frac{1}{n}\left|\sum_{i=0}^{n-1} U_i\right| \ge \varepsilon\right) \le 2\exp\left[\frac{-\varepsilon^2 n}{2\Gamma(t)}\right], \tag{64}$$

where $\Gamma(t) = (4t)^2$. From this estimate, we deduce that for all integers m

$$\begin{aligned} \Pr\left(\left|\frac{1}{K^N}\sum_{k=0}^{K^N} U_k\right| \ge \varepsilon\right) &\le \Pr(K^N \le m) + \sum_{n>m} \Pr\left(\left|\frac{1}{n}\sum_{k=0}^{n} U_k\right| \ge \varepsilon\right) \\ &\le \Pr(K^N \le m) + \sum_{n>m} \exp\left[\frac{-\varepsilon^2 n}{2\Gamma(t)}\right] \\ &\le \Pr[T^N \le t(m+2)] + \mathcal{O}\left(\exp\left[\frac{-\varepsilon^2 m}{2\Gamma(t)}\right]\right). \end{aligned} \tag{65}$$

Now choose $m = \lceil a_N + 2 \rceil$. Our assumptions on T^N and a_N imply that

$$\sum_{N=1}^{+\infty} \Pr\left(\left|\frac{1}{T^N}\sum_{k=0}^{K^N} U_k\right| \ge \varepsilon\right) < +\infty \tag{66}$$

and we obtain the conclusion by the Borel-Cantelli Lemma. *Q.E.D.*

To conclude the proof of Proposition 8 we finally use the fact that the space $\mathcal{P}$ of probability measures on $\square(S)$ is separable in the topology of *weak* convergence*. More precisely, there exists a countable family of Lipschitz continuous functions $f_i : \square(S) \to \mathbb{R}$, for $i \in \mathbb{N}$, with $\|f_i\| \le 1$, such that for every sequence $\langle \mu_n \rangle$ of probability measures μ_n on $\square(S)$, $\mu_n \to \mu$ in the weak* topology if and only if

$$\lim_{n\to\infty} \mu_n(f_i) = \mu(f_i) \tag{67}$$

for all $i \in \mathbb{N}$.

According to the above lemma, the event

$$\Omega_{i,t} = \left\{\lim_{N\to\infty} |V^N(f_i, T^N) - V^N(f_i \circ \xi_t, T^N)| = 0\right\} \tag{68}$$

has probability one, where $V^N(f_i, T^N)$ is the integral of f_i with respect to the measure $V^N(\cdot, T^N)$ defined in equation (10). Therefore also the intersection of these events, the set

$$\Omega' = \bigcap_{i\in N,\, t\in\mathbb{Q}_+} \Omega_{i,t}, \tag{69}$$

has probability one. On Ω' every limit point μ of $V^N(\cdot, T^N)$ satisfies $\mu(f_i) = \mu(f_i \circ \xi_t)$ for all $i \in \mathbb{N}$ and $t \in \mathbb{Q}_+$. Therefore $\mu(f_i) = \mu(f_i \circ \xi_t)$ for all $i \in \mathbb{N}$ and $t \in \mathbb{R}_+$.

H. *Proof of Proposition* 6

Since the chain is irreducible and aperiodic,

$$\mu^N(U) = \lim_{T\to\infty} V^N(U, T) = \lim_{T\to\infty} \Pr[X^N(T) \in U]. \tag{70}$$

It is thus sufficient to prove that $\lim_{N\to\infty} \mu^N(U) = 1$. The proof, sketched below, follows the ideas of Benaïm (1998) and Benaïm and Hirsch (1999a). However, it is far from straight-forward to deduce Proposition 6 from those papers because they do not use the notions of radius and coradius. More exactly, Benaïm (1998) proves a result similar to Proposition 6, but with U a neighborhood of the complement to a component of the "chain recurrent set," where the component meets a certain topological condition. Likewise, Benaïm and Hirsch (1999a) establish that Proposition 6 holds with U a neighborhood of the set of linearly stable equilibria, under the assumption that the mean vector field is "cooperative" and "irreducible."

The proof sketch for Proposition 6 runs as follows: Let $D = R(A) - CR(A) > 0$. For all $\varepsilon \in (0, D)$ we can find disjoint neighborhoods V_1 of $\square(S)\backslash B(A)$ and V_2 of A such that for all x in V_1 and y in V_2:

$$c(\{x\}, V_1) - c(\{y\}, V_2) \geq D - \varepsilon > 0. \tag{71}$$

Let Y^N denote the induced chain on $V = V_1 \cup V_2$ whose transition probabilities are defined by

$$\widetilde{P}_y^N(U) = \Pr[Y^N(t+1/N) \in U \mid Y^N(t) = y] = \Pr[X^N(T^N) \in U \mid X^N(0) = y], \tag{72}$$

where

$$T^N = \inf\{t \in \mathbb{T} : t > 0 : X^N(t) \in V\} \tag{73}$$

and U is any Borel subset of V.

Using (71), it can easily be shown that for some $\delta > \eta > 0$ with $\delta - \eta \approx D - \varepsilon$, and some $\alpha(N) \in \mathbb{T}$:

$$\liminf_{N\to\infty} \frac{1}{N} \ln \widetilde{P}_x^N(Y^N[\alpha(N)] \in V_2) \geq -\eta \tag{74}$$

uniformly in $x \in V_1$, and

$$\limsup_{N\to\infty} \frac{1}{N} \ln \widetilde{P}_y^N(Y^N[\alpha(N)] \in V_1) \geq -\delta \tag{75}$$

uniformly in $y \in V_2$.

The end of the proof is classical: For $i \neq j \in \{1, 2\}$ define

$$m_{ij}^N = \frac{1}{\tilde{\mu}^N(V_i)} \int_{V_j} \widetilde{P}_z^N(Y^N[\alpha(N)] \in V_j) \tilde{\mu}^N(dz), \tag{76}$$

where $\tilde{\mu}^N(\cdot) = \mu^N(\cdot)/\mu^N(V)$. By invariance of μ^N, the vector $(\tilde{\mu}^N(V_1), \tilde{\mu}^N(V_2))$ satisfies

$$m_{21}^N \tilde{\mu}^N(V_2) = m_{12}^N \tilde{\mu}^N(V_1). \tag{77}$$

Thus,

$$\tilde{\mu}^N(V_1)/\tilde{\mu}^N(V_2) \leq Ce^{(\eta-\delta)N} \tag{78}$$

for some constant C.

I. *Proof of Lemma* 1 *for Poisson Processes*

Let L^N denote the infinitesimal generator of the continuous-time process Y^N in Section 6. That is,

$$L^N(f)(x) = \lim_{t\to 0} \frac{E[Y^N(t)|Y^N(0)=x]-f(x)}{t}$$

defined for every real valued continuous function on $\square(S)$. Then

$$L^N(f)(x) = N\sum_{v\in V}\left(f\left(x+\frac{1}{N}v\right)-f(x)\right)\mu_x(v).$$

Let $f: E\to R$ be the map defined by $f(y) = \langle\theta, y-x\rangle$. By standard results in the theory of Markov processes, the process

$$f(Y^N(t))\exp\left[-\int_0^t \frac{L^N(f)(Y^N(t))}{f(Y^N(t))}\,dt\right]$$

is a martingale. Set $g(u) = e^u - u + 1$. Then

$$\frac{L^N(f)(y)}{f(y)} = \langle F(y),\theta\rangle + N\sum_{v\in V} g\left(\frac{1}{N}\langle v,\theta\rangle\right)\mu_y(v).$$

From this expression it is not hard to deduce that

$$\frac{L^N(f)(y)}{f(y)} - \langle F(y),\theta\rangle \le \frac{1}{N}\Gamma\|\theta\|_2^2$$

for some constant Γ. This makes the process

$$Z(t) = \exp\left\langle\theta, Y^N(t)-x-\int_0^t F(Y^N(t))\,dt - t\frac{1}{N}\Gamma\|\theta\|_2^2\right\rangle$$

a supermartingale, and we conclude as in the proof of Lemma 1.

REFERENCES

Anderson, S., A. de Palma, and J. Thisse (1992): *Discrete Choice Theory of Product Differentiation*. Cambridge, U.S.A.: MIT Press.

Benaïm, M. (1998): "Recursive Algorithms, Urn Processes and Chaining Number of Chain Recurrent Sets," *Ergodic Theory and Dynamical Systems*, 18, 53–87.

——— (1999): "Dynamics of Stochastic Approximation," *Le Seminaire de Probabilité*, Springer Lecture Notes in Mathematics, Vol. XXXIII, 1709. Berlin: Springer Verlag, 1–68.

Benaïm, M., and M. Hirsch (1999a): "Stochastic Approximation Algorithms with Constant Step Size whose Average is Cooperative," *Annals of Applied Probability*, 9, 216–241.

——— (1999b): "Mixed Equilibria and Dynamical Systems Arising from Fictitious Play in Perturbed Games," *Games and Economic Behavior*, 29, 36–72.

Benaïm, M., and J. Weibull (2000): "Deterministic Approximation of Stochastic Evolution in Games," IUI WP No. 534, Revised 2001; available at www.iui.se.

Binmore, K., and L. Samuelson (1997): "Muddling Through: Noisy Equilibrium Selection," *Journal of Economic Theory*, 74, 235–265.

Binmore, K., L. Samuelson, and R. Vaughn (1995): "Musical Chairs: Modelling Noisy Evolution," *Games and Economic Behavior*, 11, 1–35.

Björnerstedt, J., and J. Weibull (1996): "Nash Equilibrium and Evolution by Imitation," in *The Rational Foundations of Economic Behaviour*, ed. by K. Arrow et al. London: Macmillan.

BOMZE, I. (1986): "Non-cooperative Two-person Games in Biology: A Classification," *International Journal of Game Theory*, 15, 31–57.

BÖRGERS, T., AND R. SARIN (1997): "Learning through Reinforcement and Replicator Dynamics," *Journal of Economic Theory*, 77, 1–14.

BOYLAN, R. (1995): "Continuous Approximation of Dynamical Systems with Randomly Matched Individuals," *Journal of Economic Theory*, 66, 615–625.

CONLEY, C. C. (1978): *Isolated Invariant Sets and the Morse Index*, Regional Conference Series in Mathematics 38. Providence RI: American Mathematical Society.

CORRADI, V., AND R. SARIN (1999): "Continuous Approximations of Stochastic Evolutionary Game Dynamics," forthcoming in *Journal of Economic Theory*.

DEMICHELIS, S., AND K. RITZBERGER (2001): "From Evolutionary to Strategic Stability," Unpublished Manuscript, Institute of Advanced Study, Vienna.

DUFLO, M. (1996): *Algorithmes Stochastique*, Mathématiques et Applications 23. Berlin: Springer Verlag.

——— (1997): *Random Iterative Models.* Berlin: Springer Verlag.

ELLISON, G. (1993): "Learning, Local Interaction, and Coordination," *Econometrica*, 61, 1047–1071.

——— (2000): "Basins of Attraction, Long-run Stochastic Stability, and the Speed of Step-by-step Evolution," *Review of Economic Studies*, 67, 17–45.

GALE, J., K. BINMORE, AND L. SAMUELSON (1995): "Learning to be Imperfect: The Ultimatum Game," *Games and Economic Behavior*, 8, 56–90.

FREIDLIN, M., AND A. WENTZELL (1984): *Random Perturbations of Dynamical Systems*. Berlin: Springer Verlag.

FUDENBERG, D., AND D. KREPS (1993): "Learning Mixed Equilibria," *Games and Economic Behavior*, 5, 320–367.

FUDENBERG, D., AND D. LEVINE (1998): *Theory of Learning in Games*. Cambridge, U.S.A.: MIT Press.

HOFBAUER, J., AND J. WEIBULL (1996): "Evolutionary Selection Against Dominated Strategies," *Journal of Economic Theory*, 71, 558–573.

KANDORI, M. (1996): "Evolutionary Game Theory in Economics," in *Advances in Economics and Econometrics: Theory and Applications*, Econometric Society World Congress 7, Vol. I, ed. by D. Kreps and K. Wallis. Cambridge, U.K.: Cambridge University Press.

KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation, and Long-run Equilibria in Games," *Econometrica*, 61, 29–56.

KANIOVSKI, Y., AND P. YOUNG (1995): "Learning Dynamics in Games with Stochastic Perturbations," *Games and Economic Behavior*, 11, 330–363.

KUSHNER, H., AND G. YIN (1997): *Stochastic Approximation Algorithms and Applications*. Berlin: Springer Verlag.

MAYNARD SMITH, J. (1982): *Evolution and the Theory of Games*. Cambridge U.K.: Cambridge University Press.

MAÑÉ, R. (1987): *Ergodic Theory and Differentiable Dynamics*. Berlin: Springer Verlag.

NACHBAR, J. (1990): "Evolutionary Selection Dynamics in Games: Convergence and Limit Properties," *International Journal of Game Theory*, 19, 59–89.

RITZBERGER, K., AND J. WEIBULL (1995): "Evolutionary Selection in Normal-form Games," *Econometrica*, 63, 1371–1400.

SAMUELSON, L., AND J. ZHANG (1992): "Evolutionary Stability in Asymmetric Games," *Journal of Economic Theory*, 57, 363–391.

SANDHOLM, W. (1999): "Markov Evolution with Inexact Information," Unpublished Manuscript, University of Wisconsin.

SCHLAG, K. (1998): "Why Imitate, and if so, How? A Bounded Rational Approach to Multi-armed Bandit Problems," *Journal of Economic Theory*, 78, 130–156.

SWINKELS, J. (1993): "Adjustment Dynamics and Rational Play in Game," *Games and Economic Behavior*, 5, 455–484.

TAYLOR, P. (1979): "Evolutionarily Stable Strategies with Two Types of Player," *Journal of Applied Probability*, 16, 76–83.

TAYLOR, P., AND L. JONKER (1978): "Evolutionary Stable Strategies and Game Dynamics," *Mathematical Biosciences*, 40, 145–156.
WEIBULL, J. (1995): *Evolutionary Game Theory*. Cambridge, U.S.A.: MIT Press.
YOUNG, P. (1998a): "Individual Learning and Social Rationality," *European Economic Review*, 42, Papers and Proceedings, 651–663.
——— (1998b): *Individual Strategy and Social Structure*. Princeton: University Press.