Stochastic Imitative Game Dynamics
with Committed Agents[*]

William H. Sandholm[†]

January 16, 2012

Abstract

We consider models of stochastic evolution in two-strategy games in which agents employ imitative decision rules. We introduce committed agents: for each strategy, we suppose that there is at least one agent who plays that strategy without fail. We show that unlike the standard imitative model, the model with committed agents generates unambiguous infinite horizon predictions: the asymptotics of the stationary distribution do not depend on the order in which the mutation rate and population size are taken to their limits.

1. Introduction

Models of stochastic evolution in games describe the behavior of agents who recurrently face the same strategic interaction, updating their choices over time by applying simple myopic rules. Over very long time spans, aggregate play in these models often concentrates on a single *stochastically stable state*, which then provides a unique prediction of long run behavior.

Following Binmore et al. [5], Binmore and Samuelson [4], and Fudenberg and Imhof [11, 12], the present paper focuses on imitative rules, under which a revising agent observes the behavior of a randomly selected opponent, and then switches to the opponent's strategy with a probability that may depend on his own current payoff, his opponent's current payoff, or both.[1]

---

[1]Imitative protocols can be interpreted as models of biological natural selection, with imitation of an individual playing a given strategy being viewed instead as reproduction by an individual playing that strategy; see Nowak et al. [19], Fudenberg et al. [13], and Traulsen and Hauert [26].

In a pure imitation model, a population that reaches a monomorphic state remains there forever. Thus, to employ stochastic stability analysis in imitative models, adjustments must be made to ensure that transitions between all pairs of states are possible. The papers mentioned above accomplish this by introducing a small probability of random experimentation or mutation, which ensures that the evolutionary process is irreducible, and so admits a unique stationary distribution. The stochastically stable states are those that retain mass in this stationary distribution as the parameters defining the evolutionary process approach their limiting values.

There are two main options for specifying these limits. Starting with Kandori et al. [15] and Young [28], many authors consider small noise limits, analyzing the behavior of the stationary distribution as the mutation rate approaches zero.[2] Others, including Binmore and Samuelson [4], Young [29, Section 4.5], and Benaïm and Weibull [3], focus on large population limits, studying the behavior of the stationary distribution as the population size approaches infinity. To obtain clean comparisons between the two approaches, Binmore et al. [5] and Binmore and Samuelson [4] consider definitions of stochastic stability in which the small noise and large population limits are taken sequentially. Since the parameter in the outer limit is fixed while the inner limit is taken, it is the parameter in this inner limit that governs equilibrium selection.[3]

If the order of limits employed did not matter for equilibrium selection, then careful consideration of this modeling choice would not be necessary. However, Binmore and Samuelson [4] show that the order of limits can matter. In Hawk-Dove games, whose unique Nash equilibrium is interior, this Nash equilibrium is stochastically stable in the large population double limit. This prediction agrees with the behavior of the stochastic evolutionary process's *mean dynamic*, the differential equation defined by the process's expected motion.[4] But in the small noise double limit, the stochastically stable state must be monomorphic, with all agents playing the same strategy. For while transitions from the Nash equilibrium to a monomorphic state are quite unlikely when the population size is large, a small enough mutation rate ensures that these rare transitions are far more likely than the single mutation needed to escape a monomorphic state. Consequently,

---

[2]Fudenberg and Imhof [11, 12] and Agarwal and Lai [1] show that when it is employed in conjunction with imitation, the small noise limit is particularly tractable.

[3]A modeler's choice between these limits can reflect his assessment of whether the unlikelihood of mistakes or the largeness of the population is the driving force behind equilibrium selection. Binmore and Samuelson [4] argue that in economic contexts, experimentation and errant choices are not vanishingly rare, and so suggest that the large population limit is typically more appropriate for economic modeling. On the other hand, since genetic mutations may well be truly exceptional events, the small noise limit may be a more appropriate choice for modeling biological natural selection.

[4]As we discuss below, general results of Benaïm [2] and Benaïm and Weibull [3] imply that the the stochastically stable state must be a recurrent point of the mean dynamic.

most periods are spent at monomorphic states in the very long run.

In economic settings for which evolutionary models are appropriate—for instance, those that are one of many that an individual agent faces—agents only reconsider their choices of strategies sporadically. If there are agents for whom the setting in question is of particularly minor importance, or who are engaged in other activities that place great demands on their attention or reasoning capacities, it seems reasonable to expect that some of these agents may not bother to consider switching strategies at all.

To capture this possibility, we alter the standard model by introducing *committed agents*, who play particular strategies without fail. Doing so ensures that every strategy is always available for imitation. Because of this, the stochastic process describing the evolution of aggregate behavior is irreducible even in the absence of mutations.

Our main results show that in two-strategy games, the inclusion of committed agents in imitative models leads to unambiguous predictions about infinite-horizon behavior. Whether one takes the small noise limit first, or the large population limit first, or whether one leaves out mutations and takes only the large population limit, the prediction of infinite-horizon behavior is the same: all specifications yield the same infinite-horizon behavior as the large population limit in the model without committed agents. Indeed, this agreement is established in a very strong sense: we obtain explicit expressions for the rates of decay of the stationary distribution weights for each arrangement of the limits, and prove that these asymptotics are all identical.[5] These results reinforce Binmore and Samuelson's [4] assessment that in economic settings, a focus on large population limits may be warranted.

2. The Model

We consider games in which the members of a population of N agents choose strategies from the common finite strategy set S. In setting up the model there is no advantage to restricting the number of strategies $n = \#S$, but we will restrict attention to the two-strategy case once the analysis commences in Section 3.

We describe the population's aggregate behavior by a *population state x*, an element of the simplex $X = \{x \in \mathbf{R}_+^n : \sum_{i=1}^n x_i = 1\}$, or more specifically, the grid $\mathcal{X}^N = X \cap \frac{1}{N}\mathbf{Z}^n = \{x \in X : Nx \in \mathbf{Z}^n\}$. We identify a finite-population game with its payoff function $F^N : \mathcal{X}^N \to \mathbf{R}^n$, where $F_i^N(x) \in \mathbf{R}$ is the payoff to strategy i when the population state is $x \in \mathcal{X}^N$. Only

---

[5]This analysis builds on methods developed in Binmore and Samuelson [4], Blume [8], and Sandholm [22, 23]. The last of these papers considers stochastic evolution under a general class of noisy best response protocols, and shows that both orders of limits lead to identical predictions of infinite horizon play.

the values that the function F_i^N takes on the set $\mathcal{X}_i^N = \{x \in \mathcal{X}^N : x_i > 0\}$ are meaningful, since at the remaining states in $\mathcal{X}^N$ strategy i is unplayed.

State $x \in \mathcal{X}^N$ is a *Nash equilibrium* of F^N if no agent can obtain a higher payoff by switching strategies:

$$[x_i > 0 \Rightarrow F_i^N(x) \geq F_j^N(x + \tfrac{1}{N}(e_j - e_i))] \quad \text{for all } i, j \in S,$$

where e_i is the ith standard basis vector in $\mathbf{R}^n$. This definition accounts for the fact that after an agent switches from strategy i to strategy j, the population has one less i player and one more j player.

To consider large population limits, we must specify a notion of convergence for sequences $\{F^N\}_{N=N_0}^{\infty}$ of finite-population games. If a such a sequence converges, its limit is a continuous-population game, $F \colon X \to \mathbf{R}^n$, which we require to be a continuous function from X to $\mathbf{R}$. A natural notion of convergence for the sequence $\{F^N\}_{N=N_0}^{\infty}$ is uniform convergence, which asks that

$$\lim_{N \to \infty} \max_{x \in \mathcal{X}^N} |F^N(x) - F(x)| = 0.$$

Uniform convergence obtains when agents are matched to play a symmetric two-player normal form game $A \in \mathbf{R}^{n \times n}$, whether this occurs with self-matching ($F_i^N(x) = \sum_{j \in S} A_{ij} x_j = (Ax)_i$) or without ($F_i^N(x) = \frac{1}{N-1}(A(Nx - e_i))_i = (Ax)_i + \frac{1}{N-1}((Ax)_i - A_{ii})$).

We suppose that the game F^N is played recurrently, with agents updating their strategies over time by employing a *revision protocol* $\sigma^N \colon \mathbf{R}^n \times X \to \mathbf{R}_+^{n \times n}$. A revision protocol takes a payoff vector $\pi \in \mathbf{R}^n$ and a population state $x \in X$ as inputs, and returns as output a nonnegative matrix $\sigma^N(\pi, x)$ whose row sums equal 1:

$$\sum_{j \in S} \sigma_{ij}^N(\pi, x) = 1 \quad \text{for all } i \in S. \tag{1}$$

The scalar $\sigma_{ij}^N(\pi, x)$ is called the *conditional switch probability* from strategy i to strategy j.

A population game F^N and a revision protocol σ^N define a stochastic evolutionary process $\{X_t^N\}_{t \geq 0}$ on the state space $\mathcal{X}^N$ in the following way. Each of the N agents in the population receives revision opportunities according to independent, rate 1 Poisson alarm clocks. If a current i player receives a revision opportunity when the population state is x, he switches to strategy j with probability $\sigma_{ij}^N(F^N(x), x)$.

This paper studies the evolution of play under *imitative protocols*. In the standard setting, without committed agents, these protocols take the form

$$\sigma_{ij}^N(\pi, x) = \frac{Nx_j}{N-1} r_{ij}(\pi) \;\; \text{for } j \neq i, \tag{2}$$

with the value of $\sigma_{ii}^N(\pi, x)$ being determined by the row sum condition (1). The $\frac{Nx_j}{N-1}$ term reflects the idea that an agent who receives a revision opportunity randomly samples one of his opponents and takes her strategy as his candidate strategy. He then opts to switch to the candidate strategy with probability $r_{ij}(\pi)$. We assume throughout that

Each function $r_{ij} \colon \mathbf{R}^n \to [0, 1]$ is continuous and bounded away from zero. (B)

This assumption ensures that there is always some chance that the agent imitates the opponent he observes. Standard examples of revision protocols satisfying this condition can be found in the online appendix.

A basic feature of evolutionary processes generated by protocols of form (2) is *extinction*: unused strategies are never subsequently chosen. To respond to this difficulty, Binmore and Samuelson [4] introduce the possibility of *mutation*, replacing protocol (2) with

$$\sigma_{ij}^{N,\varepsilon}(\pi, x) = (1 - \varepsilon)\frac{Nx_j}{N-1} r_{ij}(\pi) + \frac{1}{n-1}\varepsilon \;\; \text{for } j \neq i. \tag{3}$$

Under protocol (3), a revising i player behaves as in protocol (2) with probability $1 - \varepsilon$, but with the remaining probability of ε he chooses a new strategy at random from the $n - 1$ alternatives. The resulting evolutionary process is irreducible for any positive choice of ε.

Mutations are not the only way of making imitation compatible with stochastic stability analysis. As an alternative, one can assume that in addition to the N standard agents, there are also $c_T^N = \sum_{i \in S} c_i^N$ *committed agents*. The $c_i^N \geq 1$ committed agents associated with strategy i play this strategy without fail. The presence of committed agents ensures that no strategy ever goes extinct, and so that imitative protocols satisfying condition (B) generate irreducible evolutionary processes even without mutations.

We account for the existence of committed agents by replacing revision protocol (2) with the protocol

$$\sigma_{ij}^N(\pi, x) = \frac{Nx_j + c_j^N}{N + c_T^N - 1} r_{ij}(\pi) \;\; \text{for } j \neq i, \tag{4}$$

where $x \in X^N$ represents the behavior of the standard (uncommitted) agents. To make direct comparisons with protocol (3), we also consider imitative protocols that account for committed agents and mutations:

$$\sigma_{ij}^{N,\varepsilon}(\pi, x) = (1 - \varepsilon)\frac{Nx_j + c_j^N}{N + c_T^N - 1} r_{ij}(\pi) + \frac{1}{n-1}\varepsilon \;\; \text{for } j \neq i. \tag{5}$$

Until Section 4.3, we assume that $c_i^N = c_i$ is independent of the population size, so that the fraction of committed agents becomes vanishingly small as N grows large. Section 4.3 considers the alternative assumption that c_i^N is proportional to N.

3. Preliminary Analysis

The analysis to follow focuses on play in two-strategy games. In this context, it is convenient to let $S = \{0, 1\}$ be the set of strategies, and to identify the population state x with the weight $\chi \equiv x_1$ that it places on strategy 1. With this shorthand, the set of population states becomes $\mathcal{X}^N = \{0, \frac{1}{N}, \frac{2}{N}, \ldots, 1\}$, a uniformly spaced grid in the unit interval. We write $F^N(\chi)$ for $F^N(x)$ and make similar substitutions whenever it is convenient to do so.

Stochastic evolution in a two-strategy game is described by a birth-and-death process $\{X_t^N\}_{t \geq 0}$ on the state space $\mathcal{X}^N$, with transitions only occurring between adjacent states. For the process $\{X_t^N\}$ to step upward during the next revision opportunity, this opportunity must be received by a strategy 0 player who then switches to strategy 1. Thus, the probability of an upward transition is

$$p_\chi^N = (1 - \chi)\, \sigma_{01}^N(F^N(\chi), \chi). \tag{6}$$

For a step downward to occur, the next revision opportunity must be received by a strategy 1 player who switches to strategy 0, so the probability of a downward transition is

$$q_\chi^N = \chi\, \sigma_{10}^N(F^N(\chi), \chi). \tag{7}$$

If there is a positive probability of mutations, or if committed agents are present, the process $\{X_t^N\}$ is irreducible. It thus admits a unique stationary distribution μ^N, which describes the infinite-horizon behavior of the evolutionary process in two distinct ways: it is the limiting distribution of the process, and it describes the limiting empirical distribution of the process along almost every sample path (see, e.g., Norris [18]). In the present case, this stationary distribution is described by the following well-known formula:

$$\frac{\mu_\chi^N}{\mu_0^N} = \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^N}{q_{j/N}^N} \quad \text{for } \chi \in \{\tfrac{1}{N}, \tfrac{2}{N}, \ldots, 1\}. \tag{8}$$

The value of μ_0^N is determined by the requirement that the sum of the stationary distribution weights equal 1.

Some basic conclusions about the limiting behavior of the stationary distributions μ_N

as the population size N grows large can be deduced from the *mean dynamic* associated with the collection $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$, which describes the expected motion of the process $\{X_t^N\}$ N is large. Suppose that the population games F^N converge uniformly to a limit game F, and that the revision protocols σ^N converge uniformly to a limiting protocol σ. Then in the two-strategy case, the mean dynamic is given by

$$\dot{x} = (1 - x)\, \sigma_{01}(F(x), x) - x\, \sigma_{10}(F(x, x)). \tag{9}$$

The mean dynamic presents the expected rate of change in the use of strategy 1 as the difference between two terms: the expected inflow into strategy 1 from strategy 0, and the expected outflow from strategy 1 to strategy 0.

   The mean dynamic (9) describes the behavior of the evolutionary process in the large population limit in two distinct ways. First, one can establish a finite-horizon deterministic approximation theorem: if we fix a finite time horizon T, then for N large enough, the sample paths of the process $\{X_t^N\}$ are very likely to closely track the solution to the mean dynamic (9) from the relevant initial condition. Second, one can obtain restrictions on the limiting behavior of the stationary distribution: when N is large enough, the mass in μ^N must become concentrated near the recurrent points—in the present, one-dimensional case, the rest points—of the dynamic (9).[6]

   If there are fixed numbers of committed agents, then in the large population limit, the no-mutation protocols (2) and (4) both converge uniformly to the limit protocol

$$\sigma_{ij}(\pi, x) = x_j r_{ij}(\pi) \ \text{ for } j \neq i,$$

and so generate the same mean dynamic,

$$\dot{x} = x(1 - x)\,(r_{01}(F(x)) - r_{10}(F(x)))\,. \tag{M_0}$$

Likewise, the common mean dynamic for the protocols with mutations, (3) and (5), is

$$\dot{x} = (1 - \varepsilon)\,x(1 - x)\,(r_{01}(F(x)) - r_{10}(F(x))) + \varepsilon(1 - 2x). \tag{M_ε}$$

   For ease of interpretation, we henceforth suppose that *conditional imitation probabilities are monotone*:

$$\pi_j \geq \pi_i \Leftrightarrow r_{ij}(\pi) \geq r_{ji}(\pi) \text{ for } i, j \in \{0, 1\}. \tag{10}$$

---

[6]See Kurtz [16], Benaïm [2], Benaïm and Weibull [3], and Roth and Sandholm [20].

In words: if an agent observes an opponent playing the other strategy, he is more likely to imitate if his current strategy has the lower payoff than if it has the higher payoff. Under condition (10), the rest points of the dynamic (M_0) are the Nash equilibria of F and the monomorphic states $\chi = 0$ and $\chi = 1$; if the latter are not Nash equilibria, they are unstable. The dynamic (M_ε) introduces a small force pushing toward the center of the simplex, eliminating rest points of (M_0) corresponding to non-Nash monomorphic states.

4. Results

One basic question in stochastic evolutionary game theory concerns the effects of different orders of limits on predictions of infinite-horizon play. Following Sandholm [23], we call states that retain mass in the limiting stationary distribution either *stochastically stable in the small noise double limit* or *stochastically stable in the large population double limit* according to whether the inner limit is in ε or N.

The results to follow describe the asymptotics of the stationary distribution, and *a fortiori* the stochastically stable states, for both orders of limits, and for settings both without and with committed agents. These results require the function $J: [0,1] \to \mathbf{R}$, defined in terms of the limit game $F: [0,1] \to \mathbf{R}^2$ and the conditional imitation probabilities $r_{01}(\cdot)$ and $r_{10}(\cdot)$ by

$$J(\chi) = \int_0^\chi \log \frac{r_{01}(F(y))}{r_{10}(F(y))} \, dy. \tag{11}$$

It is easy to verify that J is is a *strict Lyapunov function* for the mean dynamic (M_0): its value increases along solutions of (M_y), strictly so whenever the dynamic is not at rest. Moreoover, in light of condition (10), J is an *ordinal potential function* for the game F, in that marginally increasing the weight on the optimal strategy always increases the value of J. From this it follows immediately that local maximizers of J are Nash equilibria of F.[7]

Theorems 4.1–4.3 are stated in terms of the function $\Delta J: X \to \mathbf{R}_-$, defined by

$$\Delta J(\chi) = J(\chi) - \max_{y \in X} J(y).$$

This function is obtained from J by shifting the graph of J vertically to place its maximum value at 0.

---

[7]Compare Monderer and Shapley [17] and Sandholm [21, 23].

4.1 No committed agents

To obtain a concise statement of the results for settings without committed agents, we say that the collection $\{\alpha^\varepsilon\}_{\varepsilon \in (0, \bar{\varepsilon}]}$ is *of exact order ε*, denoted $\alpha^\varepsilon \in \Theta(\varepsilon)$, if there is an interval $[a, b] \subset (0, \infty)$ such that $\alpha^\varepsilon/\varepsilon \in [a, b]$ for all ε close enough to zero.

Theorem 4.1. *Let $\{F^N\}$ be a sequence of two-strategy population games that converges uniformly to the continuous-population game F. Suppose that agents employ an imitative protocol with mutations. Then the stationary distributions $\mu^{N,\varepsilon}$ satisfy*

(i) $\displaystyle \lim_{N \to \infty} \lim_{\varepsilon \to 0} \frac{1}{N} \log \frac{\mu_1^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = J(1)$, *and* $\dfrac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}}$ *and* $\dfrac{\mu_\chi^{N,\varepsilon}}{\mu_1^{N,\varepsilon}}$ *are in $\Theta(\varepsilon)$ when $\chi \in \mathcal{X}^N - \{0, 1\}$, and*

(ii) $\displaystyle \lim_{\varepsilon \to 0} \lim_{N \to \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_\chi^{N,\varepsilon} - \Delta J(\chi) \right| = 0.$

The proofs of this and subsequent results build on methods developed by Binmore and Samuelson [4], Blume [8], and Sandholm [22, 23] for analyzing limiting stationary distributions of birth-and-death chains. The proofs of Theorems 4.2 and 4.3 are presented in Section 6, while the proofs of Theorems 4.1 and 4.4 appear in an online appendix.

Theorem 4.1(i), which is essentially due to Binmore and Samuelson [4], shows that in the small noise double limit, all of the mass in the stationary distribution becomes concentrated on boundary states regardless of the game at hand. The sign of $J(1)$ (= $J(1) - J(0)$) determines which state is selected: state 1 if $J(1) > 0$ and state 0 if $J(1) < 0$. The magnitude of $J(1)$ determines the rate at which the mass on the other boundary state vanishes as the population size grows large. Theorem 4.1(ii) shows that in the large population double limit, the only states that can retain mass in the limiting stationary distribution are those that maximize the ordinal potential function J, which thus are Nash equilibria of the limit game. In addition, Theorem 4.1(ii) reveals that the rate at which the mass near state χ decays as the population size N grows large is determined by the value of ΔJ at that state.[8]

Theorem 4.1 implies that in games without monomorphic Nash equilibria, the small noise and large population double limits must specify different stochastically stable states. The reason for this discrepancy can be explained as follows. The only way that the process $\{X_t^{N,\varepsilon}\}$ can escape from a boundary state is by way of a mutation. If we fix the population size N and make ε extremely small, then a journey from an interior state to a boundary state—here a journey against the flow of the mean dynamic—is more likely than an escape from a boundary state by way of a single mutation. It follows that in the small noise double limit, the stationary distribution must become concentrated on the boundary states—typically, on just one of the boundary states—regardless of the expected motion

---

[8]This is so because $\frac{1}{N} \log \mu_\chi^{N,\varepsilon} = -r(\chi)$ if and only if $\mu_\chi^{N,\varepsilon} = \exp(-Nr(\chi))$.

of the system. On the other hand, if we fix a small value of ε and make N large, then as explained in Section 3, the stationary distribution $\mu^{N,\varepsilon}$ must become concentrated near the rest points of the perturbed mean dynamic (M_ε). These are only near Nash equilibria, not non-Nash boundary states. Thus, in the large population double limit, the mass in $\mu^{N,\varepsilon}$ becomes concentrated on Nash equilibria.[9]

4.2 Fixed numbers of committed agents

When committed agents are present, the stochastic evolutionary process is irreducible even without mutations. Theorem 4.2 shows that the large population limit with fixed numbers of committed agents agrees with the large population double limit from Theorem 4.1: in both cases, rates of decay of stationary distribution weights, and hence the identity of the stochastically stable state, are determined by the ordinal potential J.

Theorem 4.2. *Let $\{F^N\}$ be a sequence of two-strategy population games that converges uniformly to the continuous-population game F. Suppose that the N standard agents employ an imitative protocol without mutations, and that there are $c_i > 0$ committed agents playing strategy $i \in \{0, 1\}$. Then the sequence of stationary distributions $\{\mu^N\}$ satisfies*

$$\lim_{N\to\infty} \max_{x\in\mathcal{X}^N} \left| \tfrac{1}{N} \log \mu_x^N - \Delta J(x) \right| = 0.$$

Since there are no mutations and a vanishing fraction of committed agents, the stochastic approximation results described in Section 3 imply that as N grows large, the mass in the stationary distribution μ^N becomes concentrated on the rest points of the unperturbed mean dynamic (M_0). These rest points include any non-Nash boundary states. Nevertheless, a careful analysis of the explicit formula (8) for the stationary distribution reveals that such states cannot be selected, and moreover, that the asymptotics of the stationary distribution take the same form as in the large population double limit from Theorem 4.1(ii).

Introducing mutations brings back the question of order of limits, which is addressed by our next result.

Theorem 4.3. *Let $\{F^N\}$ be a sequence of two-strategy population games that converges uniformly to F. Suppose that the N standard agents employ an imitative protocol with mutations, and that there are $c_i > 0$ committed agents playing strategy $i \in \{0, 1\}$. Then the stationary distributions $\mu^{N,\varepsilon}$ satisfy*

---

[9]In related work, Fudenberg and Hojman [10] consider taking the population size and the noise level to their limits simultaneously, and show how the identity of the stochastically stable state depends on the relative rates at which these limits are taken.

(i) $\quad \lim\limits_{N \to \infty} \lim\limits_{\varepsilon \to 0} \max\limits_{\chi \in X^N} \left| \frac{1}{N} \log \mu_\chi^{N,\varepsilon} - \Delta J(\chi) \right| = 0 \ and$

(ii) $\quad \lim\limits_{\varepsilon \to 0} \lim\limits_{N \to \infty} \max\limits_{\chi \in X^N} \left| \frac{1}{N} \log \mu_\chi^{N,\varepsilon} - \Delta J(\chi) \right| = 0.$

Theorem 4.3 shows that when committed agents are present, the asymptotics of the stationary distribution in the small noise and large population double limits are identical. In both cases, the stochastically stable states are maximizers of the function J, and hence are Nash equilibria of the limit game F.

To establish part (i) of the theorem, we observe that for each fixed population size, as the mutation rate ε approaches zero, the transition probabilities of the process with mutations converge uniformly to those of the process without mutations. This in turn implies that the small noise double limit agrees with the limit from Theorem 4.2. For part (ii), we note that when the mutation rate ε is positive, ensuring that the evolutionary process is irreducible, the presence of a vanishing fraction of committed agents has a negligible effect on the stationary distribution. For this reason, the large population double limit here agrees with the one from Theorem 4.1(ii).

4.3 Fixed proportions of committed agents

The analysis in Section 4.2 considered cases in which the numbers of committed agents were fixed independently of the population size, so that the proportions of committed agents vanish as the population size grows large. But in some settings, it is natural to suppose that committed agents make up a nonnegligible proportion of the population. To address this possibility, we now suppose that when the population size is N, the number of agents committed to strategy j is $c_j^N = N\gamma_j > 0$, so that the proportion of committed agents in the population stays fixed as N grows.[10] If mutations are not introduced, this specification leads to the mean dynamic

$$\dot{\chi} = \frac{1}{\gamma_0 + \gamma_1} \left(\chi(1 - \chi)\left(r_{01}(F(\chi)) - r_{10}(F(\chi))\right) + (1 - \chi)\gamma_1 r_{01}(F(\chi)) - \chi\gamma_0 r_{10}(F(\chi)) \right). \qquad (\mathrm{M}_\gamma)$$

Comparing this dynamic with (M_0), we see that adding fixed proportions of committed agents introduces forces pushing away from each endpoint of the unit interval.

We now describe how the presence of fixed fractions of committed agents affects the asymptotics of the stationary distribution. To do so, we define the continuous function $\ell \colon \mathbf{R}_+ \to \mathbf{R}$ by $\ell(x) = x \log x$ for $x > 0$ and $\ell(0) = 0$, and we define $L^\gamma \colon [0, 1] \to \mathbf{R}$ by

$$L^\gamma(\chi) = \ell(\chi + \gamma_1) - \ell(\chi) - \ell(\gamma_1) + \ell(1 - \chi + \gamma_0) - \ell(1 - \chi) - \ell(1 + \gamma_0),$$

---

[10]To do this, we must only consider values of N for which each $N\gamma_j$ is an integer.

L^γ is a concave function that is maximized at $\frac{\gamma_0}{\gamma_0+\gamma_1}$, and it is symmetric about $\frac{1}{2}$ when $\gamma_0 = \gamma_1$. Finally, we define $J^\gamma \colon [0,1] \to \mathbf{R}$ and $\Delta J^\gamma \colon [0,1] \to \mathbf{R}$ by

$$J^\gamma(x) = J(x) + L^\gamma(x) \quad \text{and} \quad \Delta J^\gamma(x) = J^\gamma(x) - \max_{y \in X} J^\gamma(y)$$

It is easy to verify that under monotonicity condition (10), J^γ (and also ΔJ^γ) is a strict Lyapunov function for the mean dynamic (M_γ).

Theorem 4.4 describes the behavior of the stationary distribution in the large population limit, providing the analogue of Theorem 4.2 for settings with fixed proportions of committed agents.

Theorem 4.4. *Let $\{F^N\}$ be a sequence of two-strategy population games that converges uniformly to the continuous-population game F. Suppose that the N standard agents employ an imitative protocol without mutations, and that there are $N\gamma_i > 0$ committed agents playing strategy $i \in \{0,1\}$. Then the sequence of stationary distributions $\{\mu^N\}$ satisfies*

$$\lim_{N\to\infty} \max_{x \in X^N} \left| \tfrac{1}{N} \log \mu_x^N - \Delta J^\gamma(x) \right| = 0.$$

Theorem 4.4 shows that when fixed proportions of committed agents are present, the rates of decay of stationary distribution weights are described not by J, but rather by $J^\gamma = J + L^\gamma$. Since L^γ is concave and maximized at $\frac{\gamma_0}{\gamma_0+\gamma_1}$, both the mass in the stationary distribution and the stochastically stable state move toward the state that corresponds to the aggregate behavior of the committed agents.

If γ_0 and γ_1 are large, then $L^\gamma(x) \approx -(\ell(x)+\ell(1-x))+x\log\frac{\gamma_1}{\gamma_0}$. In this case, if the fractions of agents committed to each strategy are nearly equal, then $L^\gamma(x)$ is approximately the entropy function $h(x) = -(\ell(x) + \ell(1 - x))$, and so $J^\gamma \approx J + h$.[11] On the other hand, if $\frac{\gamma_1}{\gamma_0}$ is quite far from 1, then L^γ is approximately a linear function whose slope is large in absolute value, and so $J^\gamma(x) \approx x\log\frac{\gamma_1}{\gamma_0}$ is approximately such a function as well. In this case the stochastically stable state has nearly all active players choosing the strategy played by the vast majority of committed agents.

With Theorem 4.4 in hand, a modification of the proof of Theorem 4.3 can be used to establish the obvious extension of that order-of-limits result to settings with fixed proportions of committed agents and rare mutations.

---

[11]The entropy function plays a similar role in the analysis of large population stationary distribution asymptotics for noisy best response dynamics; see Sandholm [22].

5. Discussion

5.1 Imitative protocols, mean dynamics, and equilibrium selection

The results in Section 4.2 show that with small numbers of committed agents, the asymptotics of the stationary distribution of the stochastic evolutionary process are unambiguously described by the function J, which also serves as a Lyapunov function for the processes' mean dynamic. It does not follow, however, that stochastic evolutionary processes with the same mean dynamic generate the same equilibrium selections. In the online appendix, we consider stochastic evolutionary processes generated by two commonly studied imitative protocols: *imitation driven by dissatisfaction*, which conditions only on the payoff of the agent's current strategy, and *imitation of success*, which conditions only on the payoff of the candidate strategy.[12] Both protocols we consider generate the replicator dynamic (Taylor and Jonker [25]) as their mean dynamics, and so lead to behavior that is indistinguishable over finite time spans. Nevertheless, we show that in certain simple coordination games, the two protocols lead to different infinite-horizon equilibrium predictions: imitation of success leads to the selection of the safe equilibrium, while imitation driven by satisfaction leads to selection of the payoff dominant equilibrium.

5.2 Games with three or more strategies

Our analysis of stochastic imitative dynamics with committed agents has focused on evolution in two-strategy games. This restriction ensures that the stochastic evolutionary process is a birth-and-death process, and so that its stationary distribution is described by a simple explicit formula. Once we move beyond the two-strategy case, such formulas are generally unavailable.[13]

We conjecture that Theorem 4.3, which establishes the invariance of infinite horizon predictions to the order of limits in N and ε, should continue to hold in games with three or more strategies. To explain the issues involved, let us focus this discussion on games whose unique Nash equilibrium x^* is interior and almost globally asymptotically stable under the relevant mean dynamic (M_0) (or to be more precise, under its n-strategy analogue).[14] In this setting, the conjecture implies that if small numbers of committed

---

[12]Early analyses of these protocols include Björnerstedt and Weibull [6], Weibull [27], and Hofbauer [14].

[13]We know of only one setting in which this is possible. Building on the work of Blume [7], one can obtain an explicit formula for the stationary distribution when the game being played is an n strategy potential game, when there is exactly one committed agent for each strategy, and when the standard agents employ an imitative analogue of the logit rule. See Sandholm [24, Sec. 11.5.3] for details.

[14]This is true, for instance, if the mean dynamic is the replicator dynamic, and the limit game F is strictly stable; see Sandholm [24, Theorem 7.2.4].

agents are present, then under either orders of limits in N and ε, the limiting stationary distribution should become concentrated near x^*.

This last claim is easy to prove for the large population double limit whether the fraction of committed agents is fixed or vanishing. Here we only consider the latter case. For any $\varepsilon > 0$, results from stochastic approximation theory[15] can be used to show that the limiting distribution $\lim_{N \to \infty} \mu^{N,\varepsilon}$ is concentrated on the recurrent set of the mean dynamic with mutations, (M_ε). This dynamic is a perturbed version of the no-mutation dynamic (M_0). But while the boundary of the state space is invariant under (M_0), it is repelling under (M_ε). Hence, results on continuation of attractors[16] can be used to show that as ε approaches 0, the global attractor of (M_ε) approaches the Nash equilibrium x^*. This allows us to conclude that the limiting distribution $\lim_{\varepsilon \to 0} \lim_{N \to \infty} \mu^{N,\varepsilon}$ is a point mass at x^*.

For the small noise double limit, one can establish a version of the claim when the fractions of committed agents are small but fixed. The argument runs along the following lines: Since ε is taken to 0 first, the relevant mean dynamic for the subsequent evaluation of the large N limit is the no-mutation dynamic (M_γ). The fixed fractions of committed agents ensure that like (M_ε), (M_γ) is a perturbed version of (M_0) with a repelling boundary. Thus, if the fractions of committed agents are small, the global attractor of (M_γ) lies near x^*, and the limiting distribution $\lim_{N \to \infty} \lim_{\varepsilon \to 0} \mu^{N,\varepsilon}$ concentrates its mass on this attractor.

But with vanishing fractions of committed agents, the analysis of the small noise double limit is considerably more difficult. Again, since ε is taken to 0 first, the relevant mean dynamic for the subsequent evaluation of the large N limit is the no-mutation dynamic, in this case the dynamic (M_0). But now, general results from stochastic approximation theory tell us only that the limiting stationary distribution becomes concentrated on the set of recurrent points of (M_0). This set includes all monomorphic states, which are (unstable) rest points of (M_0). While one expects that such points cannot retain mass in the limiting stationary distribution, this cannot be proved using existing methods.[17] Thus, the analysis of the small noise double limit with vanishing fractions of committed agents remains a challenging open question.

---

[15]See Benaïm [2], Benaïm and Weibull [3], and Roth and Sandholm [20].

[16]See Sandholm [24, Theorem 9.B.5].

[17]These methods are introduced by by Benaïm [2]. A basic step in these methods is to obtain bounds on the probability of a transition from an unstable rest point to the stable one; ultimately, such bounds are used as inputs to Freidlin and Wentzell's [9] graph-theoretic approach to determining the limiting stationary distribution. In the present application, an unstable rest point of (M_0) at a monomorphic state is a *degenerate* state of the stochastic processes: it is not only a state where the expected motion of the process approaches zero, but one where the probabilities of all non-null transitions approach zero. In such degenerate cases, the methods from the theory of sample path large deviations used to obtain bounds on transition probabilities (see Benaïm [2, Sec. 6]) do not apply.

6. Proofs

The Proof of Theorem 4.2

By substituting the committed agents protocol (4) into equations (6) and (7), we obtain the one-step transition probabilities

$$p_\chi^N = (1 - \chi) \cdot \frac{N\chi + c_1}{N + c_T - 1} r_{01}(F^N(\chi)) \quad \text{and}$$

$$q_\chi^N = \chi \cdot \frac{N(1 - \chi) + c_0}{N + c_T - 1} r_{10}(F^N(\chi)).$$

Inserting these expressions into equation (8) and canceling like terms, we find that the stationary distribution of the process $\{X_t^N\}$ is given by

$$
\begin{aligned}
\frac{\mu_\chi^N}{\mu_0^N} &= \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^N}{q_{j/N}^N} \\
&= \prod_{j=1}^{N\chi} \frac{\frac{N-j+1}{N} \cdot \frac{j+c_1-1}{N+c_T-1} r_{01}(F^N(\frac{j-1}{N}))}{\frac{j}{N} \cdot \frac{N-j+c_0}{N+c_T-1} r_{10}(F^N(\frac{j}{N}))} \\
&= \prod_{j=1}^{N\chi} \frac{N-j+1}{N-j+c_0} \cdot \frac{j+c_1-1}{j} \cdot \frac{r_{01}(F^N(\frac{j-1}{N}))}{r_{10}(F^N(\frac{j}{N}))}.
\end{aligned}
\tag{12}
$$

Now recall that the finite-population games $F^N \colon X^N \to \mathbf{R}^2$ converge uniformly to the continuous limit game $F \colon [0,1] \to \mathbf{R}^2$, and that r_{01} and r_{10} are continuous and bounded away from zero by condition (B). Therefore, if we define the functions $v^N \colon [0,1] \to \mathbf{R}$ by

$$
v^N(\chi) =
\begin{cases}
\log \dfrac{r_{01}\left(F^N\left(\frac{\lceil N\chi \rceil - 1}{N}\right)\right)}{r_{10}\left(F^N\left(\frac{\lceil N\chi \rceil}{N}\right)\right)} & \text{if } \chi \in (0,1], \\[3ex]
\log \dfrac{r_{01}(F(0))}{r_{10}(F(0))} & \text{if } \chi = 0,
\end{cases}
$$

then the v^N are uniformly bounded and converge almost surely to

$$v(x) \equiv \log \frac{r_{01}(F(x))}{r_{10}(F(x))}.$$

If we define $J^N \colon X^N \to \mathbf{R}$ by

$$J^N(\chi) = \int_0^\chi v^N(y)\, dy = \frac{1}{N} \sum_{j=1}^{N\chi} \log \frac{r_{01}(F^N(\frac{j-1}{N}))}{r_{10}(F^N(\frac{j}{N}))},$$

–15–

then equation (12) implies that

$$\frac{1}{N} \log \frac{\mu_\chi^N}{\mu_0^N} = J^N(\chi) + \frac{1}{N} \sum_{j=1}^{N\chi} \log \frac{N-j+1}{N-j+c_0} + \frac{1}{N} \sum_{j=1}^{N\chi} \log \frac{j+c_1-1}{j}. \tag{13}$$

If we could show that the second and third summands in (13) converge uniformly to 0 as N approaches infinity, it would follow that

$$\lim_{N \to \infty} \max_{\chi \in X^N} \left| \frac{1}{N} \log \frac{\mu_\chi^N}{\mu_0^N} - J^N(\chi) \right| = 0. \tag{14}$$

(We note as an aside that if there is just one committed agent for each strategy, so that $c_0 = c_1 = 1$, these summands are both 0 by definition, and (14) follows trivially.) So fix N, and consider the third summand in (13). If $\chi > \frac{c_1}{N}$, then canceling like terms yields

$$\frac{1}{N} \sum_{j=1}^{N\chi} \log \frac{j+c_1-1}{j} = \frac{1}{N} \sum_{j=1}^{c_1-1} \log \frac{N\chi+j}{j} \leq \frac{c_1}{N} \log(N+1).$$

On the other hand, if $\chi \leq \frac{c_1}{N}$, then

$$\frac{1}{N} \sum_{j=1}^{N\chi} \log \frac{j+c_1-1}{j} \leq \frac{1}{N} \sum_{j=1}^{c_1} \log \frac{j+c_1-1}{j} \leq \frac{c_1}{N} \log c_1.$$

The right-hand sides of these expressions vanish as N approaches infinity, establishing the desired uniform convergence. For the second summand of (13), a similar argument shows that

$$\left| \frac{1}{N} \sum_{j=1}^{N\chi} \log \frac{N-j+1}{N-j+c_0} \right| \leq \max \left\{ \frac{c_0}{N} \log(N+1), \frac{c_0}{N} \log \left(1 + \frac{c_0}{N} \right) \right\},$$

again establishing uniform convergence to 0, and thus equation (14).

Since $J(\chi) = \int_0^\chi v(y)\, dy$ for all $\chi \in [0, 1]$, equation (14) and the bounded convergence theorem imply that

$$\lim_{N \to \infty} \max_{\chi \in X^N} \left| \frac{1}{N} \log \frac{\mu_\chi^N}{\mu_0^N} - J(\chi) \right| = \lim_{N \to \infty} \max_{\chi \in X^N} \left| J^N(\chi) - J(\chi) \right| = 0. \tag{15}$$

We obtain uniform convergence in (15) because increasing the length of the interval of integration $[0, \chi]$ that defines $J^N(\chi)$ only worsens the the bound on the rate at which $J^N(\chi)$

converges to $J(\chi)$.

Now let $\chi_*^N = \text{argmax}_{\chi \in \mathcal{X}^N} J(\chi)$ and $\chi_* = \text{argmax}_{\chi \in [0,1]} J(\chi)$. Since J is continuous, we have that $\lim_{N \to \infty} J(\chi_*^N) = J(\chi_*)$. If we can establish that

$$\lim_{N \to \infty} \tfrac{1}{N} \log \mu_{\chi_*^N}^N = 0, \tag{16}$$

then these facts and equation (15) would allow us to conclude that

$$\lim_{N \to \infty} \max_{\chi \in \mathcal{X}^N} \left| \tfrac{1}{N} \log \mu_\chi^N - \left(J(\chi) - J(\chi_*) \right) \right| \tag{17}$$

$$= \lim_{N \to \infty} \max_{\chi \in \mathcal{X}^N} \left| \left(\tfrac{1}{N} \log \frac{\mu_\chi^N}{\mu_0^N} - J(\chi) \right) - \left(\tfrac{1}{N} \log \frac{\mu_{\chi_*^N}^N}{\mu_0^N} - J(\chi_*^N) \right) - \left(J(\chi_*^N) - J(\chi_*) \right) + \tfrac{1}{N} \log \mu_{\chi_*^N}^N \right|$$

$$= 0,$$

proving the theorem.

To establish (16), first suppose to the contrary that there is a sequence $\{N_k\}$ approaching infinity along which the limit in (16) is $-\alpha < 0$. In this case, the reasoning in equation (17) implies that

$$\lim_{N_k \to \infty} \max_{\chi \in \mathcal{X}^N} \left| \tfrac{1}{N_k} \log \mu_\chi^{N_k} - \left(J(\chi) - J(\chi_*) - \alpha \right) \right| = 0.$$

Since $J(\chi) \leq J(\chi_*^N)$ for all $\chi \in \mathcal{X}^N$, it follows that for N_k far enough along the sequence, we have that $\tfrac{1}{N_k} \log \mu_\chi^{N_k} \leq -\tfrac{\alpha}{2}$ for all $\chi \in \mathcal{X}^N$, and hence that

$$\sum_{\chi \in \mathcal{X}_k^N} \mu_\chi^{N_k} = \sum_{\chi \in \mathcal{X}_k^N} \exp\left(N_k \cdot \tfrac{1}{N_k} \log \mu_\chi^{N_k} \right) \leq (N_k + 1) \exp\left(-\frac{\alpha N_k}{2} \right).$$

The last expression vanishes as k grows large, contradicting the fact that μ^{N_k} is a probability measure. Second, suppose contrary to (16) that there is a sequence $\{N_k\}$ approaching infinity along which the limit in (16) is $\alpha > 0$. Then by definition, there is a sequence $\{\delta_k\}$ converging to zero such that

$$\mu_{\chi_*^{N_k}}^{N_k} = \exp\left(N_k(\alpha + \delta_k) \right).$$

The right hand expression grows without bound as k grows large, contradicting the fact that μ^{N_k} is a probability measure. ∎

The Proof of Theorem 4.3

When there are committed agents and a mutation rate of ε as described in protocol (5), the stationary distribution of the stochastic evolutionary process takes the form

$$\frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \prod_{j=1}^{N\chi} \frac{p_{(j-1)/N}^{N,\varepsilon}}{q_{j/N}^{N,\varepsilon}} = \prod_{j=1}^{N\chi} \frac{\frac{N-j+1}{N}\left((1-\varepsilon)\frac{j+c_1-1}{N+c_T-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon\right)}{\frac{j}{N}\left((1-\varepsilon)\frac{N-j+c_0}{N+c_T-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon\right)}. \tag{18}$$

for $\chi \in \mathcal{X}^N - \{0\}$. To prove part (i) of the theorem, observe that for each $\chi \in \mathcal{X}^N - \{0\}$,

$$\lim_{\varepsilon \to 0} \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \prod_{j=1}^{N\chi} \frac{\frac{N-j+1}{N} \cdot \frac{j+c_1-1}{N+c_T-1} r_{01}(F^N(\frac{j-1}{N}))}{\frac{j}{N} \cdot \frac{N-j+c_0}{N+c_T-1} r_{10}(F^N(\frac{j}{N}))} = \frac{\mu_\chi^N}{\mu_0^N}$$

(cf equation (12)). Since $\mathcal{X}^N$ is finite for fixed N, this limit is uniform in χ, so the remainder of the proof is very similar to that of Theorem 4.2.

To prove part (ii) of the theorem, use equation (18) to show that

$$\frac{1}{N} \log \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \frac{1}{N} \sum_{j=1}^{N\chi} \left(\log \frac{N-j+1}{N} - \log \frac{j}{N} + \log \frac{(1-\varepsilon)\frac{j+c_1-1}{N+c_T-1} r_{01}(F^N(\frac{j-1}{N})) + \varepsilon}{(1-\varepsilon)\frac{N-j+c_0}{N+c_T-1} r_{10}(F^N(\frac{j}{N})) + \varepsilon} \right). \tag{19}$$

Since $0 \geq \log\left(\frac{\lceil N\chi \rceil}{N}\right) \geq \log(\chi)$ and $0 \geq \log\left(\frac{N - \lceil N\chi \rceil + 1}{N}\right) \geq \log(1 - \chi)$ for $\chi \in (0,1)$, the dominated convergence theorem implies that the Riemann sum in (19) converges to an integral. In particular, we have that

$$\lim_{N \to \infty} \frac{1}{N} \log \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \int_0^\chi \left(\log \frac{1-y}{y} + \log \frac{(1-\varepsilon)\, y\, r_{01}(F(y)) + \varepsilon}{(1-\varepsilon)(1-y)\, r_{10}(F(y)) + \varepsilon} \right) dy,$$

where the limit, which is taken over those N for which $\chi \in \mathcal{X}^N$, is uniform in χ. Bound (B) and a second application of the dominated convergence theorem then yield

$$\lim_{\varepsilon \to 0} \lim_{N \to \infty} \frac{1}{N} \log \frac{\mu_\chi^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = \int_0^\chi \left(\log \frac{1-y}{y} + \log \frac{y\, r_{01}(F(y))}{(1-y) r_{10}(F(y))} \right) dy$$

$$= \int_0^\chi \log \frac{r_{01}(F(y))}{r_{10}(F(y))}\, dy$$

$$= J(\chi).$$

This limit too is uniform in χ, since the natural bound on the distance to the limit is increasing in the range of integration $[0, \chi]$. Since all limits are uniform in χ, a variation

on the proof of Theorem 4.2 establishes that

$$\lim_{\varepsilon \to 0} \lim_{N \to \infty} \max_{\chi \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_\chi^{N,\varepsilon} - \Delta J(\chi) \right| = 0.$$

This completes the proof of Theorem 4.3. ∎

References

[1] N. Agarwal, Z. Lai, Universal stochastic dominance and equilibrium selection in imitation dynamics, unpublished manuscript, Harvard University (2010).

[2] M. Benaïm, Recursive algorithms, urn processes, and the chaining number of chain recurrent sets, Ergodic Theory and Dynamical Systems 18 (1998) 53–87.

[3] M. Benaïm, J. W. Weibull, Deterministic approximation of stochastic evolution in games, Econometrica 71 (2003) 873–903.

[4] K. Binmore, L. Samuelson, Muddling through: Noisy equilibrium selection, Journal of Economic Theory 74 (1997) 235–265.

[5] K. Binmore, L. Samuelson, R. Vaughan, Musical chairs: Modeling noisy evolution, Games and Economic Behavior 11 (1995) 1–35, erratum, 21 (1997), 325.

[6] J. Björnerstedt, J. W. Weibull, Nash equilibrium and evolution by imitation, in: K. J. Arrow, et al. (eds.), The Rational Foundations of Economic Behavior, St. Martin's Press, New York, 1996, pp. 155–181.

[7] L. E. Blume, Population games, in: W. B. Arthur, S. N. Durlauf, D. A. Lane (eds.), The Economy as an Evolving Complex System II, Addison-Wesley, Reading, MA, 1997, pp. 425–460.

[8] L. E. Blume, How noise matters, Games and Economic Behavior 44 (2003) 251–271.

[9] M. I. Freidlin, A. D. Wentzell, Random Perturbations of Dynamical Systems, 2nd ed., Springer, New York, 1998.

[10] D. Fudenberg, D. Hojman, Stochastic stability in large populations, unpublished manuscript, Harvard University (2009).

[11] D. Fudenberg, L. A. Imhof, Imitation processes with small mutations, Journal of Economic Theory 131 (2006) 251–262.

[12] D. Fudenberg, L. A. Imhof, Monotone imitation dynamics in large populations, Journal of Economic Theory 140 (2008) 229–245.

[13] D. Fudenberg, M. A. Nowak, C. Taylor, L. A. Imhof, Evolutionary game dynamics in finite populations with strong selection and weak mutation, Theoretical Population Biology 70 (2006) 352–363.

[14] J. Hofbauer, Imitation dynamics for games, unpublished manuscript, University of Vienna (1995).

[15] M. Kandori, G. J. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, Econometrica 61 (1993) 29–56.

[16] T. G. Kurtz, Solutions of ordinary differential equations as limits of pure jump Markov processes, Journal of Applied Probability 7 (1970) 49–58.

[17] D. Monderer, L. S. Shapley, Potential games, Games and Economic Behavior 14 (1996) 124–143.

[18] J. R. Norris, Markov Chains, Cambridge University Press, Cambridge, 1997.

[19] M. A. Nowak, A. Sasaki, C. Taylor, D. Fudenberg, Emergence of cooperation and evolutionary stability in finite populations, Nature 428 (2004) 646–650.

[20] G. Roth, W. H. Sandholm, Stochastic approximations with constant step size and differential inclusions, unpublished manuscript, Université de Neuchâtel and University of Wisconsin (2011).

[21] W. H. Sandholm, Potential games with continuous player sets, Journal of Economic Theory 97 (2001) 81–108.

[22] W. H. Sandholm, Simple formulas for stationary distributions and stochastically stable states, Games and Economic Behavior 59 (2007) 154–162.

[23] W. H. Sandholm, Orders of limits for stationary distributions, stochastic dominance, and stochastic stability, Theoretical Economics 5 (2010) 1–26.

[24] W. H. Sandholm, Population Games and Evolutionary Dynamics, MIT Press, Cambridge, 2010.

[25] P. D. Taylor, L. Jonker, Evolutionarily stable strategies and game dynamics, Mathematical Biosciences 40 (1978) 145–156.

[26] A. Traulsen, C. Hauert, Stochastic evolutionary game dynamics, in: H. G. Schuster (ed.), Reviews of Nonlinear Dynamics and Complexity, vol. 2, Wiley, New York, 2009, pp. 25–61.

[27] J. W. Weibull, Evolutionary Game Theory, MIT Press, Cambridge, 1995.

[28] H. P. Young, The evolution of conventions, Econometrica 61 (1993) 57–84.

[29] H. P. Young, Individual Strategy and Social Structure, Princeton University Press, Princeton, 1998.