Imitation processes with small mutations

Drew Fudenberg[a,*], Lorens A. Imhof[b]

[a]*Department of Economics, Harvard University, Cambridge, MA 02138, USA*
[b]*Department of Economics, Bonn University, D-53113 Bonn, Germany*

Abstract

   This note characterizes the impact of adding rare stochastic mutations to an "imitation dynamic," meaning a process with the properties that absent strategies remain absent, and non-homogeneous states are transient. The resulting system will spend almost all of its time at the absorbing states of the no-mutation process. The work of Freidlin and Wentzell [Random Perturbations of Dynamical Systems, Springer, New York, 1984] and its extensions provide a general algorithm for calculating the limit distribution, but this algorithm can be complicated to apply. This note provides a simpler and more intuitive algorithm. Loosely speaking, in a process with K strategies, it is sufficient to find the invariant distribution of a $K \times K$ Markov matrix on the K homogeneous states, where the probability of a transit from "all play i" to "all play j" is the probability of a transition from the state "all agents but 1 play i, 1 plays j" to the state "all play j".
© 2005 Elsevier Inc. All rights reserved.

1. Introduction

   Many papers in economics and evolutionary game theory study various sorts of "imitation dynamics," according to which agents are more likely to adopt strategies that are popular and/or successful. [1] We study these dynamics in a discrete-time, discrete-state

---

   * Corresponding author. Fax: +1 617 495 4341.
    E-mail addresses: dfudenberg@harvard.edu (D. Fudenberg), limhof@uni-bonn.de (L.A. Imhof).

[1] See e.g. [3,8,9,25].

Markov chain. Under the starkest version of these dynamics, it is impossible for agents to adopt a strategy that is not currently in use, so that any "homogeneous" state where all agents use the same strategy is absorbing. Moreover, it is typically the case that all of the interior states are transient, so that the dynamics converges to one of the homogeneous states. Our first result imposes these two conditions; we later relax them to allow a wider class of imitation processes.

Although all of the homogeneous states are absorbing, this does not mean that they are all equally plausible, as they may differ in their robustness to a small probability of mutation or experimentation. One way to formalize the idea that some of the homogeneous states are "more persistent" than others is to assume that a small mutation term makes the system ergodic, and then analyze the limit, as the mutation probability goes to zero, of the invariant distributions. It is easily seen that the resulting system will spend almost all of its time at the absorbing states of the underlying no-mutation process; the work of Freidlin and Wentzell [12] and its extensions provide a general algorithm for calculating the limit distribution. [2] This idea was first applied to equilibrium selection in games by Foster and Young [10,11] in a stochastic differential equation model; Kandori et al. [15] and Young [27] are early applications to equilibrium selection in discrete models.

As we explain below, most of our assumptions are fairly standard and are satisfied by many models in the literature. However, existing techniques for discrete-state models have focused on determining the support of the limit distribution by using order-of-magnitude comparisons for the transitions between the various recurrent classes of the no-mutation process. [3] In contrast, our assumptions imply that nearly every absorbing state can, loosely speaking, be exited with a single mutation, so that the limit distribution may assign non-zero probability to multiple states. Nevertheless, we are able to provide a simple and intuitive algorithm for computing the limit distribution by computing the invariant distribution of an embedded Markov chain on a smaller state space.

The simplest version of our result applies to the "anonymous" case, where the state is simply the number of agents playing each strategy. Here, in a game with K strategies, it is sufficient to find the invariant distribution of a $K \times K$ Markov matrix on the K homogeneous states, where the probability of a transit from "all play i" to "all play j" is the probability of a mutation from i to j times the probability (under the no-mutation dynamics) of a transition from the state "all agents but 1 play i, 1 plays j" to the state "all play j." [4] Although the result was inspired by the ideas of Freidlin and Wentzell, the proof is virtually self-contained and uses only the law of large numbers.

Under our conditions, the details of the dynamics at strictly interior points have no additional impact on the limit distribution. That is, any two processes that meet our conditions and also behave the same way along "edges" where only two strategies are present must lead to the same limit distribution. Intuitively, when mutations are of order ε, the process

---

[2] Duflo [4, pp. 255–258] gives a lucid and easy to use summary of the basic results. Kifer [20] gives a fairly general discrete-time formulation.

[3] This is true for example of Kandori et al. [15] and Young [27], and the results surveyed in Chapter 5 of Fudenberg and Levine [14] and in Chapter 7 of Samuelson [23].

[4] Karandikar et al. [18] use a related construction to characterize the non-degenerate limit distribution of a model with two satisfying players and a continuum of states.

will spend $1 - O(\varepsilon)$ of the time at the homogeneous states, which are the vertices of the state space, $O(\varepsilon)$ time on the edges, and $o(\varepsilon)$ time at interior points. A second implication of our results is that the ratio of the mutation probabilities will matter, even when this ratio is bounded away from zero and infinity; this is related to the fact that a single mutation is enough to cause a transition from one homogeneous state to another.

The more general version of our result allows for arbitrary finite state spaces, and also a wider class of imitation processes, where some absent strategies can be re-introduced. The larger state space allows us to include Ellison's [6] model of local interaction, where the state is the vector of actions being played at each location, the wider class of processes allows us to include best-response-type dynamics as in [24,27] in games with a dominated strategy.

2. Anonymous populations

This section states our assumptions and conclusion for the case of an anonymous population; Section 4 considers a more general state space. Consider a game with K pure strategies denoted by $1, \ldots, K$. We fix a number N of agents, and suppose that the states in the state space $\mathcal{S}$ represent the number of agents playing each strategy. Thus, $\mathcal{S} = \{(x_1, \ldots, x_K) : x_i \in \{0, 1, \ldots, N\}, x_1 + \cdots + x_K = N\}$. For $\varepsilon \geqslant 0$, let $\{X^\varepsilon(t) : t = 0, 1, \ldots\}$ be a homogeneous Markov chain with state space $\mathcal{S}$ and transition probabilities $p_\varepsilon(s, s')$, $s, s' \in \mathcal{S}$. The chain describes the composition of a population, which is assumed to be of constant size. The parameter ε determines the size of mutation rates, as specified below. The chain $\{X^0(t)\}$ describes the evolution without mutations; it has transition probabilities $p_0(s, s')$.

Assumption 1. Under p_0, if at any time a strategy is not played, it will remain absent.

This implies that for each i the homogeneous state in which all players use strategy i is absorbing. We will denote these states by s_i, where $s_1 = (N, 0, \ldots, 0)$, $s_2 = (0, N, 0, \ldots, 0)$, etc.

Assumption 2. Under p_0, all states in $\mathcal{S} \setminus \{s_1, \ldots, s_K\}$ are transient.

Remark. Assumptions 1 and 2 are consistent with any process with the properties that (a) extinct strategies are never re-introduced, and (b) at any state where at least two strategies are being used, there is a positive probability of a transition to a state where the number of agents using the most popular strategy is increased. In particular, in the presence of Assumption 1, Assumption 2 holds if the share of every non-extinct strategy has a positive probability of increasing. Equivalently, given Assumption 1, Assumption 2 holds if starting at any state, there is a positive probability of eventual absorption. These two assumptions are satisfied in the model of Kandori et al. [15] if we specify that the homogeneous states are absorbing (which their assumptions allow) and that all agents adjust their actions in each period (which they call the "best reply dynamic.") Young [27] and Sandholm [24] study variants of the best-response dynamic that satisfy Assumptions 1 and 2 in 2×2 coordination games, but need not do so in general, because extinct strategies are re-introduced if they are a best response to the state, and because there could be interior cycles.

We assume that the transition probabilities $p_\varepsilon(s, s')$ depend continuously on ε, and that for every $\varepsilon > 0$, $\{X^\varepsilon(t)\}$ is irreducible. Let $\pi_\varepsilon = \{\pi_\varepsilon(s)\}_{s \in \mathcal{S}}$ denote the unique invariant distribution of $\{X^\varepsilon(t)\}$, $\varepsilon > 0$. A standard argument implies that for every state $s \in \mathcal{S} \setminus \{s_1, \ldots, s_K\}$, $\lim_{\varepsilon \to 0} \pi_\varepsilon(s) = 0$.[5] We will show how to compute the limit distribution $\lim_{\varepsilon \to 0} \pi_\varepsilon$ by working with a Markov process on state space $\{s_1, \ldots, s_K\}$.

For $j \neq i$ let $s_{i/j}$ denote the almost pure state in which all players use i except for one, which uses j.

Assumption 3. For $i = 1, \ldots, K$,

$$\lim_{\varepsilon \to 0} \frac{p_\varepsilon(s_i, s_{i/j})}{\varepsilon} = \mu_{ij} \geqslant 0, \quad j \neq i. \tag{1}$$

Assumption 4. For $i = 1, \ldots, K$,

$$\lim_{\varepsilon \to 0} \frac{p_\varepsilon(s_i, s)}{\varepsilon} = 0 \quad \text{for all } s \in \mathcal{S} \setminus \{s_i, s_{i/1}, \ldots, s_{i/i-1}, s_{i/i+1}, \ldots, s_{i/K}\}.$$

Combined, these assumptions imply that each transition has a well-defined limiting order, that the probability that a single mutant invades a pure population is $O(\varepsilon)$, while the probability that two or more mutants invade simultaneously is $o(\varepsilon)$. We do allow that for some pure population states the invasion probabilities of certain mutant types are $o(\varepsilon)$. That is, we allow that $\mu_{ij} = 0$ for some i, j; the precise requirements are given below in connection with the matrix Λ. Note that Assumptions 3 and 4 apply only to transitions from the homogeneous states s_i; there is no restriction on the rate with which $p_\varepsilon(s, s')$ converges to $p_0(s, s')$ for $s \in \mathcal{S} \setminus \{s_1, \ldots, s_K\}$.

In biological applications it may be natural to assume that all of the mutation probabilities $p_\varepsilon(s_i, s_{i/j})$ are of the same order, so that $\mu_{ij} > 0$ for all i, j with $i \neq j$. However, in the context of learning models, where mutations may be interpreted as mistakes, certain types of mistakes may not occur or more generally may become infinitely less likely than others. In this case, $p_\varepsilon(s_i, s_{i/j}) = 0$ for some i, j and all $\varepsilon > 0$, and so $\mu_{ij} = 0$.

Under the no-mutation process, absent strategies remain absent. Thus if $X^0(0) = s_{i/j}$, then $\{X^0(t)\}$ will be absorbed in s_i or s_j. Let ρ_{ij} denote the probability that the no-mutation process will be absorbed in s_j. That is, ρ_{ij} is the probability that a population of i-players invaded by a single j-player evolves without mutations to a population of all j-players. Define a $K \times K$ matrix $\Lambda = (\Lambda_{ij})$ by

$$\Lambda_{ij} = \mu_{ij}\rho_{ij}, \quad j \neq i, \quad \Lambda_{ii} = 1 - \sum_{j \neq i} \mu_{ij}\rho_{ij}.$$

---

[5] This is an implication of Theorem 4 of Young [27], but it can be shown directly with a much shorter argument. Assume there exists some state $\bar{s} \in \mathcal{S} \setminus \{s_1, \ldots, s_K\}$ such that $\pi_\varepsilon(\bar{s}) \not\to 0$ as $\varepsilon \to 0$. Then there is a sequence $(\varepsilon(k))_{k=1}^\infty \subset (0, \infty)$ with $\varepsilon(k) \to 0$ and a distribution $\pi^* = \{\pi^*(s)\}_{s \in \mathcal{S}}$ with $\pi^*(\bar{s}) > 0$ such that $\pi_{\varepsilon(k)}(s) \to \pi^*(s)$ for every $s \in \mathcal{S}$. Since $\pi_{\varepsilon(k)}(s) = \sum_{r \in \mathcal{S}} \pi_{\varepsilon(k)}(r) p_{\varepsilon(k)}(r, s)$ for every k, it follows that $\pi^*(s) = \sum_{r \in \mathcal{S}} \pi^*(r) p_0(r, s)$, that is, π^* is an invariant distribution for the limit process. But then $\pi^*(\bar{s}) > 0$ implies that $\bar{s}$ is recurrent under the no-mutation process (see [5, Theorem 4.5, p. 306]) which yields a contradiction.

For $i \neq j$, Λ_{ij} is the probability that in a population of i-players a single j-mutant occurs times the probability that this mutant takes over the whole population so that only j-players remain. Every row sum of Λ is 1, and, rescaling the parameter ε if necessary, we may also assume that every diagonal entry of Λ is non-negative. [6] Thus Λ can be regarded as the transition matrix of a Markov chain with state space $\{s_1, \ldots, s_K\}$. The following condition on Λ is the key to our result:

Assumption 5. There is a unique vector $\lambda = (\lambda_1, \ldots, \lambda_K)$ with

$$\lambda\Lambda = \lambda, \quad \lambda_1 + \cdots + \lambda_K = 1, \quad \lambda_1 \geqslant 0, \ldots, \lambda_K \geqslant 0. \tag{2}$$

Condition (2) means that λ is an invariant distribution of the chain. In particular, if for every pair of distinct strategies i, j, there exists a sequence $i_0, i_1, \ldots, i_m$ with $i_0 = i$ and $i_m = j$ such that $\Lambda_{i_k, i_{k+1}} > 0$ for $k = 0, \ldots, m-1$, then Λ is an irreducible matrix and there exists a unique λ satisfying (2).

An even simpler sufficient condition is the following:

Assumption 5′. For all pairs of different strategies i, j, $\mu_{ij} > 0$ and $\rho_{ij} > 0$.

Assumption 5 implies that there is at most one state that cannot be exited with a single mutation. Assumption 5′ is stronger: it says that there is no such state, as a single mutation can cause a transition from any basin to any other. This assumption is satisfied in any imitation model where every strategy that is played by at least one agent has a positive probability of increasing its share. In a more restricted setting, Noldeke and Samuelson [21] show that if a state is in the support of the limit distribution, then so is any other state in an absorbing set that can be reached with a single mutation; when Λ is irreducible, their condition applies to every homogeneous state. Assumption 5 rules out the cases studied by Kandori et al. [15] and Young [27], where $\rho_{ij} = 0$ for all pairs i, j, and Λ is a diagonal matrix.

The goal of this paper is to characterize the limit, as the mutation rate goes to zero, of the invariant distributions of the perturbed chains. Theorem 1 says that this limit distribution is the unique invariant distribution of Assumption 5. Thus, the limit distribution can be computed by solving a system of K linear equalities.

Theorem 1. *In the anonymous-population model, under Assumptions* 1–5,

$$\lim_{\varepsilon \to 0} \pi_\varepsilon(s_i) = \lambda_i \; for \; i = 1, \ldots, K.$$

Proof. This is a consequence of the more general Theorem 2. $\quad\square$

Note that in general multiplying one of the μ_{ij} by a positive constant, leaving the other μ's unchanged, will change Λ and hence change the limit distribution. This is in contrast

---

[6] Replacing X^ε with $\tilde{X}^\varepsilon = X^{c\varepsilon}$ and p_ε with $\tilde{p}_\varepsilon = p_{c\varepsilon}$ does not change the limit behavior of the ergodic distributions. From (1), $\lim_{\varepsilon \to 0} \tilde{p}_\varepsilon(s_i, s_{i/j})/\varepsilon = c\mu_{ij} = \tilde{\mu}_{ij}$ and so, if $c > 0$ is sufficiently small, $\tilde{\Lambda}_{ii} = 1 - \sum_{j \neq i} \tilde{\mu}_{ij}\rho_{ij} \geqslant 0$ for $i = 1, \ldots, K$.

to, e.g., Kandori et al. [15] where the ratio of the mutation probabilities has no effect on the limit distribution so long as it is bounded away from 0 and infinity.[7] This dependence on the ratio of the two sorts of mutations is a consequence of the fact that for fixed N the limit distribution assigns positive probability to more than one point, which in turn is related to the fact that a single mutation is enough to cause a transition from one homogeneous state to another. The "bandwagon properties" analyzed by Kandori and Rob [17] lead to the conclusion that it is sufficient to analyze direct transitions from one homogeneous state to another, but their assumptions are rather different, as they impose restrictions on the payoff matrix of the game, and do not assume that a single mutation can upset a homogeneous steady state.

3. Examples

Example 1. Consider a game with two strategies, i.e. $K = 2$, and suppose $\mu_{12} = \mu_{21} = \mu > 0$ and that ρ_{12} and ρ_{21} are strictly positive as well. Then

$$\Lambda = \begin{bmatrix} 1 - \mu\rho_{12} & \mu\rho_{12} \\ \mu\rho_{21} & 1 - \mu\rho_{21} \end{bmatrix}$$

is irreducible, and strictly positive for small enough μ. Hence, it has unique invariant distribution

$$\lambda = \left(\frac{\rho_{21}}{\rho_{12} + \rho_{21}}, \frac{\rho_{12}}{\rho_{12} + \rho_{21}} \right)$$

and π_ε converges weakly to the probability measure on $\mathcal{S}$ that puts mass $\frac{\rho_{21}}{\rho_{12}+\rho_{21}}$ and $\frac{\rho_{12}}{\rho_{12}+\rho_{21}}$ on s_1 and s_2, respectively.

Example 2. As a less trivial example, we consider the case of three pure strategies, 1, 2, 3. We suppose that the no-mutation process is the frequency-dependent Moran process studied by Fudenberg et al. [13], where the payoff matrix is $(a_{ij})_{i,j=1}^3$, and each pure strategy is a strict Nash equilibrium.[8] Suppose the population size is N and $\mu_{ij} = \mu \in (0, \frac{1}{2})$ for all $i \neq j$. Using standard results on birth–death chains (see e.g. [19, p. 113, Exercise 4]), Fudenberg et al. [13] show that the absorption probabilities ρ_{ji} are given by

$$\rho_{ji} = \frac{1}{1 + \sum_{k=1}^{N-1} \prod_{l=1}^{k} \frac{g_l(i, j)}{f_l(i, j)}},$$

---

[7] Bergin and Lipman [2] show that allowing the ratio of mutation probabilities to converge to 0 or infinity does change the limit distribution in the model of Kandori et al. [15].

[8] Fudenberg et al. [13] compute the limit distribution of this process in general 2×2 games, and compute the limit of the limit distribution as the population becomes infinite in these games and in 3×3 coordination games. Taylor et al. [26] compare the absorption probabilities of the no-mutation version of the process to the benchmark of neutral selection; Nowak et al. [22] apply that calculation to the two strategies "All D" and "Tit for Tat" in the finitely repeated prisoner's dilemma.

where

$$f_l(i, j) = \frac{a_{ii}(l-1) + a_{ij}(N-l)}{N-1}, \quad g_l(i, j) = \frac{a_{ji}l + a_{jj}(N-l-1)}{N-1},$$

and the simple functional form comes from the fact that the ratio of the transition probabilities is simply the ratio of the prevailing payoffs f and g.

The matrix Λ is

$$\Lambda = \begin{bmatrix} 1 - \mu(\rho_{12} + \rho_{13}) & \mu\rho_{12} & \mu\rho_{13} \\ \mu\rho_{21} & 1 - \mu(\rho_{21} + \rho_{23}) & \mu\rho_{23} \\ \mu\rho_{31} & \mu\rho_{32} & 1 - \mu(\rho_{31} + \rho_{32}) \end{bmatrix}.$$

This matrix is irreducible, and its unique invariant distribution is $\lambda = [\gamma_1 + \gamma_2 + \gamma_3]^{-1}$ $(\gamma_1, \gamma_2, \gamma_3)$, where

$$\gamma_1 = \rho_{21}\rho_{31} + \rho_{21}\rho_{32} + \rho_{31}\rho_{23},$$
$$\gamma_2 = \rho_{31}\rho_{12} + \rho_{12}\rho_{32} + \rho_{32}\rho_{13},$$
$$\gamma_3 = \rho_{21}\rho_{13} + \rho_{12}\rho_{23} + \rho_{13}\rho_{23}.$$

We conclude that the ergodic distribution π_ε converges weakly to the distribution on $\mathcal{S}$ that puts mass λ_i at s_i, $i = 1, 2, 3$. Thus when the mutation probabilities are small, X^ε spends time at the states s_1, s_2, s_3 roughly in proportions $\lambda_1 : \lambda_2 : \lambda_3$ and almost no time at other states.

Consider the pay-off matrix

$$A = \begin{pmatrix} 1.0 & 1.0 & 2.0 \\ 0.1 & 2.5 & 1.0 \\ 0.1 & 1.0 & 4.0 \end{pmatrix}$$

and suppose $\mu_{ij} = 0.1$ for all $i \neq j$. Using the above expression for λ, we obtain that for $N = 10$, $\lambda = (0.716, 0.061, 0.223)$, whereas for $N = 50$, $\lambda = (0.082, 0.0, 0.918)$. Suppose now that mutations correspond to mistakes and that players of strategy 2 are less likely to make mistakes. For instance, suppose that $\mu_{21} = \mu_{23} = 0.001$ and that the other μ_{ij} remain unchanged. Then, for $N = 10$, $\lambda = (0.101, 0.867, 0.032)$, but for $N = 50$, $\lambda = (0.082, 0.004, 0.914)$.

4. General imitation processes

We now extend our result to cases where not all homogeneous states are absorbing, and to larger state spaces that allow e.g. the state to encode the location of each agent. To motivate this latter extension, consider the following example.

Example 3. Specialize the local interaction model of Ellison [6] to the case where each agent only interacts with his two closest neighbors. (This paper studies the myopic best reply dynamics, but this is a form of imitation dynamics in coordination games.) If there is an odd number of agents, the homogeneous states are the only recurrent classes of the

no-mutation process.[9] However, this example is not covered by Theorem 1, since the state space tracks the strategy played at each location, so a single mutation can send the state "all A" to any one of N different configurations.

To allow for the system to depend on the location or identity of the agents playing each strategy, we now consider an arbitrary finite state space $\mathcal{S}$, and a family of homogeneous Markov chains $\{X^\varepsilon(t) : t = 0, 1, \ldots\}$ on $\mathcal{S}$ with transition probabilities $p_\varepsilon(s, s')$, $s, s' \in \mathcal{S}$. As before, the chain $\{X^0(t)\}$ describes the evolution without mutations; it has transition probabilities $p_0(s, s')$.

Assumption 6. Under p_0, there exists at least one absorbing state.

This parallels Assumption 1, but is weaker, even with the state space of Section 2, as it allows some homogeneous states not to be absorbing. This might be the case if there was immigration of some strategies even when they are extinct, perhaps because these strategies are being played by agents in a different but connected population.

Denote by s_i, $i = 1, \ldots, K$, the states that are absorbing under p_0.

Assumption 7. All states in $\mathcal{S} \setminus \{s_1, \ldots, s_K\}$ are transient under p_0.

This parallels Assumption 2; in that section the absorbing states were the homogeneous states. Assume that the transition probabilities $p_\varepsilon(s, s')$ depend continuously on ε, and that for every $\varepsilon > 0$, $\{X^\varepsilon(t)\}$ is irreducible. Let $\pi_\varepsilon = \{\pi_\varepsilon(s)\}_{s \in \mathcal{S}}$ denote the unique invariant distribution of $\{X^\varepsilon(t)\}$, $\varepsilon > 0$. Again, $\lim_{\varepsilon \to 0} \pi_\varepsilon(s) = 0$ for every $s \in \mathcal{S} \setminus \{s_1, \ldots, s_K\}$.

Assumption 8. For every $i = 1, \ldots, K$ and every $s \in \mathcal{S} \setminus \{s_i\}$, $\mu_{is} = \lim_{\varepsilon \to 0} \frac{p_\varepsilon(s_i, s)}{\varepsilon}$ exists.

Let S_i denote the set of states s in $\mathcal{S} \setminus \{s_i\}$ such that $\mu_{is} > 0$; these are the states that can be reached from s_i "with one mutation." In the model of Section 2, these were a subset of the states $s_{i/j}$ where all agents but 1 play i. Thus Assumptions 3 and 4 together imply Assumption 8, and also give additional structure on the sets S_i. Working on an arbitrary state space, we do not have much a priori idea of what the absorbing states or the sets S_i might look like; and instead define them from the limit of the transition probabilities, so Assumption 8 does not imply Assumptions 3 and 4.

Let ρ_{sj} denote the probability that starting at s the no-mutation process will be absorbed in s_j. (With this notation, the absorption probabilities ρ_{ij} used in Section 2 are denoted by $\rho_{s_{i/j}, j}$.) Define a $K \times K$ matrix $\Lambda = (\Lambda_{ij})$ by $\Lambda_{ij} = \sum_{s \in S_i} \mu_{is} \rho_{sj}$ for $i \neq j$ and $\Lambda_{ii} = 1 - \sum_{j \neq i} \Lambda_{ij}$. As before, every row sum of Λ is 1, and, rescaling the parameter ε if necessary, we may also assume that every diagonal entry of Λ is non-negative. Thus, Λ can be regarded as the transition matrix of a Markov chain with state space $\{s_1, \ldots, s_K\}$.

---

[9] With an even number of players, there is also a two-cycle.

Assumption 9. There is a unique vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_K)$ with

$$\boldsymbol{\lambda}\Lambda = \boldsymbol{\lambda}, \quad \lambda_1 + \cdots + \lambda_K = 1, \quad \lambda_1 \geqslant 0, \ldots, \lambda_K \geqslant 0.$$

Note that Assumption 9 reduces to Assumption 5 for the model of Section 2.

As with Theorem 1, Theorem 2 says that the limit distribution is the invariant distribution corresponding to Λ. The first part of the proof shows that it is sufficient to compute the invariant distribution of the embedded Markov chain on the homogeneous states, the second part of the proof calculates what this distribution is. Theorem 4 of Young [27] determines which states belong to the support of the limit distribution by analyzing the first-order terms in the invariant distribution of the embedded chain. [10] Our theorem provides a simple means to directly calculate the limit distribution for processes that meet our conditions.

Theorem 2. *In the general imitation model, under Assumptions 6–9,*

$$\lim_{\varepsilon \to 0} \pi_\varepsilon(s_i) = \lambda_i \quad \text{for } i = 1, \ldots, K.$$

Proof. For $\varepsilon \geqslant 0$, let $\{Y^\varepsilon(t)\}$ be the Markov chain obtained from observing $\{X^\varepsilon(t)\}$ only when it is in $\{s_1, \ldots, s_K\}$. Specifically, $Y^\varepsilon(t) = X^\varepsilon(\tau_t^\varepsilon)$, $t = 0, 1, \ldots$, where

$$\tau_0^\varepsilon = \min\{u \geqslant 0 : X^\varepsilon(u) \in \{s_1, \ldots, s_K\}\},$$
$$\tau_{t+1}^\varepsilon = \min\{u > \tau_t^\varepsilon : X^\varepsilon(u) \in \{s_1, \ldots, s_K\}\}, \quad t \geqslant 0.$$

Note that $P\{\tau_t^\varepsilon < \infty\} = 1$ for every $t \geqslant 0$. For $\varepsilon > 0$, $\{X^\varepsilon(t)\}$ is irreducible, and so $\{Y^\varepsilon(t)\}$ is irreducible as well. Let $\mathbf{v}_\varepsilon = (v_\varepsilon(s_1), \ldots, v_\varepsilon(s_K))$ denote the ergodic distribution of $\{Y^\varepsilon(t)\}$. According to the strong law of large numbers for Markov chains (see e.g. [5, Theorems 4.6 and 5.1, p. 307 and 312]), for $i = 1, \ldots, K$,

$$\lim_{T \to \infty} \frac{\#\{t \leqslant T : X^\varepsilon(t) = s_i\}}{T + 1} = \pi_\varepsilon(s_i), \quad \lim_{T \to \infty} \frac{\#\{t \leqslant T : Y^\varepsilon(t) = s_i\}}{T + 1} = v_\varepsilon(s_i),$$

where $\#A$ denotes the number of elements of the set A. For every T,

$$\#\{t \leqslant T : X^\varepsilon(t) = s_i\} \leqslant \#\{t \leqslant T : Y^\varepsilon(t) = s_i\},$$

and so

$$\pi_\varepsilon(s_i) \leqslant v_\varepsilon(s_i), \quad i = 1, \ldots, K, \ \varepsilon > 0.$$

Since $\sum_{i=1}^K \pi_\varepsilon(s_i) \to 1$ as $\varepsilon \to 0$, it follows that if $v_\varepsilon(s_1), \ldots, v_\varepsilon(s_K)$ converge as $\varepsilon \to 0$, then so do $\pi_\varepsilon(s_1), \ldots, \pi_\varepsilon(s_K)$ and the respective limits must coincide. It remains to show that $\mathbf{v}_\varepsilon \to \boldsymbol{\lambda}$.

---

[10] See Kandori and Rob [16] for a simpler version of the argument. The result is the discrete-time, discrete-state version of Freidlin and Wentzell's continuous-time, continuum-state analysis; Kifer [20] treats the case of discrete-time and a continuum of states. Ellison [7] gives a fairly general sufficient condition for the limit distribution to be concentrated on a single point.

To determine the limit behavior of the transition probabilities $q_\varepsilon(s_i, s_j)$ of the embedded chain $\{Y^\varepsilon(t)\}$ note first that

$$q_\varepsilon(s_i, s_j) = P\left\{X^\varepsilon(\tau_1^\varepsilon) = s_j | X^\varepsilon(0) = s_i\right\}$$

$$= \sum_{s \in \mathcal{S}} p_\varepsilon(s_i, s) P\left\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\right\}. \qquad (3)$$

Consider the case $i \neq j$. Then

$$P\left\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s_i\right\} = P\left\{X^\varepsilon(0) = s_j | X^\varepsilon(0) = s_i\right\} = 0.$$

Hence, by (3) and Assumption 8,

$$\lim_{\varepsilon \to 0} \frac{q_\varepsilon(s_i, s_j)}{\varepsilon} = \lim_{\varepsilon \to 0} \sum_{s \in \mathcal{S} \setminus \{s_i\}} \frac{p_\varepsilon(s_i, s)}{\varepsilon} P\left\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\right\}$$

$$= \lim_{\varepsilon \to 0} \sum_{s \in \mathcal{S} \setminus \{s_i\}} \mu_{is} P\left\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\right\}$$

$$= \lim_{\varepsilon \to 0} \sum_{s \in \mathcal{S}_i} \mu_{is} P\left\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\right\}. \qquad (4)$$

For every $s^* \in \mathcal{S}_i$,

$$\lim_{\varepsilon \to 0} P\left\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s^*\right\}$$

$$= \lim_{\varepsilon \to 0} \sum_{t=0}^{\infty} P\left\{X^\varepsilon(t) = s_j, \tau_0^\varepsilon = t | X^\varepsilon(0) = s^*\right\}. \qquad (5)$$

Assumption 7, that all states in $\mathcal{S} \setminus \{s_1, \ldots, s_K\}$ are transient for $\{X^0(t)\}$, implies that there exist $T < \infty$ and $\delta > 0$ such that $P\{\tau_0^0 < T | X^0(0) = s\} > 2\delta$ for all $s \in \mathcal{S} \setminus \{s_1, \ldots, s_K\}$. The exit probability $P\{\tau_0^\varepsilon < T | X^\varepsilon(0) = s\}$ can be expressed as a finite sum of finite products of one-step transition probabilities, and since the transition probabilities are continuous functions of ε, so is the exit probability.

Therefore, for some $\varepsilon_0 > 0$, $P\{\tau_0^\varepsilon < T | X^\varepsilon(0) = s\} > \delta$ for all $s \in \mathcal{S} \setminus \{s_1, \ldots, s_K\}$ and all $0 \leqslant \varepsilon \leqslant \varepsilon_0$. Consequently, $P\{\tau_0^\varepsilon \geqslant mT | X^\varepsilon(0) = s\} \leqslant (1-\delta)^m$ for all $s \in \mathcal{S} \setminus \{s_1, \ldots, s_K\}$, $m = 0, 1, \ldots$, and all $0 \leqslant \varepsilon \leqslant \varepsilon_0$. In particular, if $\lfloor t/T \rfloor$ denotes the largest integer less than or equal to t/T,

$$P\left\{X^\varepsilon(t) = s_j, \tau_0^\varepsilon = t | X^\varepsilon(0) = s^*\right\} \leqslant P\left\{\tau_0^\varepsilon \geqslant \lfloor t/T \rfloor T | X^\varepsilon(0) = s^*\right\}$$

$$\leqslant (1-\delta)^{\lfloor t/T \rfloor} \leqslant (1-\delta)^{-1+t/T}$$

for every $t = 0, 1, \ldots$ and $0 \leqslant \varepsilon \leqslant \varepsilon_0$. Thus, by the Weierstrass M-test, the series in (5) converges uniformly on $[0, \varepsilon_0]$, so that we can interchange limit and infinite sum, see

e.g. Apostol [1], Theorems 9.6 and 9.7, p. 223. Hence,

$$\lim_{\varepsilon \to 0} P\left\{X^{\varepsilon}(\tau_0^{\varepsilon}) = s_j | X^{\varepsilon}(0) = s^*\right\} = \sum_{t=0}^{\infty} \lim_{\varepsilon \to 0} P\left\{X^{\varepsilon}(t) = s_j, \tau_0^{\varepsilon} = t | X^{\varepsilon}(0) = s^*\right\}$$

$$= \sum_{t=0}^{\infty} P\left\{X^0(t) = s_j, \tau_0^0 = t | X^0(0) = s^*\right\}$$

$$= P\left\{X^0(\tau_0^0) = s_j | X^0(0) = s^*\right\}$$

$$= \rho_{s^* j}.$$

Inserting this into (4) we obtain that

$$\lim_{\varepsilon \to 0} \frac{q_{\varepsilon}(s_i, s_j)}{\varepsilon} = \sum_{s \in \mathcal{S}_i} \mu_{is} \rho_{sj} = \Lambda_{ij}, \quad i \neq j.$$

Therefore, since $\sum_{j=1}^{n} q_{\varepsilon}(s_i, s_j) = 1$,

$$\lim_{\varepsilon \to 0} \frac{q_{\varepsilon}(s_i, s_i) - 1}{\varepsilon} = \lim_{\varepsilon \to 0} \sum_{j \neq i} \frac{-q_{\varepsilon}(s_i, s_j)}{\varepsilon} = -\sum_{j \neq i} \Lambda_{ij} = \Lambda_{ii} - 1.$$

In matrix notation,

$$\lim_{\varepsilon \to 0} \frac{1}{\varepsilon}[Q_{\varepsilon} - I_n] = \Lambda - I_n.$$

For every $\varepsilon > 0$, $\mathbf{v}_{\varepsilon}$ is the unique non-negative solution to

$$\mathbf{v}_{\varepsilon}\frac{1}{\varepsilon}[Q_{\varepsilon} - I_n] = 0, \quad v_{\varepsilon}(s_1) + \cdots + v_{\varepsilon}(s_K) = 1,$$

and, by Assumption 9, $\boldsymbol{\lambda}$ is the unique non-negative solution to

$$\boldsymbol{\lambda}[\Lambda - I_n] = 0, \quad \lambda_1 + \cdots + \lambda_K = 1.$$

It is now obvious that $\mathbf{v}_{\varepsilon} \to \boldsymbol{\lambda}$. □

Example 3 (*revisited*). The absorbing states are the homogeneous points s_1 (all play A) and s_2 (all play B); the states in S_1 and S_2 are those with $N - 1$ A's and $N - 1$ B's, respectively. Suppose that strategy A is risk-dominant. In a 2×2 game with 2 neighbors, a single mutation is not enough to leave the basin of the risk-dominant equilibrium, that is $\rho_{s1} = 1$ for all $s \in S_1$, so $\Lambda_{12} = \sum_{s \in S_1} \mu_{1s} \rho_{s2} = 0$. Moreover, $\rho_{s1} = 1$ for all $s \in S_2$, so $\Lambda_{21} > 0$, and

$$\Lambda = \begin{bmatrix} 1 & 0 \\ \Lambda_{21} & 1 - \Lambda_{21} \end{bmatrix}$$

has a unique invariant distribution, namely (1,0). Thus, we recover Ellison's conclusion that the limit distribution assigns probability 1 to the risk-dominant equilibrium. Moreover, the

same conclusion obtains provided that $\rho_{s1} > 0$ for some $s \in S_2$, while the limit distribution gives positive probability to both absorbing states if $\rho_{s1} < 1$ for some $s \in S_1$ and $\rho_{s1} > 0$ for some $s \in S_2$.

Acknowledgments

References

[1] T.M. Apostol, Mathematical Analysis, Addison-Wesley, Reading, MA, 1974.

[2] J. Bergin, B. Lipman, Evolution with state dependent mutations, Econometrica 64 (1996) 943–956.

[3] J. Bjonerstedt, J. Weibull, Nash equilibrium and evolution by imitation, in: K. Arrow et al. (Eds.), The Rational Foundations of Economic Analysis, MacMillan, London, 1995.

[4] M. Duflo, Algorithmes Stochastiques, Springer, New York, 1996.

[5] R. Durrett, Probability: Theory and Examples, Duxbury Press, Belmont, 1996.

[6] G. Ellison, Learning, local interaction, and coordination, Econometrica 61 (1993) 1047–1071.

[7] G. Ellison, Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution, Rev. Econ. Stud. 67 (2000) 17–45.

[8] G. Ellison, D. Fudenberg, Rules of thumb for social learning, J. Polit. Economy 101 (1993) 612–643.

[9] G. Ellison, D. Fudenberg, Word of mouth communication and social learning, Quart. J. Econ. 110 (1995) 93–126.

[10] D. Foster, P. Young, Stochastic evolutionary game dynamics, Theoretical Population Biol. 38 (1990) 219–232.

[11] D. Foster, P. Young, A correction to the paper "Stochastic evolutionary game dynamics", Theoretical Population Biol. 51 (1997) 77–78.

[12] M.I. Freidlin, A.D. Wentzell, Random Perturbations of Dynamical Systems, Springer, New York, 1984.

[13] D. Fudenberg, L. Imhof, M. Nowak, C. Taylor, Stochastic evolution as a generalized Moran Process, Mimeo, 2004.

[14] D. Fudenberg, D. Levine, The Theory of Learning in Games, MIT Press, Cambridge, MA, 1998.

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

[16] M. Kandori, R. Rob, Evolution of equilibria in the long run, J. Econ. Theory 65 (1995) 383–414.

[17] M. Kandori, R. Rob, Bandwagon effects and long run technology choice, Games Econ. Behav. 22 (1998) 30–60.

[18] R. Karandikar, D. Mookherjee, D. Ray, F. Vega-Redondo, Evolving aspirations and cooperation, J. Econ. Theory 80 (1998) 292–331.

[19] S. Karlin, H. Taylor, A First Course in Stochastic Processes, Academic Press, New York, 1975.

[20] Y. Kifer, A discrete-time version of the Wentzell–Freidlin theory, Ann. Probability 18 (1990) 1676–1692.

[21] G. Noldeke, L. Samuelson, An evolutionary analysis of backward and forward induction, Games Econ. Behav. 5 (1993) 425–454.

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

[23] L. Samuelson, Evolutionary Games and Equilibrium Selection, MIT Press, Cambridge, MA, 1997.

[24] W. Sandholm, Simple and clever decision rules for a model of evolution, Econ. Lett. 61 (1998) 165–170.

[25] K. Schlag, Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, J. Econ. Theory 78 (1998) 130–156.

[26] C. Taylor, D. Fudenberg, A. Sasaki, M.A. Nowak, Evolutionary game dynamics in finite populations, Bull. Math. Biol. 66 (2004) 1621–1644.

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