



Imitation Processes with Small Mutations

Citation

Fudenberg, Drew, and Lorenz A. Imhof. 2006. Imitation processes with small mutations. *Journal of Economic Theory* 131, no. 1: 251-262.

Published Version

<http://dx.doi.org/10.1016/j.jet.2005.04.006>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3190369>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Imitation Processes with Small Mutations

Drew Fudenberg
Department of Economics,
Harvard University
Cambridge, MA 02138

Lorens A. Imhof
Institute of Statistics
Aachen University
52056 Aachen
Germany

September 9, 2004

Abstract

This note characterizes the impact of adding rare stochastic mutations to an “imitation dynamic,” meaning a process with the properties that any state where all agents use the same strategy is absorbing, and all other states are transient. The work of Freidlin and Wentzell [10] and its extensions implies that the resulting system will spend almost all of its time at the absorbing states of the no-mutation process, and provides 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 .”

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 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. This does not mean that all of these states are equally plausible, as

¹See e.g. [3], [7], [8], [18].

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. The work of Freidlin and Wentzell [10] and its subsequent applications and extensions² shows that the resulting system will spend almost all of its time at the absorbing states of the underlying no-mutation process, and provides a general algorithm for calculating the limit distribution, but this algorithm can be complicated to apply, especially in games with more than two strategies.

This note provides a simpler and more intuitive algorithm for imitation processes with the property that a single mutation is enough for there to be a positive probability of a transition away from any homogeneous state. The simplest version of the 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 transition from the state “all agents but 1 play i , 1 plays j ” to the state “all play j .” The more general version of our result allows for larger state spaces, as in Ellison’s [5] model of local interaction.

One implication of the result is that the properties of the no-mutation process at strictly interior points have no impact on the limit distribution. Intuitively, when mutations are of order ε , the process 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” where only two strategies have positive probability, and $o(\varepsilon)$ time at interior points. A second implication 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.

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 $1, \dots, K$. We fix a number N of agents, and suppose that the state space S is simply the number of agents playing each strategy. For $\varepsilon \geq 0$, let $\{X^\varepsilon(t) : t = 0, 1, \dots\}$ 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 as-

²[4] pp. 255-258 gives a lucid and easy to use summary of the basic results. [15] gives a fairly general discrete-time formulation. [12] and [20] are early applications to equilibrium selection in games.

sumed 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 $p_0(s, s')$. We assume that under p_0 , if at any time a strategy is not played, it will remain absent; thus for each i the homogeneous state s_i in which all players use strategy i is absorbing. We further assume that all other states in \mathcal{S} are transient under p_0 .

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 $s \in \mathcal{S} \setminus \{s_1, \dots, s_K\}$, $\lim_{\varepsilon \rightarrow 0} \pi_\varepsilon(s) = 0$.³ We will show how to compute the limit distribution $\lim_{\varepsilon \rightarrow 0} \pi_\varepsilon$ by working with a Markov process on state space $\{s_1, \dots, s_K\}$.

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 1: For $i = 1, \dots, K$,

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

Assumption 2: For $i = 1, \dots, K$,

$$\lim_{\varepsilon \rightarrow 0} \frac{p_\varepsilon(s_i, s)}{\varepsilon} = 0 \quad \text{for all } s \in \mathcal{S} \setminus \{s_i, s_{i/1}, \dots, s_{i/i-1}, s_{i/i+1}, \dots, s_{i/K}\}. \quad (2)$$

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 .

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

³This is an implication of Theorem 4 of [20], but it can be shown directly with a much shorter argument. Assume there exists some state $\bar{s} \in \mathcal{S} \setminus \{s_1, \dots, s_K\}$ such that $\pi_\varepsilon(\bar{s}) \not\rightarrow 0$ as $\varepsilon \rightarrow 0$. Then there is a sequence $(\varepsilon(k))_{k=1}^\infty \subset (0, \infty)$ with $\varepsilon(k) \rightarrow 0$ and a distribution $\pi^* = \{\pi^*(s)\}_{s \in \mathcal{S}}$ with $\pi^*(\bar{s}) > 0$ such that $\pi_{\varepsilon(k)}(s) \rightarrow \pi^*(s)$ 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 Theorem 4.5 in Durrett [9], page 306) which yields a contradiction.

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

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

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.⁴ Thus Λ can be regarded as the transition matrix of a Markov chain with state space $\{s_1, \dots, s_K\}$. The following condition on Λ is the key to our result:

Assumption 3: There is a unique vector $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_K)$ with

$$\boldsymbol{\lambda}\Lambda = \boldsymbol{\lambda}, \quad \lambda_1 + \dots + \lambda_K = 1, \quad \lambda_1 \geq 0, \dots, \lambda_K \geq 0. \quad (3)$$

Condition (3) means that $\boldsymbol{\lambda}$ 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, \dots, i_m with $i_0 = i$ and $i_m = j$ such that $\Lambda_{i_k, i_{k+1}} > 0$ for $k = 0, \dots, m-1$, then Λ is an irreducible matrix and there exists a unique $\boldsymbol{\lambda}$ satisfying (3). An even simpler sufficient condition is that for all pairs of different strategies i, j , $\mu_{ij} > 0$ and $\rho_{ij} > 0$, so that a single mutation can cause a transition from any basin to any other. In a more restricted setting, Noldeke and Samuelson [16] show that if a state is in the support of the limit distribution, then so is any other state that can be reached with a single mutation; when Λ is irreducible, their condition applies to every homogeneous state. Conversely, Assumption 3 implies that there is at most one state that cannot be exited with a single mutation. The assumption therefore rules out the cases studied by [12] and [20], where $\rho_{ij} = 0$ for all pairs i, j , and Λ is a diagonal matrix. Theorem 1 says that the limit distribution is the unique invariant distribution of Assumption 3.

Theorem 1 *In the anonymous-population model, under Assumptions 1 through 3, $\lim_{\varepsilon \rightarrow 0} \pi_\varepsilon(s_i) = \lambda_i$ for $i = 1, \dots, K$.*

Proof. This is a consequence of the more general theorem 2.

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 to previous work such as Kandori et al [12] where

⁴Replacing 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), $\lim_{\varepsilon \rightarrow 0} \tilde{p}_\varepsilon(s_i, s_i/j)/\varepsilon = c\mu_{ij} = \tilde{\mu}_{ij}$ and so, if $c > 0$ is sufficiently small, $\tilde{\Lambda}_{ii} = 1 - \sum_{j \neq i} \tilde{\mu}_{ij}\rho_{ij} \geq 0$ for $i = 1, \dots, K$.

the ratio of the mutation probabilities has no effect on the limit distribution so long as it is bounded away from 0 and infinity.⁵ 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 [14] 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 [11], where the payoff matrix is $(a_{ij})_{i,j=1}^3$, and each pure strategy is a strict Nash equilibrium.⁶ Suppose the population size is N and $\mu_{ij} = \mu \in (0, \frac{1}{2})$ for all $i \neq j$. Then the absorption probabilities ρ_{ji} are given by

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

⁵Bergin 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 [12].

⁶[11] compute the limit distribution of this process in general 2x2 games, and compute the limit of the limit distribution as the population becomes infinite in these games and in 3x3 coordination games.

[19] compare the absorption probabilities of the no-mutation version of the process to the benchmark of neutral selection; [17] apply that calculation to the two strategies "All D" and "Tit for Tat" in the finitely repeated prisoner's dilemma.

where

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

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 $\boldsymbol{\lambda} = [\gamma_1 + \gamma_2 + \gamma_3]^{-1}(\gamma_1, \gamma_2, \gamma_3)$, 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 $\boldsymbol{\lambda}$, we obtain that for $N = 10$, $\boldsymbol{\lambda} = (.716, .061, .223)$, whereas for $N = 50$, $\boldsymbol{\lambda} = (.082, .0, .918)$. Suppose now that $\mu_{21} = \mu_{23} = .001$ and that the other μ_{ij} remain unchanged. Thus, under a learning model, players of strategy 2 are less likely to make mistakes than 1- or 3-players. Then, for $N = 10$, $\boldsymbol{\lambda} = (.101, .867, .032)$, but for $N = 50$, $\boldsymbol{\lambda} = (.082, .004, .914)$.

Example 3. Consider the local interaction model of [5] and specialize to the case of an odd number of agents (so that the only steady states of the no-mutation process are homogeneous) and that each agent only interacts with his two closest neighbors. Ellison looks at the myopic best reply dynamics, but this is a form of imitation dynamics in coordination games, and in the case under consideration, the homogeneous states are the only recurrent classes of the no-mutation process.⁷ 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. It is true that all of these configurations are in a sense identical, but that is not true for the states where 2 agents play B and the rest play A . However, the example is covered by the more general version of the theorem that we present in the next section.

⁷With an even number of players, there is also a two-cycle.

4 General Imitation Processes

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 $\{X^\varepsilon(t) : t = 0, 1, \dots\}$ 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 $p_0(s, s')$. Denote by s_i , $i = 1, \dots, K$, the states that are absorbing under p_0 ; we assume that all other states in \mathcal{S} are transient under p_0 . (In section 2 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, $\lim_{\varepsilon \rightarrow 0} \pi_\varepsilon(s) = 0$ for every $s \in \mathcal{S} \setminus \{s_1, \dots, s_K\}$.

Assumption 4: For every $i = 1, \dots, K$ and every $s \in \mathcal{S} \setminus \{s_i\}$, $\mu_{is} = \lim_{\varepsilon \rightarrow 0} \frac{p_\varepsilon(s_i, s)}{\varepsilon}$ 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 1 and 2 together imply Assumption 4, 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.

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 $\{s_1, \dots, s_K\}$.

Assumption 5: There is a unique vector $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_K)$ with

$$\boldsymbol{\lambda} \Lambda = \boldsymbol{\lambda}, \quad \lambda_1 + \dots + \lambda_K = 1, \quad \lambda_1 \geq 0, \dots, \lambda_K \geq 0.$$

Note that Assumption 5 reduces to Assumption 3 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 [20] 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.⁸ Our theorem provides a simple means to directly calculate the limit distribution.

Theorem 2 *In the general imitation model, under Assumptions 4 and 5,*

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

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

$$\begin{aligned} \tau_0^\varepsilon &= \min\{u \geq 0 : X^\varepsilon(u) \in \{s_1, \dots, s_K\}\}, \\ \tau_{t+1}^\varepsilon &= \min\{u > \tau_t^\varepsilon : X^\varepsilon(u) \in \{s_1, \dots, s_K\}\}, \quad t \geq 0. \end{aligned}$$

Note that $P\{\tau_t^\varepsilon < \infty\} = 1$ for every $t \geq 0$. For $\varepsilon > 0$, $\{X^\varepsilon(t)\}$ is irreducible, and so $\{Y^\varepsilon(t)\}$ is irreducible as well. Let $\nu_\varepsilon = (\nu_\varepsilon(s_1), \dots, \nu_\varepsilon(s_K))$ denote the ergodic distribution of $\{Y^\varepsilon(t)\}$. According to the strong law of large numbers for Markov chains [see e.g. Theorems 4.6 and 5.1 in Durrett [9], pages 307 and 312], for $i = 1, \dots, K$,

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

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

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

and so

$$\pi_\varepsilon(s_i) \leq \nu_\varepsilon(s_i), \quad i = 1, \dots, K, \quad \varepsilon > 0.$$

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

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

$$\begin{aligned} q_\varepsilon(s_i, s_j) &= P\{X^\varepsilon(\tau_1^\varepsilon) = s_j | X^\varepsilon(0) = s_i\} \\ &= \sum_{s \in \mathcal{S}} p_\varepsilon(s_i, s) P\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\}. \end{aligned} \quad (4)$$

Consider the case $i \neq j$. Then

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

⁸See [13] 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; [15] treats the case of discrete-time and a continuum of states. [6] gives a fairly general sufficient condition for the limit distribution to be concentrated on a single point.

Hence, by (4) and Assumption 4,

$$\begin{aligned}
\lim_{\varepsilon \rightarrow 0} \frac{q_\varepsilon(s_i, s_j)}{\varepsilon} &= \lim_{\varepsilon \rightarrow 0} \sum_{s \in \mathcal{S} \setminus \{s_i\}} \frac{p_\varepsilon(s_i, s)}{\varepsilon} P\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\} \\
&= \lim_{\varepsilon \rightarrow 0} \sum_{s \in \mathcal{S} \setminus \{s_i\}} \mu_{is} P\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\} \\
&= \lim_{\varepsilon \rightarrow 0} \sum_{s \in \mathcal{S}_i} \mu_{is} P\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s\}. \tag{5}
\end{aligned}$$

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

$$\lim_{\varepsilon \rightarrow 0} P\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s^*\} = \lim_{\varepsilon \rightarrow 0} \sum_{t=0}^{\infty} P\{X^\varepsilon(t) = s_j, \tau_0^\varepsilon = t | X^\varepsilon(0) = s^*\}. \tag{6}$$

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

$$\begin{aligned}
P\{X^\varepsilon(t) = s_j, \tau_0^\varepsilon = t | X^\varepsilon(0) = s^*\} &\leq P\{\tau_0^\varepsilon \geq \lfloor t/T \rfloor T | X^\varepsilon(0) = s^*\} \\
&\leq (1 - \delta)^{\lfloor t/T \rfloor} \leq (1 - \delta)^{-1+t/T}
\end{aligned}$$

for every $t = 0, 1, \dots$ and $0 \leq \varepsilon \leq \varepsilon_0$. Thus, by the Weierstrass M-test, the series in (6) 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, page 223. Hence

$$\begin{aligned}
\lim_{\varepsilon \rightarrow 0} P\{X^\varepsilon(\tau_0^\varepsilon) = s_j | X^\varepsilon(0) = s^*\} &= \sum_{t=0}^{\infty} \lim_{\varepsilon \rightarrow 0} P\{X^\varepsilon(t) = s_j, \tau_0^\varepsilon = t | X^\varepsilon(0) = s^*\} \\
&= \sum_{t=0}^{\infty} P\{X^0(t) = s_j, \tau_0^0 = t | X^0(0) = s^*\} \\
&= P\{X^0(\tau_0^0) = s_j | X^0(0) = s^*\} \\
&= \rho_{s^*j}.
\end{aligned}$$

Inserting this into (5) we obtain that

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

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

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

In matrix notation,

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

For every $\varepsilon > 0$, ν_ε is the unique non-negative solution to

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

and, by Assumption 5, λ is the unique non-negative solution to

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

It is now obvious that $\nu_\varepsilon \rightarrow \lambda$. \square

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 $\rho_{s_1} = 1$ for all $s \in S_1$, so $\Lambda_{12} = \sum_{s \in S_1} \mu_{1s} \rho_{s2} = 0$. Moreover, $\rho_{s_1} = 1$ 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 $\rho_{s_1} > 0$ for some $s \in S_2$, while the limit distribution gives positive probability to both absorbing states if $\rho_{s_1} < 1$ for some $s \in S_1$.

Acknowledgments

We thank Glenn Ellison, Michihiro Kandori, and Larry Samuelson for helpful comments.

References

- [1] T. M. Apostol, *Mathematical Analysis*, Addison-Wesley, Reading, 1974.
- [2] J. Bergin and B. Lipman, Evolution with state dependent mutations, *Econometrica* 64 (1996), 943-956.
- [3] J. Bjonerstedt and J. Weibull, Nash equilibrium and evolution by imitation. In: *The Rational Foundations of Economic Analysis*, ed. K. Arrow et al, McMillan, London 1995.

- [4] M. Dufflo, *Algorithmes Stochastiques*, Springer, New York, 1996.
- [5] G. Ellison, Learning, local interaction, and coordination, *Econometrica* 61 (1993), 1047-1071.
- [6] G. Ellison, Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution, *Review of Economic Studies* 67 (2000), 17-45.
- [7] G. Ellison and D. Fudenberg, Rules of thumb for social learning, *Journal of Political Economy* 101 (1993), 612-643.
- [8] G. Ellison and D. Fudenberg, Word of mouth communication and social learning, *Quarterly Journal of Economics* 110 (1995), 93-126.
- [9] R. Durrett, *Probability: Theory and Examples*, Duxbury Press, Belmont, 1996.
- [10] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, Springer, New York, 1984.
- [11] D. Fudenberg, L. Imhof, M. Nowak, and C. Taylor, Stochastic evolution as a generalized Moran process, mimeo, 2004.
- [12] M. Kandori, G. Mailath, and R. Rob, Learning, mutation, and long run equilibria in games, *Econometrica* 61 (1993), 29-56.
- [13] M. Kandori and R. Rob, Evolution of equilibria in the long run, *Journal of Economic Theory* 65 (1995), 383-414.
- [14] M. Kandori and R. Rob, Bandwagon effects and long run technology choice, *Games and Economic Behavior* 22 (1998), 30-60.
- [15] Y. Kifer, A discrete-time version of the Wentzell-Freidlin theory, *Annals of Probability* 18 (1990), 1676-1692.
- [16] G. Noldeke and L. Samuelson, An evolutionary analysis of backward and forward induction, *Games and Economic Behavior* 5 (1993), 425-454.
- [17] M. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg, Emergence of cooperation and evolutionary stability in finite populations, *Nature* 428 (2004), 646-650.
- [18] K. Schlag, Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, *Journal of Economic Theory* 78 (1998), 130-156.
- [19] C. Taylor, D. Fudenberg, M. Nowak, and A. Sasaki, Evolutionary game dynamics in finite populations, *Bulletin of Mathematical Biology*, in press.
- [20] P. Young, The evolution of conventions, *Econometrica* 61 (1993), 57-84.