



Random matching in adaptive dynamics [☆]

Glenn Ellison ^a, Drew Fudenberg ^{b,*}, Lorenz A. Imhof ^c

^a *Massachusetts Institute of Technology, USA*

^b *Harvard University, USA*

^c *Bonn University, Germany*

Received 5 July 2007

Available online 8 May 2008

Abstract

This paper studies the effect of randomness in per-period matching on the long-run outcome of non-equilibrium adaptive processes. If there are many matchings between each strategy revision, the randomness due to matching will be small; our question is when a very small noise due to matching has a negligible effect. We study two different senses of this idea, and provide sufficient conditions for each. The less demanding sense corresponds to sending the matching noise to zero while holding fixed all other aspects of the adaptive process. The second sense in which matching noise can be negligible is that it does not alter the *limit distribution* obtained as the limit of the invariant distributions as an exogenous “mutation rate” goes to zero.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C7; C62

Keywords: Equilibrium selection; Imitation process; Repeated matching; Stochastic matching process; Invariant distribution; Limit distribution; Markov chain

1. Introduction

This paper studies how the relative frequency of strategic interactions and strategy revision influences the long-run outcome of non-equilibrium adaptive processes. Many analyses of these processes abstract away from any randomness in the per-period matching process by assuming either that each agent plays each other agent exactly once in each period, or that agents are independently matched an infinite number of times; in either case, the result is that each agent faces the true distribution of opponents’ play. This is true for example of Kandori et al. (1993), Bergin and Lipman (1996) and Binmore and Samuelson (1997). Neither of these motivations are compelling as stated; a more plausible reason for ignoring the randomness due to the matching process is that the neglected noise is small and thus has little effect. The first claim—that the randomness is small—will follow from the law of large numbers when there are sufficiently many pairings between each strategy revision; our focus is thus on when the second claim then follows, i.e. when a very small noise due to matching has a negligible effect.

[☆] This work was supported by NSF Grants SES-0550897 and SES-0646816.

* Corresponding author at: Harvard University, Cambridge, MA, USA.

E-mail addresses: gellison@mit.edu (G. Ellison), dfudenberg@harvard.edu (D. Fudenberg), limhof@uni-bonn.de (L.A. Imhof).

We study two different senses of this idea. The first, and less demanding sense corresponds to sending the matching noise to zero while holding fixed all other aspects of the adaptive process. The second sense in which matching noise can be negligible is that it does not alter the *limit distribution* obtained as the limit of the invariant distributions as an exogenous “mutation rate” goes to zero. When applied to a model with mutations, the difference between these two senses is in the order of limits: the first sense asks for continuity of e.g. the invariant distribution in the matching noise holding the mutation rate fixed, whereas the second sense asks for continuity of the limit distribution in the matching noise.

Theorem 1 provides a sufficient condition for the first form of continuity, which is roughly that the transition probabilities of the no-noise adjustment process are continuous in the realized per-period payoffs. Under this condition, if the no-noise process is irreducible, its invariant distribution is a continuous function of the noise at the zero matching-noise limit, and for systems with absorbing states the absorption probabilities are continuous. Theorem 2 extends the former conclusion to finite time horizons. To put Theorem 1 into perspective, recall that Robson and Vega-Redondo (1996) show that the limit distribution in Kandori et al.’s (1993) adjustment model selects the Pareto-dominant equilibrium in 2×2 coordination games if players are randomly paired and players observe and respond to the realized distribution of payoffs (even if players are rematched any finite number m times before adjusting their strategies, so that the noise due to matching is small). This contrasts with Kandori et al.’s (1993) finding that the risk-dominant equilibrium is selected when players observe the true state so there is no matching noise at all. For generic payoff matrices, the Robson and Vega-Redondo (1996) adjustment process satisfies the condition of Theorem 1, so the difference between their findings and those of Kandori et al. (1993) comes from taking the no-mutation limit before taking the limit on the number of rounds.

This motivates our investigation of when the second sort of continuity is satisfied: When is the limit distribution a continuous function of the matching noise? Theorem 3 gives a sufficient condition: In addition to the continuity requirement of Theorem 1, it requires essentially that any transition probabilities that are approaching zero as the mutation rate vanishes do so at similar rates when there is no matching noise or a small amount of matching noise. This condition is not satisfied in Robson and Vega-Redondo (1996), but is satisfied in the frequency-dependent Moran process studied by Fudenberg et al. (2006) and also by the frequency-dependent Wright–Fisher process studied by Imhof and Nowak (2006). Theorem 4 provides a less restrictive sufficient condition, by noting that the conditions on the asymptotics of the transition probabilities need only apply at a subset of the transitions. Theorem 5 derives implications of the results for a class of “smooth imitation processes.”

Theorems 6 and 7 turn to the issue of how the frequency of interaction influences the basins of attraction, $\mathcal{D}^{(m)}$ and \mathcal{D}^* , of the processes with m and infinitely many rounds of matching per period. We focus on the case where the processes have the same recurrent classes for every m . Theorem 6 replaces the continuity condition of Theorems 1 and 2 with a mild form of monotonicity that is satisfied by Robson and Vega-Redondo (1996), and shows that $\mathcal{D}^{(m)} = \mathcal{D}^{(1)} \subseteq \mathcal{D}^*$. An example of Robson and Vega-Redondo shows that $\mathcal{D}^{(m)}$ can be a proper subset of \mathcal{D}^* ; this helps further explain the source of the discontinuity in the limit distribution at $m = \infty$. Theorem 7 reimposes the continuity condition to get a sharper result on the relationship of $\mathcal{D}^{(m)}$ and \mathcal{D}^* .

2. The model

Suppose the evolution of a population is described by a homogeneous Markov chain $\{X_t^{(\theta)}: t = 0, 1, \dots\}$ with finite state space \mathcal{S} and transition probabilities $p_{ij}^{(\theta)}$. The parameter θ is an element of some parameter set Θ with accumulation point $\theta^* \notin \Theta$ and relates to the number of rounds that an underlying game is played. We assume that $R_t^{(\theta)}$ rounds are played in period t , where $R_0^{(\theta)}, R_1^{(\theta)}, \dots$ are i.i.d. finite random variables such that

$$\lim_{\theta \rightarrow \theta^*} P(R_t^{(\theta)} \geq M) = 1 \quad \text{for all } M < \infty, \quad (1)$$

and for each t , $R_t^{(\theta)}$ and $X_t^{(\theta)}$ are independent.

Example 1. (a) Condition (1) is satisfied if the number of rounds per period is deterministic, $R_t^{(\theta)} \equiv m$, and $m = \theta \rightarrow \theta^* = \infty$, $\Theta = \mathbb{N}$.

(b) Suppose that in any given period, the game is not played at all with probability $1 - \theta$, and after each round that has taken place there will be another round in the same period with probability θ , where $\theta \in \Theta = (0, 1)$. Then $R_t^{(\theta)}$ is

geometrically distributed with parameter $1 - \theta$, that is, $P(R_t^{(\theta)} = r) = \theta^r (1 - \theta)$, $r = 0, 1, 2, \dots$, and condition (1) is met when $\theta \rightarrow \theta^* = 1$.

For each period t , let $Y_{t,r}$ be a random finite-dimensional vector that describes the payoffs in round r , $r = 1, 2, \dots$. The distribution of $Y_{t,r}$ may depend on $X_t^{(\theta)}$, but not on $R_t^{(\theta)}$. Given $X_t^{(\theta)}$, the vectors $Y_{t,1}, Y_{t,2}, \dots$ are assumed to be i.i.d. with finite expectation. (Note that the support of the $Y_{t,r}$ is not restricted to be finite, but it typically will be finite if the game has only finitely many possible outcomes.) Their common conditional distribution is assumed to be the same for all θ and t . Thus for every state $i \in \mathcal{S}$ there is an integrable random vector Z_i such that

$$P\{Y_{t,r} \in A | X_t^{(\theta)} = i\} = P\{Z_i \in A\} \tag{2}$$

for all r, t and θ , and all Borel-measurable sets A .

Assume that for all $i, j \in \mathcal{S}$, there is a function f_{ij} such that on the set $\{X_t^{(\theta)} = i\} \cap \{R_t^{(\theta)} > 0\}$,

$$P\{X_{t+1}^{(\theta)} = j | R_t^{(\theta)}, Y_{t,1}, Y_{t,2}, \dots\} = f_{ij} \left(\frac{1}{R_t^{(\theta)}} \sum_{r=1}^{R_t^{(\theta)}} Y_{t,r} \right). \tag{3}$$

Note that no assumption is made on the conditional probability on the set $\{R_t^{(\theta)} = 0\}$. Clearly, by (1), $P(R_t^{(\theta)} = 0) \rightarrow 0$ as $\theta \rightarrow \theta^*$. Let $\{X_t^* : t = 0, 1, 2, \dots\}$ be the Markov chain with transition probabilities

$$p_{ij}^* = P\{X_{t+1}^* = j | X_t^* = i\} = f_{ij}(EZ_i), \quad i, j \in \mathcal{S}.$$

This is the Markov chain obtained by assuming that there are infinitely many rounds during each period.

Note that transition probabilities of the form (3) can also occur when matching is deterministic but choices of agents are stochastic. For example, consider a two-player game with two pure strategies A and B and suppose a finite population consists of n types of agents, where agents of type i play A with probability p_i and B with probability $1 - p_i$. Suppose further that in every round, every individual plays exactly once against every other individual. Then realized payoffs are stochastic and it is natural to assume transition probabilities of the form (3).

3. Invariant distributions and absorption probabilities

In this section we show that a continuity condition on the adjustment process implies that long-run and finite-horizon outcomes are continuous as the number of rounds of random matching goes to infinity.

3.1. A theorem on long-run behavior

Theorem 1.

- (a) If f_{ij} is continuous at EZ_i , then $\lim_{\theta \rightarrow \theta^*} p_{ij}^{(\theta)} = p_{ij}^*$.
- (b) Suppose that $\lim_{\theta \rightarrow \theta^*} p_{ij}^{(\theta)} = p_{ij}^*$ for all $i, j \in \mathcal{S}$. If (p_{ij}^*) has a unique invariant distribution v^* and $v^{(\theta)}$ is an invariant distribution of $(p_{ij}^{(\theta)})$, then

$$\lim_{\theta \rightarrow \theta^*} v^{(\theta)} = v^*.$$

- (c) Suppose that $\lim_{\theta \rightarrow \theta^*} p_{ij}^{(\theta)} = p_{ij}^*$ for all $i, j \in \mathcal{S}$. Suppose that the chains $\{X_t^{(\theta)}\}$, $\theta \in \Theta$, and $\{X_t^*\}$ have a common set \mathcal{A} of absorbing states and that the other states are transient states of $\{X_t^*\}$. For $j \in \mathcal{A}$ and $i \in \mathcal{S} \setminus \mathcal{A}$, let $\rho_{ij}^{(\theta)}$ denote the probability that $\{X_t^{(\theta)}\}$ gets absorbed at j if $X_0^{(\theta)} = i$, and let ρ_{ij}^* denote the corresponding absorption probability for $\{X_t^*\}$. Then

$$\lim_{\theta \rightarrow \theta^*} \rho_{ij}^{(\theta)} = \rho_{ij}^*.$$

Proof. (a) By (3), we have on $\{R_t^{(\theta)} > 0\}$,

$$\begin{aligned} P\{X_{t+1}^{(\theta)} = j | X_t^{(\theta)}\} &= E\{1_{\{X_{t+1}^{(\theta)}=j\}} | X_t^{(\theta)}\} \\ &= E\{E[1_{\{X_{t+1}^{(\theta)}=j\}} | X_t^{(\theta)}, R_t^{(\theta)}, Y_{t,1}, Y_{t,2}, \dots] | X_t^{(\theta)}\} \\ &= E\left\{f_{X_t^{(\theta)},j}\left(\frac{1}{R_t^{(\theta)}} \sum_{r=1}^{R_t^{(\theta)}} Y_{t,r}\right) \middle| X_t^{(\theta)}\right\}. \end{aligned}$$

Thus, in view of (2),

$$\begin{aligned} p_{ij}^{(\theta)} &:= P\{X_{t+1}^{(\theta)} = j | X_t^{(\theta)} = i\} \\ &= P(R_0^{(\theta)} = 0)q_{ij} + P(R_0^{(\theta)} > 0)E\left[f_{ij}\left(\frac{1}{R_0^{(\theta)}} \sum_{r=1}^{R_0^{(\theta)}} Z_{i,r}\right) \middle| R_0^{(\theta)} > 0\right], \end{aligned}$$

where $q_{ij} = P\{X_{t+1}^{(\theta)} = j | X_t^{(\theta)} = i, R_t^{(\theta)} = 0\}$ and $Z_{i,1}, Z_{i,2}, \dots$ are i.i.d. copies of Z_i , which are also independent of $R_0^{(\theta)}$. Hence

$$p_{ij}^{(\theta)} = P(R_0^{(\theta)} = 0)q_{ij} + \sum_{k=1}^{\infty} P(R_0^{(\theta)} = k)E f_{ij}\left(\frac{1}{k} \sum_{r=1}^k Z_{i,r}\right). \tag{4}$$

Because the $Z_{i,j}$ are i.i.d. with finite mean, the strong law of large numbers implies that $(Z_{i,1} + \dots + Z_{i,k})/k \rightarrow EZ_i$ almost surely (see e.g. Billingsley, 1995, p. 282). And since f_{ij} is continuous at EZ_i ,

$$\lim_{k \rightarrow \infty} f_{ij}\left(\frac{1}{k} \sum_{r=1}^k Z_{i,r}\right) = f_{ij}(EZ_i) = p_{ij}^* \quad \text{almost surely.}$$

It follows by Lebesgue’s dominated convergence theorem that

$$\lim_{k \rightarrow \infty} E f_{ij}\left(\frac{1}{k} \sum_{r=1}^k Z_{i,r}\right) = p_{ij}^*. \tag{5}$$

Thus, for every $\epsilon > 0$, there exists M_ϵ such that

$$\left|E f_{ij}\left(\frac{1}{k} \sum_{r=1}^k Z_{i,r}\right) - p_{ij}^*\right| \leq \epsilon \quad \text{for all } k > M_\epsilon. \tag{6}$$

Therefore,

$$\begin{aligned} |p_{ij}^{(\theta)} - p_{ij}^*| &\leq P(R_0^{(\theta)} = 0)|q_{ij} - p_{ij}^*| + \sum_{k=1}^{\infty} P(R_0^{(\theta)} = k) \left|E f_{ij}\left(\frac{1}{k} \sum_{r=1}^k Z_{i,r}\right) - p_{ij}^*\right| \\ &\leq \sum_{k=0}^{M_\epsilon} P(R_0^{(\theta)} = k) + \epsilon \sum_{k=M_\epsilon+1}^{\infty} P(R_0^{(\theta)} = k) \\ &\leq P(R_0^{(\theta)} \leq M_\epsilon) + \epsilon. \end{aligned}$$

The first inequality follows from (4) and the triangle inequality. For the second inequality we used (6) and that $|q_{ij} - p_{ij}^*| \leq 1$ and $|E f_{ij}((1/k) \sum_{r=1}^k Z_{i,r}) - p_{ij}^*| \leq 1$. It now follows by assumption (1) and the fact that $\epsilon > 0$ was arbitrary that

$$\lim_{\theta \rightarrow \theta^*} p_{ij}^{(\theta)} = p_{ij}^*. \tag{7}$$

(b) Write $\Pi^{(\theta)} = (p_{ij}^{(\theta)})_{i,j \in \mathcal{S}}$ and $\Pi^* = (p_{ij}^*)_{i,j \in \mathcal{S}}$. If the invariant distributions $v^{(\theta)}$ do not converge to v^* as $\theta \rightarrow \theta^*$, then, because the probability simplex is compact, the sequence has an accumulation point $w \neq v^*$. Because $\Pi^{(\theta)}$ converges to Π^* as $\theta \rightarrow \theta^*$ and $v^{(\theta)}\Pi^{(\theta)} = v^{(\theta)}$ for all θ , it follows that $w\Pi^* = w$, which would contradict the assumption that v^* is the unique invariant probability vector for Π^* .

(c) Fix a common absorbing state j . Since all non-absorbing states are transient states of $\{X_t^*\}$, it follows from (7) that all non-absorbing states are also transient states of $\{X_t^{(\theta)}\}$ when θ is close enough to θ^* (that is, when θ is large enough in the case where $\theta^* = \infty$). Therefore, the absorption probabilities are uniquely determined by

$$\rho_{ij}^{(\theta)} = p_{ij}^{(\theta)} + \sum_{s \in \mathcal{T}} p_{is}^{(\theta)} \rho_{sj}^{(\theta)}, \quad \rho_{ij}^* = p_{ij}^* + \sum_{s \in \mathcal{T}} p_{is}^* \rho_{sj}^*, \quad i \in \mathcal{T},$$

where $\mathcal{T} = \mathcal{S} \setminus \mathcal{A}$, see e.g. Karlin and Taylor (1975, p. 90). It now follows as in the proof of (b) that $\rho_{ij}^{(\theta)} \rightarrow \rho_{ij}^*$. \square

Remark 1. (a) Theorem 1(a) continues to hold if one replaces the independence assumption on $Y_{t,1}, Y_{t,2}, \dots$ by the assumption that these vectors satisfy the weak law of large numbers. In this case, $f_{ij}((Z_{i,1} + \dots + Z_{i,k})/k) \rightarrow f_{ij}(EZ_i)$ in distribution and this implies (5).

(b) The perturbation bounds reviewed by Cho and Meyer (2001) show that the invariant distributions $v^{(\theta)}$ converge at least as quickly as the transition matrices $\Pi^{(\theta)}$. These bounds also imply our result in Theorem 1(b), but the much simpler self-contained argument in the present proof is sufficient for our purposes.

(c) We use Theorem 1 in one of the proofs of Theorem 5.

The continuity assumption of Theorem 1(a) is obviously satisfied when the f_{ij} are continuous everywhere. Even when this is not the case, Theorem 1(a) can be used to obtain results for adjustment dynamics in generic games if the f_{ij} are continuous except on lower-dimensional sets of observed payoffs, as is the case under any dynamics that are “Darwinian” in the sense of Kandori et al. (1993).

To see this, consider a symmetric 2-player game with pure strategies $1, \dots, n$ and payoff matrix A , and let $Y_{t,r}$ be a vector giving the average realized payoff of each strategy in the r th matching. For any population state i , the expected payoff $EZ_i^{(k)}$ to strategy k can be expressed as $a_k s_{ki}$, where a_k is the k th row of A and $s_{ki} \in \mathbb{R}^n \setminus \{0\}$ is determined by the distribution of the random matchings.¹ If D_{ij} is the set of points of discontinuity of f_{ij} and D_{ij} has Lebesgue measure 0, then the set of payoff matrices

$$\mathcal{A}_{ij} = \{A \in \mathbb{R}^{n \times n}: [a_1 s_{1i}, \dots, a_n s_{ni}]^T \in D_{ij}\}$$

has Lebesgue measure 0 as well. Thus for all $A \in \cap_{i,j} \mathcal{A}_{ij}^C$, the assumption of Theorem 1(a) is satisfied. Thus, by Theorem 1(a) and (b) we obtain the following result.

Corollary 1. *Suppose the Z_i are realized payoffs from random matchings in a symmetric two-player game and the f_{ij} are continuous almost everywhere. Then, for generic payoff matrices, the invariant distributions v satisfy*

$$\lim_{\theta \rightarrow \theta^*} v^{(\theta)} = v^*, \tag{8}$$

provided the invariant distribution is unique when there are infinitely many rounds in each period.

We show in Example 3 that (8) may fail for a three-dimensional set of 2×2 payoff matrices.

3.2. Illustrative examples

The following examples show that some form of continuity assumption is necessary for the conclusion that matching noise becomes irrelevant in the limit of infinitely many observations. We begin with a simple but artificial example.

¹ If strategy k does not occur in state i , $Z_i^{(k)}$ may be undefined. In this case, we set $Z_i^{(k)} = a_k s_{ki}$ for some fixed $s_{ki} \neq 0$.

Example 2. Suppose there are two actions 0 and 1, and in every round every player gets payoff $\sqrt{2}$ with probability $1/\sqrt{2}$ and 0 with probability $1 - 1/\sqrt{2}$, independent of his action and the state of the population. Suppose the dynamics are that players switch from 0 to 1 if and only if they get a payoff of 1 and never switch from 1 to 0.

In the limit model with an infinite number of rounds of matching, all players play action 1 after the first period. Hence, the unique invariant distribution, v^* places probability one on all agents' playing action 1. For any finite number m of rounds, in contrast, there cannot be a switch because $\sqrt{2}$ is irrational. Hence, any distribution over the two actions is an invariant distribution of $(p_{ij}^{(m)})$.

Next, we present a more natural example: transition rules are discontinuous because agents adopt the strategy that provided the highest payoff in previous period. Again, there is a discontinuity in long-run behavior.

Example 3. Consider a symmetric 2-player game with pure strategies 1, 2 and positive payoff matrix

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Assume that

$$2b = c + d, \quad a + b > 2c. \tag{9}$$

Consider a population of size 3 and suppose that in every period there are exactly m rounds. Let $X_t^{(m)}$ denote the number of agents using strategy 1 at time t . At each round, two individuals are drawn at random to play the game, the remaining individual receives no payoff in that round. For every period t and round r , let $Y_{t,r}^{(i)}$ denote the average payoff to agents using strategy i . Set $Y_{t,r}^{(i)} = 0$ if strategy i is not present at time t . If $X_t^{(m)} = 1$, the probability that the two agents that use strategy 2 are chosen is $\frac{1}{3}$, and the probability for a mixed pair is $\frac{2}{3}$. Hence

$$P\{Y_{t,r}^{(1)} = 0, Y_{t,r}^{(2)} = d | X_t^{(m)} = 1\} = \frac{1}{3}, \quad P\left\{Y_{t,r}^{(1)} = b, Y_{t,r}^{(2)} = \frac{c}{2} | X_t^{(m)} = 1\right\} = \frac{2}{3}. \tag{10}$$

Similarly,

$$P\{Y_{t,r}^{(1)} = a, Y_{t,r}^{(2)} = 0 | X_t^{(m)} = 2\} = \frac{1}{3}, \quad P\left\{Y_{t,r}^{(1)} = \frac{b}{2}, Y_{t,r}^{(2)} = c | X_t^{(m)} = 2\right\} = \frac{2}{3}.$$

If after m rounds, the average payoff of agents using strategy 1 is larger than the average payoff of the other agents, all agents switch to strategy 1. Otherwise, they all switch to strategy 2. (Agents do not switch if only one strategy is currently present.) After this adjustment step, every agent changes (independently of the others) to the other strategy with probability ϵ , $0 < \epsilon < \frac{1}{2}$. Thus for $i, j = 0, \dots, 3$,

$$f_{ij}(y^{(1)}, y^{(2)}) = \begin{cases} \binom{3}{j} (1 - \epsilon)^j \epsilon^{3-j}, & \text{if } y^{(1)} > y^{(2)}, \\ \binom{3}{j} \epsilon^j (1 - \epsilon)^{3-j}, & \text{otherwise.} \end{cases}$$

For any number of rounds m ,

$$p_{0j}^{(m)} = p_{0j}^* = \binom{3}{j} \epsilon^j (1 - \epsilon)^{3-j}, \quad p_{3j}^{(m)} = p_{3j}^* = \binom{3}{j} (1 - \epsilon)^j \epsilon^{3-j}.$$

In view of (10), $EZ_1 = (\frac{2}{3}b, \frac{1}{3}(c + d))$. By (9), the two components of EZ_1 coincide, and so

$$p_{1j}^* = f_{1j}(EZ_1) = \binom{3}{j} \epsilon^j (1 - \epsilon)^{3-j}.$$

Let $(Z_{11}^{(1)}, Z_{11}^{(2)})$, $(Z_{12}^{(1)}, Z_{12}^{(2)})$, ... be i.i.d. copies of Z_1 . The random variables $Z_{1r}^{(1)} - Z_{1r}^{(2)}$, $r = 1, 2, \dots$, have mean zero and (since $d > 0$), positive and finite variance. Therefore, by the central limit theorem,

$$P \left\{ \frac{Z_{11}^{(1)} + \dots + Z_{1m}^{(1)}}{m} > \frac{Z_{11}^{(2)} + \dots + Z_{1m}^{(2)}}{m} \right\} \\ = P \left\{ \frac{1}{\sqrt{m}} \sum_{r=1}^m (Z_{1r}^{(1)} - Z_{1r}^{(2)}) > 0 \right\} \rightarrow \frac{1}{2}$$

as $m \rightarrow \infty$.² It follows that

$$\lim_{m \rightarrow \infty} p_{1j}^{(m)} = \binom{3}{j} \frac{1}{2} [(1 - \epsilon)^j \epsilon^{3-j} + \epsilon^j (1 - \epsilon)^{3-j}];$$

in particular, $\lim_{m \rightarrow \infty} p_{1j}^{(m)} \neq p_{1j}^*$.

By (9), $E Z_2^{(1)} = \frac{1}{3}(a + b) > \frac{2}{3}c = E Z_2^{(2)}$, and it follows that

$$\lim_{m \rightarrow \infty} p_{2j}^{(m)} = p_{2j}^* = \binom{3}{j} (1 - \epsilon)^j \epsilon^{3-j}.$$

It can now be verified that for fixed $\epsilon > 0$, the invariant distribution $v^{(m,\epsilon)}$ of $\{X_t^{(m)}\}$ does not converge to the invariant distribution $v^{(\infty,\epsilon)}$ of $\{X_t^*\}$. In fact,

$$\lim_{m \rightarrow \infty} v^{(m,\epsilon)} = [O(\epsilon), O(\epsilon^2), O(\epsilon), 1 + O(\epsilon)]$$

and

$$v^{(\infty,\epsilon)} = \left[\frac{1}{2} + O(\epsilon), O(\epsilon), O(\epsilon), \frac{1}{2} + O(\epsilon) \right].$$

Note that the difference between these two distributions does not vanish as $\epsilon \rightarrow 0$:

$$\lim_{\epsilon \rightarrow 0} \lim_{m \rightarrow \infty} v^{(m,\epsilon)} = [0, 0, 0, 1] \neq \lim_{\epsilon \rightarrow 0} v^{(\infty,\epsilon)} = \left[\frac{1}{2}, 0, 0, \frac{1}{2} \right].$$

These limit results hold in particular for the following payoff matrices, which satisfy (9):

$$\begin{pmatrix} 3 & 4 \\ 2 & 6 \end{pmatrix}, \quad \begin{pmatrix} 5 & 4 \\ 2 & 6 \end{pmatrix}, \quad \begin{pmatrix} 7 & 4 \\ 2 & 6 \end{pmatrix}, \quad \begin{pmatrix} 1 & 4 \\ 2 & 6 \end{pmatrix}, \quad \begin{pmatrix} 5 & 2 \\ 3 & 1 \end{pmatrix}.$$

In the first case, strategy 2 is risk- and Pareto-dominant, in the second case, strategy 1 is risk-dominant and strategy 2 is Pareto-dominant, and in the third case, strategy 1 is risk- and Pareto-dominant. (Condition (9) cannot hold if strategy 1 is Pareto-dominant and strategy 2 is risk-dominant.) In the final two cases, one strategy is strictly dominant.

The following example illustrates the effect of smoothing the transition functions.

Example 4. Modify the model of the previous example so that agents switch to strategy 1 if the difference of the average payoff of agents using strategy 1 and the average payoff of the other agents exceeds a random threshold. Otherwise, all agents switch to strategy 2. Then mutations may occur as described above. Suppose the threshold is normally distributed with expectation zero and variance $\sigma^2 > 0$. Denote the corresponding distribution function by F_{σ^2} . Then the function f_{ij} in (3) is given by the continuous function

$$f_{ij}^{(\sigma^2)}(y^{(1)}, y^{(2)}) = F_{\sigma^2}(y^{(1)} - y^{(2)}) \binom{3}{j} (1 - \epsilon)^j \epsilon^{3-j} \\ + F_{\sigma^2}(y^{(2)} - y^{(1)}) \binom{3}{j} \epsilon^j (1 - \epsilon)^{3-j}.$$

² See Billingsley (1995, p. 357).

Clearly, if $y^{(1)} \neq y^{(2)}$, then $f_{ij}^{(\sigma^2)}(y^{(1)}, y^{(2)}) \rightarrow f_{ij}(y^{(1)}, y^{(2)})$ as $\sigma^2 \rightarrow 0$. Again, for every fixed $\sigma^2 > 0$, $p_{0j}^{(m)} = p_{0j}^*$, $p_{3j}^{(m)} = p_{3j}^*$ and $p_{2j}^{(m)} \rightarrow p_{2j}^*$. But now,

$$p_{1j}^* = f_{1j}^{(\sigma^2)}(EZ_1) = \binom{3}{j} \frac{1}{2} [(1 - \epsilon)^j \epsilon^{3-j} + \epsilon^j (1 - \epsilon)^{3-j}],$$

so that $p_{1j}^{(m)} \rightarrow p_{1j}^*$. Thus, with smoothing, the invariant distribution of $(p_{ij}^{(m)})$ converges to that of (p_{ij}^*) .

3.3. A theorem on finite-horizon behavior

Theorem 1 implies in the case of irreducibility that the long-run behavior of $\{X_t^{(\theta)}\}$ is similar to that of $\{X_t^*\}$, provided θ is close enough to θ^* . To investigate the similarity for any finite time-horizon, suppose that $\{X_t^{(\theta)}\}$ and $\{X_t^*\}$ have the same initial distribution q . Then the distribution of $X_t^{(\theta)}$ and X_t^* are given by $p^{(\theta)}(t) := q(\Pi^{(\theta)})^t$ and $p^*(t) := q(\Pi^*)^t$, respectively. Thus under the continuity condition of Theorem 1, it follows from (7) that $p^{(\theta)}(t) \rightarrow p^*(t)$ for every t . The following theorem shows that the convergence is uniform in t if $\{X_t^*\}$ is irreducible and aperiodic.

For probability vectors $p = (p_i)_{i \in \mathcal{S}}$, $q = (q_i)_{i \in \mathcal{S}}$ let $\|p - q\| = \frac{1}{2} \sum_{i \in \mathcal{S}} |p_i - q_i|$ be their variation distance.

Theorem 2. *Suppose that $\{X_t^*\}$ is irreducible and aperiodic and that for every $i, j \in \mathcal{S}$, f_{ij} is continuous at EZ_i . Then, for every common initial distribution for the chains $\{X_t^{(\theta)}\}$ and $\{X_t^*\}$,*

$$\lim_{\theta \rightarrow \theta^*} \sup_{t=0,1,2,\dots} \|p^{(\theta)}(t) - p^*(t)\| = 0.$$

Proof. It follows from (7) that there is a neighborhood U_0 of θ^* such that $\{X_t^{(\theta)}\}$ is also irreducible and aperiodic for all $\theta \in U_0 \cap \Theta$. Let $\epsilon > 0$. By Theorem 1(a) and (b), there exists a neighborhood U_1 of θ^* with $U_1 \subset U_0$ such that

$$\|v^{(\theta)} - v^*\| < \frac{\epsilon}{2} \quad \text{for all } \theta \in U_1 \cap \Theta. \tag{11}$$

Moreover, there exist constants $\beta(\theta)$, $\beta^* < 1$ such that for $t = 0, 1, 2, \dots$,

$$\|p^{(\theta)}(t) - v^{(\theta)}\| \leq 2[\beta(\theta)]^t, \quad \theta \in U_0 \cap \Theta, \quad \|p^*(t) - v^*\| \leq 2(\beta^*)^t, \tag{12}$$

and $\beta(\theta) \rightarrow \beta^*$, see e.g. Iosifescu (1980, p. 126). Consequently, there exists $T < \infty$ and another neighborhood U_2 of θ^* with $U_2 \subset U_1$ such that for all $\theta \in U_2 \cap \Theta$ and all $t > T$,

$$\|p^{(\theta)}(t) - p^*(t)\| \leq \|p^{(\theta)}(t) - v^{(\theta)}\| + \|v^{(\theta)} - v^*\| + \|v^* - p^*(t)\| < \epsilon.$$

The first inequality follows from the triangle inequality, the second one from (11) and the fact that by (12), $\|p^{(\theta)}(t) - v^{(\theta)}\| < \frac{\epsilon}{4}$ and $\|v^* - p^*(t)\| < \frac{\epsilon}{4}$ for t large enough. To complete the proof note that by (7) there is a neighborhood U_3 of θ^* such that if $\theta \in U_3 \cap \Theta$, $\|p^{(\theta)}(t) - p^*(t)\| < \epsilon$ for $t = 0, \dots, T$. Altogether, we have shown that $\sup_{t=0,1,2,\dots} \|p^{(\theta)}(t) - p^*(t)\| \leq \epsilon$ for all $\theta \in U_2 \cap U_3 \cap \Theta$. \square

Remark 2. Note that in Example 3, $\{X_t^*\}$ is irreducible and aperiodic, but $p_{1j}^{(m)} \not\rightarrow p_{1j}^*$, and so $p^{(m)}(1)$ will not approach $p^*(1)$ if the common initial distribution is given by $P(X_0^{(m)} = 1) = P(X_0^* = 1) = 1$.

4. Limit distributions

In this section we consider evolutionary processes with mutations where the updating depends on the outcomes of a repeated game as in Section 2. The mutations ensure that the processes have unique invariant distributions. In such models it is standard to describe long-run behavior with infrequent mutations by characterizing the *limit distribution* obtained as the limit of the invariant distribution as the mutation rate goes to zero.

In these models analysts typically think that both mutation probabilities and matching noise are very small; which we can idealize by sending them both to zero. The paper of Robson and Vega-Redondo (1996), combined with our results, shows that the order of limits can matter. Robson and Vega-Redondo analyze a model in which players are

repeatedly randomly matched to play a symmetric 2×2 coordination game in which strategy A is Pareto optimal and strategy B is risk-dominant. They consider an emulation-learning rule with ϵ mutations: in each period there are m rounds of random matching and each player adopts the strategy that achieved the highest payoff with independent probability $1 - \epsilon$.

With $m = \infty$ the Robson and Vega-Redondo (1996) model coincides with the Kandori et al. (1993) model, and hence the “long-run equilibrium” (the limit as $\epsilon \rightarrow 0$ of $v^{(*, \epsilon)}$) places probability one on all players playing the risk-dominant equilibrium. For any finite m , in contrast, Robson and Vega-Redondo show that the long-run equilibrium places probability one on the Pareto optimal equilibrium.

The transition rule in the Robson and Vega-Redondo model need not satisfy our continuity condition, but it does satisfy it for generic payoff functions.³ Hence, Theorem 1 implies that for any fixed mutation rate ϵ we have $v^{(m, \epsilon)} \rightarrow v^{(*, \epsilon)}$, i.e. long-run behavior when m is sufficiently large (with “sufficient” depending on ϵ) is similar to long-run behavior with m infinite. This shows that the difference between the long-run equilibrium of the Kandori et al. (1993) model and the long-run equilibrium of the Robson and Vega-Redondo (1996) model is not due to a discontinuity at $m = \infty$ as in Example 1. Instead, it is due to a difference in the order of limits, as Robson and Vega-Redondo say: $\lim_{m \rightarrow \infty} \lim_{\epsilon \rightarrow 0} v^{(m, \epsilon)} \neq \lim_{\epsilon \rightarrow 0} \lim_{m \rightarrow \infty} v^{(m, \epsilon)}$. Which order of limits is more relevant will depend on the magnitudes of ϵ and m . Intuitively, the difference between the Robson and Vega-Redondo model and the Kandori et al. model is that in the Robson and Vega-Redondo model a transition from one equilibrium to the other can be facilitated by an unrepresentative draw of the random matching. The probability of such draws goes to zero as m gets large, so the unrepresentative matching draws is the main source of transitions when ϵ is small compared to $1/m$.

Remark 3. Theorem 20 of Vega-Redondo (1996, p. 140) asserts that for a Kandori et al. (1993)-style dynamic with m rounds of matching in a 2×2 coordination game, where one equilibrium is risk dominant and the other is Pareto-dominant, $\lim_{\epsilon \rightarrow 0} \lim_{m \rightarrow \infty} v^{(m, \epsilon)}$ is concentrated on the risk-dominant equilibrium for all population sizes N greater than some \bar{N} . The theorem does not impose a continuity condition; the suggested proof assumes that for any fixed population size, $\lim_{m \rightarrow \infty} v^{(m, \epsilon)} = v^{(\infty, \epsilon)}$, and thus that $\lim_{\epsilon \rightarrow 0} \lim_{m \rightarrow \infty} v^{(m, \epsilon)} = \lim_{\epsilon \rightarrow 0} v^{(\infty, \epsilon)}$, which does not follow from the assumptions of the theorem. However, Corollary 1 shows that the suggested proof is in fact correct for generic payoff matrices. Moreover, we conjecture that the theorem is true as stated, because of its hypothesis of a “sufficiently large” population.

4.1. Continuity theorems

In this subsection we prove a general theorem showing that limit distributions will be continuous as $\theta \rightarrow \theta^*$ if the transition probabilities $p_{ij}^{\theta, \epsilon}$ are continuous and in addition any transition probabilities that are approaching zero as $\epsilon \rightarrow 0$ in the θ^* model do so at similar rates when θ is close to θ^* . One application of this result is that the limit distributions of imitation processes like that of Robson and Vega-Redondo (1996) will be continuous as the number of rounds of play goes to infinity if players react smoothly to payoff differences.

Theorem 3. Consider a family of Markov transition matrices $(p_{jk}^{\theta, \epsilon})$ on a common state space \mathcal{S} . For each $\theta \in \Theta \cup \{\theta^*\}$, suppose that $(p_{jk}^{\theta, \epsilon})$ is irreducible with invariant distribution $v^{(\theta, \epsilon)}$ when $\epsilon > 0$. Suppose that the limit distribution of the θ^* -process, $\lambda^{\theta^*} \equiv \lim_{\epsilon \rightarrow 0} v^{(\theta^*, \epsilon)}$ is well-defined. Suppose also that the transition probabilities satisfy

$$\lim_{\theta \rightarrow \theta^*} \lim_{\epsilon \rightarrow 0} \frac{p_{jk}^{\theta, \epsilon}}{p_{jk}^{\theta^*, \epsilon}} = 1 \tag{13}$$

for all $j, k \in \mathcal{S}$, where all 0/0 fractions are taken to be equal to 1. Then, for θ sufficiently close to θ^* the limit distributions $\lambda^\theta = \lim_{\epsilon \rightarrow 0} v^{(\theta, \epsilon)}$ are also well-defined and satisfy

$$\lim_{\theta \rightarrow \theta^*} \lambda^\theta = \lambda^{\theta^*}.$$

³ We only require continuity when payoffs are equal to their expectation under an infinite number of matchings. Generically, a player will not be indifferent when facing a mixed strategy of the form $Ak/N + B(N - k)/N$.

In particular, we have

$$\lim_{\theta \rightarrow \theta^*} \lim_{\epsilon \rightarrow 0} v^{(\theta, \epsilon)} = \lim_{\epsilon \rightarrow 0} \lim_{\theta \rightarrow \theta^*} v^{(\theta, \epsilon)}.$$

Proof. It is a standard result in this literature (e.g. Freidlin and Wentzell, 1984) that the invariant distribution $v^{(\theta, \epsilon)}$ satisfies

$$v_i^{(\theta, \epsilon)} = \frac{\sum_{z \in Z_i} (\prod_{(j,k) \in z} P_{j,k}^{\theta, \epsilon})}{\sum_{z \in Z} (\prod_{(j,k) \in z} P_{j,k}^{\theta, \epsilon})},$$

where Z_i is the set of all i -trees on \mathcal{S} and Z is the union of the Z_s over all $s \in \mathcal{S}$. (An i -tree is a set of ordered pairs of states describing a directed graph in which there is a unique directed edge out of every element other than i and in which there is a path from every state other than i to i .) Hence,

$$\frac{v_i^{(\theta, \epsilon)}}{v_i^{(\theta^*, \epsilon)}} = \frac{\sum_{z \in Z_i} (\prod_{(j,k) \in z} P_{j,k}^{\theta, \epsilon})}{\sum_{z \in Z_i} (\prod_{(j,k) \in z} P_{j,k}^{\theta^*, \epsilon})} \cdot \frac{\sum_{z \in Z} (\prod_{(j,k) \in z} P_{j,k}^{\theta^*, \epsilon})}{\sum_{z \in Z} (\prod_{(j,k) \in z} P_{j,k}^{\theta, \epsilon})}.$$

Each of the fractions on the right side of the above expression is of the form $\sum_{y \in Y} a_y / \sum_{y \in Y} b_y$. Such a ratio is bounded below by $\min_{y \in Y} a_y / b_y$ and above by $\max_{y \in Y} a_y / b_y$. Every i -tree on \mathcal{S} contains $|\mathcal{S}| - 1$ directed edges. Hence, the first of the two fractions is bounded below by $(\min_{j,k \in \mathcal{S}} P_{j,k}^{\theta, \epsilon} / P_{j,k}^{\theta^*, \epsilon})^{|\mathcal{S}| - 1}$ and bounded above by the same expression with max in place of min. The second fraction has similar bounds. Our assumption on the convergence of the transition probabilities thereby implies that $\lim_{\theta \rightarrow \theta^*} \lim_{\epsilon \rightarrow 0} v_i^{(\theta, \epsilon)} / v_i^{(\theta^*, \epsilon)} = 1$. This gives $\lambda_i^\theta \rightarrow \lambda_i^{\theta^*}$ as $\theta \rightarrow \theta^*$ for all i . The final conclusion of the theorem about the order of limits follows from Theorem 1(b): The hypothesis about the limit of the ratios implies that the transition probabilities are continuous as $\theta \rightarrow \theta^*$. Hence Theorem 1(b) implies $v^{(\theta^*, \epsilon)} = \lim_{\theta \rightarrow \theta^*} v^{(\theta, \epsilon)}$, so that $\lambda^{\theta^*} = \lim_{\epsilon \rightarrow 0} v^{(\theta^*, \epsilon)} = \lim_{\epsilon \rightarrow 0} \lim_{\theta \rightarrow \theta^*} v^{(\theta, \epsilon)}$. \square

Remark 4. The Robson and Vega-Redondo (1996) model does not satisfy the assumption of Theorem 3. Consider the state i where two of the agents play the action of the Pareto-optimal equilibrium, while $N - 2$ of them play the other action, where $N \geq 4$. When $m = \infty$, a transition to state j where all agents play the Pareto-optimal action has probability ϵ^N . For finite m , $p_{ij}^{m, \epsilon} = g(m)(1 - \epsilon)^N + (1 - g(m))\epsilon^N$, where $g(m)$ is the probability that the random matching is such that the Pareto-optimal action has the higher realized payoff. $g(m)$ converges to zero as $m \rightarrow \infty$, so the $p_{ij}^{m, \epsilon}$ is continuous in the $m \rightarrow \infty$ limit. The ratio of the two transition probabilities, however, is $p_{ij}^{m, \epsilon} / p_{ij}^{\infty, \epsilon} = 1 - g(m) + g(m)(1 - \epsilon)^N / \epsilon^N$. This diverges as $\epsilon \rightarrow 0$ for every value of m , so the hypothesis of Theorem 3 is not satisfied.

In models of evolution with mutations, some transitions are typically assumed to be much less likely than others when ϵ is small, e.g. a transition requiring two simultaneous mutations may have probability ϵ^2 , whereas a transition requiring a single mutation has probability ϵ . Young (1993) and others have noted that the calculation of the limit distribution can be simplified by ignoring some relatively unlikely transitions and focusing on what are known as “minimum cost trees.” Theorem 3 can similarly be strengthened by noting that the condition on the ratios of the transition probabilities need only hold for a subset of the transitions.

We will say that an i -tree z_i is negligible (for small ϵ) if

$$\frac{\prod_{(j,k) \in z_i} P_{j,k}^{\theta, \epsilon}}{\max_{z \in Z_i} (\prod_{(j,k) \in z} P_{j,k}^{\theta, \epsilon})} \rightarrow 0 \quad \text{as } \epsilon \rightarrow 0$$

for all $\theta \in \Theta \cup \{\theta^*\}$.

Theorem 4. For every $\theta \in \Theta \cup \{\theta^*\}$ and $\epsilon > 0$ let $(p_{jk}^{\theta, \epsilon})_{j,k \in \mathcal{S}}$ be irreducible with invariant distribution $v^{(\theta, \epsilon)}$. Suppose that $\lim_{\theta \rightarrow \theta^*} p_{jk}^{\theta, \epsilon} = p_{jk}^{\theta^*, \epsilon}$ for all $\epsilon > 0$ and $j, k \in \mathcal{S}$ and that $\lambda^{\theta^*} = \lim_{\epsilon \rightarrow 0} v^{(\theta^*, \epsilon)}$ exists. Let

$$\mathcal{I} = \{i \in \mathcal{S} : \lim_{\epsilon \rightarrow 0} v_i^{(\theta, \epsilon)} = 0 \text{ for all } \theta \in \Theta \cup \{\theta^*\}\}.$$

Suppose that if $i \in \mathcal{S} \setminus \mathcal{I}$ and z is a nonnegligible i -tree, then condition (13) is satisfied for all $(j, k) \in z$. Then the conclusions of Theorem 3 hold.

Proof. The argument for this extension is straightforward. It suffices to show that $\lambda_i^\theta \rightarrow \lambda_i^{\theta^*}$ for $i \in \mathcal{S} \setminus \mathcal{I}$. In the proof of Theorem 3 we noted that

$$\frac{v_i^{(\theta, \epsilon)}}{v_i^{(\theta^*, \epsilon)}} = \frac{\sum_{z \in Z_i} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})}{\sum_{z \in Z_i} (\prod_{(j,k) \in z} p_{j,k}^{\theta^*, \epsilon})} \cdot \frac{\sum_{z \in Z} (\prod_{(j,k) \in z} p_{j,k}^{\theta^*, \epsilon})}{\sum_{z \in Z} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})}.$$

Let W_i denote the set of nonnegligible i -trees and let $W = \cup_{i \in \mathcal{S} \setminus \mathcal{I}} W_i$. Then for all $\theta \in \Theta \cup \{\theta^*\}$,

$$\frac{\sum_{z \in W_i} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})}{\sum_{z \in Z_i} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})} \rightarrow 1, \quad \frac{\sum_{z \in W} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})}{\sum_{z \in Z} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})} \rightarrow 1 \text{ as } \epsilon \rightarrow 0.$$

It now follows that for $i \in \mathcal{S} \setminus \mathcal{I}$,

$$\lim_{\theta \rightarrow \theta^*} \lim_{\epsilon \rightarrow 0} \frac{v_i^{(\theta, \epsilon)}}{v_i^{(\theta^*, \epsilon)}} = \lim_{\theta \rightarrow \theta^*} \lim_{\epsilon \rightarrow 0} \frac{\sum_{z \in W_i} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})}{\sum_{z \in W_i} (\prod_{(j,k) \in z} p_{j,k}^{\theta^*, \epsilon})} \cdot \frac{\sum_{z \in W} (\prod_{(j,k) \in z} p_{j,k}^{\theta^*, \epsilon})}{\sum_{z \in W} (\prod_{(j,k) \in z} p_{j,k}^{\theta, \epsilon})} = 1.$$

Thus, as $\lim_{\epsilon \rightarrow 0} v_i^{(\theta^*, \epsilon)} = \lambda_i^{\theta^*}$, $\lambda_i^\theta = \lim_{\epsilon \rightarrow 0} v_i^{(\theta, \epsilon)}$ exists for θ sufficiently close to θ^* , and $\lim_{\theta \rightarrow \theta^*} \lambda_i^\theta = \lambda_i^{\theta^*}$. \square

4.2. Smooth imitation processes

We now discuss a class of smoothed imitation processes as an illustration of when the order of limits does not matter. Consider a population of size N and a game with pure strategies $1, \dots, K$. Let $\mathcal{S} = \{(x_1, \dots, x_K) : x_i \in \{0, \dots, N\}, x_1 + \dots + x_K = N\}$, where x_i is the number of agents that use strategy i . Assume that for each fixed $\epsilon \geq 0$, the family $(\{X_t^{(\theta, \epsilon)} : \theta \in \Theta\})$ satisfies the assumptions of Section 2 with certain functions $f_{ij}^{(\epsilon)}$ and that $\{X_t^{(\theta^*, \epsilon)}\}$ is the corresponding Markov chain for infinitely many rounds. Here ϵ corresponds to the mutation rate, as specified below, and $\{X_t^{(\theta, 0)}\}$ and $\{X_t^{(\theta^*, 0)}\}$ are the no-mutation processes.

Following Fudenberg and Imhof (2006), we say that the no-mutation process is an *imitation process* if (a) $p_{jk}^{\theta, 0} = 0$ whenever there exists a strategy i that is present in state k but not in state j and (b) every state where two or more actions are played is transient. These conditions mean that absent strategies will never be (re-)introduced without mutations and that, starting from any initial state, the process will be absorbed in a state where all agents use the same strategy.

Mutations are typically introduced into processes like these to capture the idea that extinct strategies can arise via mutations. For $i = 1, \dots, K$, let s_i denote the homogeneous state in which every agent plays i , and for $i, j = 1, \dots, K$ with $i \neq j$, let $s_{i/j}$ denote the state in which every agent plays i except for one, who plays j . We will say that $\{X_t^{(\theta, \epsilon)}\}$ is a *standard mutation extension* of $\{X_t^{(\theta, 0)}\}$ if the transition probabilities are continuous in ϵ and for $s \neq s_i$ we have

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} P\{X_{t+1}^{(\theta, \epsilon)} = s \mid X_t^{(\theta, \epsilon)} = s_i\} = \begin{cases} \mu_{ij}^{(\theta)} > 0 & \text{for } s = s_{i/j}, j \neq i, \\ 0 & \text{otherwise.} \end{cases}$$

This condition implies that the probability that a single mutant invades a homogeneous population is exactly of order ϵ , and that the probability that two or more mutants invade at the same time is $o(\epsilon)$. Fudenberg and Imhof (2006) show that for small $\epsilon > 0$, the standard mutation extension of an imitation process spends most of the time at the states s_1, \dots, s_K ; moreover, transitions between these states occur mainly along the edges of the state space, where only two pure strategies are present.

Let $\rho_{ij}^{(\theta)}$ be the probability that $\{X_t^{(\theta, 0)}\}$ will be absorbed in s_j when $X_0^{(\theta, 0)} = s_{i/j}$. Define a $K \times K$ matrix $\Lambda^{(\theta)} = (\Lambda_{ij}^{(\theta)})$ by

$$\Lambda_{ij}^{(\theta)} = \mu_{ij}^{(\theta)} \rho_{ij}^{(\theta)} / M, \quad j \neq i, \quad \Lambda_{ii}^{(\theta)} = 1 - \sum_{j \neq i} \mu_{ij}^{(\theta)} \rho_{ij}^{(\theta)} / M,$$

where M is some constant chosen sufficiently large so that all of the diagonal elements are strictly positive. This matrix can be regarded as the transition matrix of a Markov chain with states s_1, \dots, s_K , where the probability of a transition from s_i to $s_j \neq s_i$ is given by the probability that a mutant using j invades a population in state s_i and the probability that this mutant will take over the population in the absence of further mutations. We assume that $\Lambda^{(\theta)}$ is irreducible.⁴ Intuitively, this condition is met if there are enough edges along which transitions between homogeneous states may occur to ensure the all states of the reduced process on $\{s_1, \dots, s_K\}$ are recurrent. This is the case if, under the no-mutation process, at every state in $\mathcal{S} \setminus \{s_1, \dots, s_K\}$, every existing strategy has a positive chance of increasing its share, as for example under the aspiration-and-imitation process of Binmore and Samuelson (1997) and the frequency-dependent Moran process considered in Fudenberg et al. (2006). For these processes, $\Lambda^{(\theta)}$ is a strictly positive stochastic matrix. The $\Lambda^{(\theta)}$ matrix is not irreducible for the Robson and Vega-Redondo process, because players are assumed to emulate the more successful strategy with probability one, and the strategy that is used by only one player will have a lower payoff regardless of the realization of the random matching. Although, in general, it may be difficult to compute the absorptions probabilities $\rho_{ij}^{(\theta)}$, to check irreducibility of $\Lambda^{(\theta)}$ it is enough to know for all pairs $i \neq j$ whether $\rho_{ij}^{(\theta)}$ is positive.

We now show that limit distributions are continuous as the number of rounds of matching increases in such smoothed imitation processes.

Theorem 5. *Suppose that for every $\theta \in \Theta \cup \{\theta^*\}$ the no-mutation process $\{X_t^{(\theta,0)}\}$ is an imitation process and that $\{X_t^{(\theta,\epsilon)}\}$ is a standard mutation extension of $\{X_t^{(\theta,0)}\}$. Suppose that for all $\epsilon > 0$ and every $\theta \in \Theta \cup \{\theta^*\}$, $\{X_t^{(\theta,\epsilon)}\}$ is irreducible; denote the unique invariant distribution by $v^{(\theta,\epsilon)}$. Suppose also that $\Lambda^{(\theta)}$ is irreducible for every $\theta \in \Theta \cup \{\theta^*\}$. Finally, suppose that for $i \neq j$, $\mu_{ij}^{(\theta)} \rightarrow \mu_{ij}^{(\theta^*)}$ as $\theta \rightarrow \theta^*$ and that all the functions $f_{ij}^{(\epsilon)}(z)$ are continuous in z . Then the limit distributions*

$$\lambda^\theta = \lim_{\epsilon \rightarrow 0} v^{(\theta,\epsilon)}, \quad \theta \in \Theta, \quad \lambda^{\theta^*} = \lim_{\epsilon \rightarrow 0} v^{(\theta^*,\epsilon)}$$

exist and $\lim_{\theta \rightarrow \theta^*} \lambda^\theta = \lambda^{\theta^*}$, i.e. $\lim_{\theta \rightarrow \theta^*} \lim_{\epsilon \rightarrow 0} v^{(\theta,\epsilon)} = \lim_{\epsilon \rightarrow 0} \lim_{\theta \rightarrow \theta^*} v^{(\theta,\epsilon)}$.

Proofs. We give two different short proofs, one drawing on Theorem 1 and the Fudenberg and Imhof (2006) result on imitation processes, the other based on Theorem 1(a) and the least-cost-trees analysis of Theorems 3 and 4.

First proof: The assumptions of Theorem 1 of Fudenberg and Imhof (2006) are satisfied, so it follows that for each θ the limit distribution exists, and that the limit distributions are all concentrated on the homogeneous states s_1, \dots, s_K . Moreover, the probabilities the limit distribution assigns to the homogeneous states are given by $\eta_1^{(\theta)}, \dots, \eta_K^{(\theta)}$, where $\eta^{(\theta)}$ is the unique invariant distribution of $\Lambda^{(\theta)}$. From Theorem 1 (a) and (c), $\Lambda^{(\theta)} \rightarrow \Lambda^{(\theta^*)}$ as $\theta \rightarrow \theta^*$, and it follows by Theorem 1(b) that $\eta^{(\theta)} \rightarrow \eta^{(\theta^*)}$.

Second proof: As noted above, the limit distribution of the θ^* process is concentrated on the pure states s_i . The assumption that $\Lambda^{(\theta)}$ is irreducible implies that for any θ (including θ^*), there is an s_i -tree with probability that is $O(\epsilon^{K-1})$.⁵ The nonnegligible trees consist of transitions from each s_i to some $s_{i/j}$ together with transitions that have positive probability in the no-mutation process. Part (a) of Theorem 1 implies that the transition probabilities are continuous as $\theta \rightarrow \theta^*$, which implies that ratios of transition probabilities of the latter type converge to one (because the denominator is bounded away from 0). Our definition of a standard mutation extension ensures that the former also have ratios that converge in the $\epsilon \rightarrow 0$ limit to constants that approach one as $\theta \rightarrow \theta^*$, so we can apply the argument of Theorem 4. \square

⁴ Together with our assumption that $\Lambda_{ii}^{(\theta)} > 0$ for all i , this also ensures that $\Lambda^{(\theta)}$ is aperiodic.

⁵ Such a tree can be explicitly constructed as follows. Pick some pure state $s_j \neq s_i$. Choose a path from s_j to s_i that is not self-intersecting and that has positive probability in $\Lambda^{(\theta)}$. Add edges to the tree corresponding to the transitions in this path. If this path does not include all pure states, pick another pure state and a path from this state to s_i , and add edges to the tree corresponding to the transitions on this path until s_i or another state for which a transition has already been defined is reached. Once transitions out of all pure states have been defined, pick transitions that lead from each transitory state to a pure state, and add these to the tree as needed.

Under the assumptions of Theorem 5, $\lambda_{s_i}^{(\theta)} > 0$ for all homogeneous states s_i and $\lambda_s^{(\theta)} = 0$ for all other states s . In this case one may be interested in how the weight given to each homogeneous state depends on the payoff matrix, and on the limiting behavior as the population size or some other parameter changes. Fudenberg and Imhof (2008) study the asymptotic behavior of the limit distribution in large populations for a class of imitation processes and provide conditions under which a given strategy is selected in the sense that as the population grows to infinity, the limit distribution converges to a point mass on the state where only this strategy is used.

5. Basins of attraction

A common approach to determining the limit distribution for an evolutionary process with small mutation rates is to determine the basins of attraction of the corresponding no-mutation process and then calculate the limit distribution by a mutation counting argument. For this reason it may be interesting to know the relationship between the basins of attraction and the frequency of interaction, as in the next result, which will let us show that the difference between the limit distributions in the models of Kandori et al. (1993) and Robson and Vega-Redondo (1996) is a consequence of the different sizes of the respective basins.

We now assume that there are always exactly m rounds in each period.

If $\Omega \subset \mathcal{S}$ is the union of one or more recurrent (i.e. closed and irreducible) classes of $\{X_t^{(m)}\}$, the basin of attraction of Ω for $\{X_t^{(m)}\}$ is (see Ellison, 2000)

$$\mathcal{D}^{(m)}(\Omega) = \{s \in \mathcal{S} \mid P\{\exists t_0 \text{ s.t. } X_t^{(m)} \in \Omega \forall t \geq t_0 \mid X_0^{(m)} = s\} = 1\}.$$

Let $\mathcal{D}^*(\Omega)$ denote the basin of attraction of Ω for $\{X_t^*\}$.

Theorem 6. *Suppose that the chains $\{X_t^{(m)}\}$, $m = 1, 2, \dots$, and $\{X_t^*\}$ have the same recurrent classes, and that for every $i \in \mathcal{S}$, the support of the distribution of Z_i is finite.*

Let $\Omega \subset \mathcal{S}$ be the union of one or more of the recurrent classes.

(a) *For every $m \in \mathbb{N}$,*

$$\mathcal{D}^{(m)}(\Omega) \subseteq \mathcal{D}^{(1)}(\Omega).$$

(b) *If for every $i, j \in \mathcal{S}$, the set $\{x: f_{i,j}(x) = 0\}$ is convex, then for every $m \in \mathbb{N}$,*

$$\mathcal{D}^{(m)}(\Omega) = \mathcal{D}^{(1)}(\Omega) \subseteq \mathcal{D}^*(\Omega).$$

Proof. Let Ω' denote the set of all recurrent states in $\mathcal{S} \setminus \Omega$. If $\Omega' = \emptyset$, then $\mathcal{D}^{(m)}(\Omega) = \mathcal{D}^*(\Omega) = \mathcal{S}$ for all m . Assume now that $\Omega' \neq \emptyset$.

(a) We show that $\mathcal{S} \setminus \mathcal{D}^{(1)}(\Omega) \subseteq \mathcal{S} \setminus \mathcal{D}^{(m)}(\Omega)$. Let $s \in \mathcal{S} \setminus \mathcal{D}^{(1)}(\Omega)$. Then there is a path (i_1, i_2, \dots, i_k) such that $i_1 = s$, $i_k \in \Omega'$, and

$$p_{i_j, i_{j+1}}^{(1)} = E f_{i_j, i_{j+1}}(Z_{i_j}) > 0 \quad \text{for } j = 1, \dots, k - 1.$$

Thus for every $j = 1, \dots, k - 1$, there exists z_{i_j} such that $P\{Z_{i_j} = z_{i_j}\} > 0$ and $f_{i_j, i_{j+1}}(z_{i_j}) > 0$. Let $Z_{i_j, 1}, \dots, Z_{i_j, m}$ be i.i.d. copies of Z_{i_j} . Then

$$P\{Z_{i_j, 1} = Z_{i_j, 2} = \dots = Z_{i_j, m} = z_{i_j}\} = [P\{Z_{i_j} = z_{i_j}\}]^m > 0,$$

and it follows that

$$p_{i_j, i_{j+1}}^{(m)} = E f_{i_j, i_{j+1}}\left(\frac{Z_{i_j, 1} + \dots + Z_{i_j, m}}{m}\right) \geq [P\{Z_{i_j} = z_{i_j}\}]^m f_{i_j, i_{j+1}}(z_{i_j}) > 0,$$

where we used that $f_{i_j, i_{j+1}}$ is nonnegative. Hence $P\{X_{k-1}^{(m)} \in \Omega' \mid X_0^{(m)} = s\} \geq \prod_{j=1}^{k-1} p_{i_j, i_{j+1}}^{(m)} > 0$, which shows that $s \notin \mathcal{D}^{(m)}(\Omega)$.

(b) Let $N_{i,j} := \{x: f_{i,j}(x) = 0\}$ be convex for all $i, j \in \mathcal{S}$. We show that $\mathcal{S} \setminus \mathcal{D}^{(m)}(\Omega) \subseteq \mathcal{S} \setminus \mathcal{D}^{(1)}(\Omega)$. Let $s \in \mathcal{S} \setminus \mathcal{D}^{(m)}(\Omega)$. Then there is a path (i_1, i_2, \dots, i_k) such that $i_1 = s, i_k \in \Omega'$, and

$$p_{i_j, i_{j+1}}^{(m)} = E f_{i_j, i_{j+1}} \left(\frac{Z_{i_j,1} + \dots + Z_{i_j,m}}{m} \right) > 0 \quad \text{for } j = 1, \dots, k-1.$$

This implies that

$$P \left\{ \frac{Z_{i_j,1} + \dots + Z_{i_j,m}}{m} \in N_{i_j, i_{j+1}} \right\} < 1. \tag{14}$$

Since $N_{i_j, i_{j+1}}$ is convex, it follows that $P\{Z_{i_j} \in N_{i_j, i_{j+1}}\} < 1$. Indeed, otherwise $Z_{i_j,1}, \dots, Z_{i_j,m} \in N_{i_j, i_{j+1}}$ almost surely, so that also the convex combination $\frac{1}{m}(Z_{i_j,1} + \dots + Z_{i_j,m})$ lies in $N_{i_j, i_{j+1}}$ almost surely, which contradicts (14). Therefore, $p_{i_j, i_{j+1}}^{(1)} = E f_{i_j, i_{j+1}}(Z_{i_j}) > 0$ for $j = 1, \dots, k-1$. Hence $s \notin \mathcal{D}^{(1)}(\Omega)$. In view of part (a), it follows that $\mathcal{D}^{(m)}(\Omega) = \mathcal{D}^{(1)}(\Omega)$.

We next show that $\mathcal{S} \setminus \mathcal{D}^*(\Omega) \subseteq \mathcal{S} \setminus \mathcal{D}^{(1)}(\Omega)$. Let $s \in \mathcal{S} \setminus \mathcal{D}^*(\Omega)$. Then there is a path (i_1, i_2, \dots, i_k) such that $i_1 = s, i_k \in \Omega'$, and

$$p_{i_j, i_{j+1}}^* = f_{i_j, i_{j+1}}(E Z_{i_j}) > 0 \quad \text{for } j = 1, \dots, k-1.$$

That is, $E Z_{i_j} \notin N_{i_j, i_{j+1}}$, and so, as $N_{i_j, i_{j+1}}$ is convex, $P\{Z_{i_j} \in N_{i_j, i_{j+1}}\} < 1$. Therefore, $p_{i_j, i_{j+1}}^{(1)} = E f_{i_j, i_{j+1}}(Z_{i_j}) > 0$ for $j = 1, \dots, k-1$. Hence $s \notin \mathcal{D}^{(1)}(\Omega)$. \square

Remark 5. The convexity assumption in part (b) of Theorem 6 is satisfied in models derived from two-action games if f_{ij} is of the form $f_{ij}(x_1, x_2) > 0$ for $x_1 > x_2$ and $f_{ij}(x_1, x_2) = 0$ otherwise. It would also be satisfied under any process that was such that either $f_{ij}(x) = 0$ for all x or $f_{ij}(x) > 0$ for all x . The condition is, however, more demanding in models involving games with more than two actions. For example, it would not be satisfied for the best-response dynamic in a 3×3 game if strategy A was the best response to A , strategy B was the best response to B , but C was the best response to a mixture of A and B .

The next example shows that the conclusion of Theorem 6(b) is false if the convexity assumption is dropped:

Example 5. Suppose that the common state space of the chains $\{X_t^{(m)}\}, m = 1, 2, \dots, \{X_t^*\}$ is $S = \{1, 2, 3\}$, and that the states 1 and 3 are absorbing states of all of these chains. If $X_0^{(m)}$ [resp. X_0^*] is in state 2, the transition depends on the realization of m (resp. infinitely many) flips of a fair coin with possible values 0 and 1: When the average of these realizations is z , the chain switches to state 1 with probability $f_{21}(z)$ and to state 3 otherwise, where $f_{21}(x) = 1 - 2|x - \frac{1}{2}|$. Here $N_{21} = \{x: f_{21}(x) = 0\} = \{0, 1\}$ is not convex, and if $m > 1$, then

$$P \left\{ \frac{Z_{21} + \dots + Z_{2m}}{m} \in N_{21} \right\} < 1,$$

but this does not imply that $P\{Z_{21} \in N_{21}\} < 1$. We have

$$\mathcal{D}^{(1)}(\{3\}) = \{2, 3\}, \quad \mathcal{D}^{(m)}(\{3\}) = \{3\} \quad \text{for all } m \geq 2 \text{ and } \mathcal{D}^*(\{3\}) = \{3\}.$$

Thus the inclusion in Theorem 6(a) is strict for $m \geq 2$, but neither the claimed equality nor the claimed inclusion in Theorem 6(b) holds.

The following example of Robson and Vega-Redondo (1996, Section 2) shows that $\mathcal{D}^{(m)}(\Omega)$ can be a proper subset of $\mathcal{D}^*(\Omega)$, that is $\mathcal{D}^{(m)}(\Omega) \subset \mathcal{D}^*(\Omega)$ and $\mathcal{D}^{(m)}(\Omega) \neq \mathcal{D}^*(\Omega)$. The convexity assumption in Theorem 6(b) is satisfied in the example.

Example 6. Consider a population of size N , where N is even and $N \geq 4$. In each round, all N individuals are paired at random to play a symmetric 2×2 game with payoff matrix

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} = \begin{pmatrix} \sqrt{3} & 0 \\ 1 & 1 \end{pmatrix}.$$

Let $X_t^{(m)}$ denote the number of individuals using strategy 1. The payoff vector in period t , round r is given by

$$Y_{t,r} = \begin{bmatrix} Y_{t,r}^{(1)} \\ Y_{t,r}^{(2)} \end{bmatrix},$$

where $Y_{t,r}^{(i)}$ is the (sub-population) average payoff of all players that use strategy i , where $Y_{t,r}^{(i)} = 0$ if strategy i is not used. To calculate the distribution of $Y_{t,r}$ suppose $X_t^{(m)} = x$ and $0 < x < N$. Clearly,

$$Y_{t,r}^{(2)} = 1.$$

Let $C_{t,r}$ denote the random number of cross-pairings. The distribution of $C_{t,r}$ depends only on x . There are

$$\begin{aligned} \frac{X_t^{(m)} - C_{t,r}}{2} & \quad (1, 1)\text{-pairs,} \\ C_{t,r} & \quad \text{mixed pairs,} \\ \frac{N - X_t^{(m)} - C_{t,r}}{2} & \quad (2, 2)\text{-pairs.} \end{aligned}$$

Thus

$$Y_{t,r}^{(1)} = \frac{1}{X_t^{(m)}} [a_{11}(X_t^{(m)} - C_{t,r}) + a_{12}C_{t,r}] = \frac{\sqrt{3}}{X_t^{(m)}} [X_t^{(m)} - C_{t,r}].$$

The (no-mutation) process is defined as follows: if

$$\frac{Y_{t,1}^{(1)} + \dots + Y_{t,m}^{(1)}}{m} > \frac{Y_{t,1}^{(2)} + \dots + Y_{t,m}^{(2)}}{m},$$

then $X_{t+1}^{(m)} = N$, otherwise $X_{t+1}^{(m)} = 0$. Thus, for $i = 0, \dots, N$, $f_{i,j}(y^{(1)}, y^{(2)}) \equiv 0$ for $1 \leq j \leq N - 1$, and

$$f_{i,0}(y^{(1)}, y^{(2)}) = \begin{cases} 1, & y^{(1)} < y^{(2)}, \\ 0, & y^{(1)} > y^{(2)}, \end{cases} \quad f_{i,N}(y^{(1)}, y^{(2)}) = \begin{cases} 0, & y^{(1)} < y^{(2)}, \\ 1, & y^{(1)} > y^{(2)}. \end{cases}$$

For every $m < \infty$,

$$\mathcal{D}^{(m)}(\{0\}) = \{0, 1\}, \quad \mathcal{D}^*(\{0\}) \supseteq \{0, 1, 2\}. \tag{15}$$

To see this note that if only one agent plays 1, he is always in a mixed pair and receives payoff 0, while all 2-players receive payoff 1. Thus $1 \in \mathcal{D}^{(m)}(\{0\})$ and $1 \in \mathcal{D}^*(\{0\})$. If two agents play 1, the probability that they are matched in all the m rounds is $1/(N - 1)^m > 0$. If this happens, they receive $\sqrt{3}$, which is more than 1, the payoff to strategy 2. Thus $2 \notin \mathcal{D}^{(m)}(\{0\})$. However, if two agents play 1 and there are infinitely many rounds, their payoff is $\sqrt{3}/(N - 1) < 1$, and it follows that $2 \in \mathcal{D}^*(\{0\})$.

Example 7. Note that when mutations are added (as in Robson and Vega-Redondo (1996)) the example satisfies the continuity assumption of Theorems 1 and 2: The behavior rule is only discontinuous when the realized payoffs to the two actions are equal, which is impossible because the payoff to 2 is irrational.

Theorem 7. Suppose that the chains $\{X_t^{(m)}\}$, $m = 1, 2, \dots$, and $\{X_t^*\}$ have the same recurrent classes. Let $\Omega \subset \mathcal{S}$ be the union of one or more of the recurrent classes. Suppose that for all $i, j \in \mathcal{S}$, f_{ij} is continuous at EZ_i . Then there exists m_0 such that for every number of rounds $m \geq m_0$,

$$\mathcal{D}^{(m)}(\Omega) \subseteq \mathcal{D}^*(\Omega) \tag{16}$$

and the inclusion is strict if and only if there exists $i \in \mathcal{D}^*(\Omega)$ and $j \in \mathcal{S} \setminus \mathcal{D}^{(m)}(\Omega)$ with $p_{ij}^{(m)} > 0$.

Proof. By Theorem 1(a), the present continuity assumption on f_{ij} implies that for all i, j , $p_{ij}^{(m)} \rightarrow p_{ij}^*$ as $m \rightarrow \infty$. It follows that there is m_0 such that

$$p_{ij}^* > 0 \implies p_{ij}^{(m)} > 0 \quad \text{for all } m \geq m_0. \tag{17}$$

Now if $i \notin \mathcal{D}^*(\Omega)$, then there must exist some recurrent class Ω' with $\Omega \cap \Omega' = \emptyset$ such that Ω' can be reached by $\{X_t^*\}$ starting from i . It therefore follows from (17), that for $m \geq m_0$, Ω' can also be reached by $\{X_t^{(m)}\}$ starting from i . Hence $i \notin \mathcal{D}^{(m)}(\Omega)$. This proves inclusion (16).

Finally, if (16) holds, then this inclusion is strict if and only if there exists some state $i \in \mathcal{D}^*(\Omega)$ with $i \notin \mathcal{D}^{(m)}(\Omega)$. The proof is completed by observing that $i \in \mathcal{D}^{(m)}(\Omega)$ if and only if $p_{ij}^{(m)} = 0$ for all $j \in \mathcal{S} \setminus \mathcal{D}^{(m)}(\Omega)$. \square

Note that Example 5 also shows that the inclusion stated in Theorem 7 need not hold for every m , but only for $m \geq m_0$; in the example $m_0 = 2$.

Example 8. In the setting of Example 6, (15) continues to hold for every m , if $f_{ij}(y^{(1)}, y^{(2)}) \equiv 0$ for $1 \leq j \leq N - 1$, and with some small $\delta > 0$,

$$f_{i,0}(y^{(1)}, y^{(2)}) = \begin{cases} 1, & y^{(1)} < y^{(2)} - \delta, \\ 0, & y^{(1)} > y^{(2)} + \delta, \end{cases} \quad f_{i,N}(y^{(1)}, y^{(2)}) = \begin{cases} 0, & y^{(1)} < y^{(2)} - \delta, \\ 1, & y^{(1)} > y^{(2)} + \delta, \end{cases}$$

and all f_{ij} arbitrarily smooth. This shows that despite smoothness of the functions f_{ij} there need not be equality in (16) even if m is large.

The next example shows that the inclusion (16) need not hold even when m is large if we merely assume that Ω is a recurrent class of $\{X_t^{(m)}\}$ for all $m = 1, 2, \dots$ and of $\{X_t^*\}$ and that the f_{ij} are continuous.

Example 9. Consider a symmetric 2-player game with pure strategies 1, 2, 3 and payoff matrix

$$\begin{pmatrix} 1 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 2 & 2 \\ \frac{1}{2} & 0 & 0 \end{pmatrix}.$$

The population consists of N agents, which may be of type A or B . Type A agents play strategy 1. Type B agents play strategy 2 or 3, each with probability $\frac{1}{2}$. Every time a type B agents plays, he makes a new randomized decision independent of everything else. Let $X_t^{(m)}$ and X_t^* denote the number of type A agents in period t . If the average payoff p of an agent is at least 1, he keeps his type, otherwise he changes to the opposite type with probability $1 - p$. That is, the probability of keeping the type is given by the continuous function $\min\{1, p\}$.

If $X_t^{(m)} = N$ or $X_t^* = N$, there will be only pairs of type A agents. Thus everyone receives every time payoff 1, so that no-one will change his type. Hence N is an absorbing state of $\{X_t^{(m)}\}$ for every m and of $\{X_t^*\}$. If $1 \leq X_t^{(m)} \leq N - 1$, then in every round there is a positive probability that a mixed pair is formed, so that both agents receive payoff $\frac{1}{2}$. Using this fact, one can show that $p_{i,i+1}^{(m)} > 0$ and $p_{i,i-1}^{(m)} > 0$ for all m . Similarly, if $1 \leq X_t^* \leq N - 1$, then the average payoff of every agent is less than 1, and it follows that $p_{i,i+1}^* > 0$ and $p_{i,i-1}^* > 0$. If $X_t^{(m)} = 0$, there will be only B pairs and there is a positive probability that every agent plays strategy 2 in each of the m rounds. In this case, all average payoffs are 0 and every agent switches his type. Hence $p_{0N}^{(m)} > 0$ for all m . However, if there are infinitely many rounds and $X_t^* = 0$, the average payoff to every agent is 1, so that no-one switches. That is, $p_{00}^* = 1$. Consequently,

$$\mathcal{D}^{(m)}(\{N\}) = \{0, \dots, N\}, \quad \mathcal{D}^*(\{N\}) = \{N\}.$$

Acknowledgments

We thank Fernando Vega-Redondo for helpful conversations and the editor and two referees for their comments.

References

Bergin, J., Lipman, B.L., 1996. Evolution with state-dependent mutations. *Econometrica* 64, 943–956.
 Billingsley, P., 1995. *Probability and Measure*. Wiley, New York.
 Binmore, K., Samuelson, L., 1997. Muddling through: Noisy equilibrium selection. *J. Econ. Theory* 74, 235–265.
 Cho, G.E., Meyer, C.D., 2001. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra Appl.* 335, 137–150.

- Ellison, G., 2000. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Rev. Econ. Stud.* 67, 17–45.
- Freidlin, M., Wentzell, A., 1984. *Random Perturbations of Dynamical Systems*. Springer-Verlag, New York.
- Fudenberg, D., Imhof, L.A., 2006. Imitation processes with small mutations. *J. Econ. Theory* 131, 251–262.
- Fudenberg, D., Imhof, L.A., 2008. Monotone imitation dynamics in large populations. *J. Econ. Theory* 140, 229–245.
- Fudenberg, D., Nowak, M., Taylor, C., Imhof, L.A., 2006. Evolutionary game dynamics in finite populations with strong selection and weak mutations. *Theoretical Population Biol.* 70, 352–363.
- Imhof, L.A., Nowak, M.A., 2006. Evolutionary game dynamics in a Wright–Fisher process. *J. Math. Biol.* 52, 667–681.
- Iosifescu, M., 1980. *Finite Markov Processes and their Applications*. Wiley, New York.
- Kandori, M., Mailath, G., Rob, R., 1993. Learning, mutation, and long-run equilibria in games. *Econometrica* 61, 29–56.
- Karlin, S., Taylor, H.M., 1975. *A First Course in Stochastic Processes*. Academic Press, San Diego.
- Robson, A.J., Vega-Redondo, F., 1996. Efficient equilibrium selection in evolutionary games with random matching. *J. Econ. Theory* 70, 65–92.
- Vega-Redondo, F., 1996. *Evolution, Games, and Economic Behaviour*. Oxford Univ. Press, Oxford.
- Young, H.P., 1993. The evolution of conventions. *Econometrica* 61, 57–84.