

MIT Open Access Articles

Random matching in adaptive dynamics

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ellison, Glenn, Drew Fudenberg, and Lorens A. Imhof. "Random Matching in Adaptive Dynamics." Games and Economic Behavior 66, no. 1 (May 2009): 98–114.

As Published: <http://dx.doi.org/10.1016/j.geb.2008.04.010>

Publisher: Elsevier

Persistent URL: <http://hdl.handle.net/1721.1/101251>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-NoDerivatives





HIER

Harvard Institute of Economic Research

Discussion Paper Number 2126

Random Matching in Adaptive Dynamics

by

Glenn Ellison, Drew Fudenberg
and Lorens A. Imhof

December 2006

HARVARD UNIVERSITY
Cambridge, Massachusetts

This paper can be downloaded without charge from:
<http://post.economics.harvard.edu/hier/2006papers/2006list.html>

The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=951469>

Random matching in adaptive dynamics

Glenn Ellison, Drew Fudenberg, Lorens A. Imhof

November 9, 2006

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 doesn't 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 ergodic 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.

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, Mailath and Rob [1993] (KMR), Bergin and Lipman [1996] and Binmore and Samuelson [1997]. Neither of these motivations is 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.

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 doesn't 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 ergodic 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 ergodic, its invariant distribution is a continuous function of the noise at the zero matching-noise limit, and for non-ergodic systems 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] (RV) show that the limit distribution in KMR's 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 KMR's 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 RV adjustment process satisfies the condition of Theorem 1, so the difference between their findings and those of KMR 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 RV, but is satisfied in the frequency-dependent Moran process studied by Fudenberg, Imhof, Nowak, and Taylor [2006] and also by the frequency-dependent Wright-Fisher process studied by Imhof and Nowak [2006]. More generally, the condition is satisfied for the class of imitation process covered by our Theorem 4.

Theorems 5 and 6 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 5 replaces the continuity condition of Theorems 1 and 2 with a mild form of monotonicity that is satisfied by RV, and shows that $\mathcal{D}^{(m)} = \mathcal{D}^{(1)} \subseteq \mathcal{D}^*$. An example of RV 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 6 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 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

$$P^{Y_{t,r}|X_t^{(\theta)}=i} = P^{Z_i} \quad \text{for all } r, t \text{ and } \theta. \quad (2)$$

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

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

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 $\{R_t^{(\theta)} > 0\}$,

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

Thus, in view of (2),

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

where $q_{ij} = P \left\{ X_{t+1}^{(\theta)} = j | X_t^{(\theta)} = i, R_t^{(\theta)} = 0 \right\}$ 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

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

By the strong law of large numbers, $(Z_{i,1} + \dots + Z_{i,k})/k \rightarrow EZ_i$ almost surely. 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

$$\lim_{k \rightarrow \infty} E f_{ij} \left(\frac{1}{k} \sum_{r=1}^k Z_{i,r} \right) = p_{ij}^*. \quad (4)$$

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

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

Therefore,

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

It now follows by assumption (1) and the fact that $\epsilon > 0$ was arbitrary that

$$\lim_{\theta \rightarrow \theta^*} p_{ij}^{(\theta)} = p_{ij}^*. \quad (5)$$

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 $\theta \rightarrow \theta^*$, the sequence has some accumulation point $w \neq v^*$. Because $\Pi^{(\theta)}$ is continuous in θ and $v^{(\theta)}\Pi^{(\theta)} = v^{(\theta)}$ for all θ , this would imply that $w\Pi^* = w$. This contradicts 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 (5) 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, page 90]. It now follows as in the proof of b) that $\rho_{ij}^{(\theta)} \rightarrow \rho_{ij}^*$. \square

Remark 1 a) *Theorem 1a) 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 (4).*

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 1b), but the much simpler self-contained argument in the present proof is sufficient for our purposes.*

3.2 Illustrative examples

The following examples show that the continuity assumption in Theorem 1 is crucial. We begin with a simple artificial example in which the f_{ij} functions are not continuous and long-run behavior differs between models with a finitely and infinitely many rounds of matching.

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 payoff matrix

$$\begin{pmatrix} 2 & 2 \\ 1 & 3 \end{pmatrix}.$$

Consider of 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

$$P\left\{Y_{t,r}^{(1)} = 0, Y_{t,r}^{(2)} = 3 \mid X_t^{(m)} = 1\right\} = \frac{1}{3}, \quad P\left\{Y_{t,r}^{(1)} = 2, Y_{t,r}^{(2)} = \frac{1}{2} \mid X_t^{(m)} = 1\right\} = \frac{2}{3}. \quad (6)$$

Similarly,

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

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 (6), $EZ_1 = (\frac{4}{3}, \frac{4}{3})$, and so

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

Let $(Z_{11}^{(1)}, Z_{11}^{(2)})$, $(Z_{12}^{(1)}, Z_{12}^{(2)})$, \dots be i.i.d. copies of Z_1 . By the central limit theorem,

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

as $m \rightarrow \infty$. 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, $\lim_{m \rightarrow \infty} p_{1j}^{(m)} \neq p_{1j}^*$. A similar calculation shows 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 of $\{X_t^{(m)}\}$ does not converge to the invariant distribution of $\{X_t^*\}$.

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*

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

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 1a) implies in the case of ergodicity 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 (5) 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. Note that in Example 3, $p_{1j}^{(m)} \not\rightarrow p_{1j}^*$, and so the distribution of $X_1^{(m)}$ will not approach that of X_1^* if $P(X_0^{(m)} = 1) = P(X_0^* = 1) > 0$.

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 (5) 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), there exists a neighborhood U_1 of θ^* with $U_1 \subset U_0$ such that

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

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

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

and $\beta(\theta) \rightarrow \beta^*$, see e.g. [10, 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 $t > T$,

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

To complete the proof note that by (5) there is a neighborhood U_3 of θ^* such that if $\theta \in U_3 \cap \Theta$, $\|p^{(\theta)}(t) - p^*(t)\| < \epsilon$ for $t = 0, \dots, T$. \square

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 ergodic 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] (RV), combined with our results, shows that the order of limits can matter. RV 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 RV model coincides with the KMR 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, RV show that the long-run equilibrium places probability one on the Pareto optimal equilibrium.

The transition rule in the RV model need not satisfy our continuity condition, but it does satisfy it for generic payoff functions.¹ 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 KMR model and the long-run equilibrium of the RV model is not from a discontinuity at $m = \infty$ as in Example 1, but rather is due to a difference in the order of limits: $\lim_{m \rightarrow \infty} \lim_{\epsilon \rightarrow 0} v^{(m,\epsilon)} \neq \lim_{\epsilon \rightarrow 0} \lim_{m \rightarrow \infty} v^{(m,\epsilon)}$. Which order of limits is more relevant will depend on the magnitudes of ϵ and m . Intuitively, the difference between the RV model and the KMR model is that in the RV 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 m .

¹We only require continuity when payoffs are equal to their expectation under an infinite number of matching. Generically, a player will not be indifferent when facing a mixed strategy of the form $Ak/N + B(N - k)/N$. If this did happen, we could smooth the best response functions when the two strategies have payoffs within some small σ without affecting RV’s analysis.

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 RV 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 ergodic 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*

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

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*

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

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

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

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,

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

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 1b): The hypothesis about the limit of the ratios implies that the transition probabilities are continuous as $\theta \rightarrow \theta^*$. Hence Theorem 1b) 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 2 *The RV 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.

Specifically, we will say that state i is a long-run equilibrium for θ near θ^* if every neighborhood of θ^* contains a θ for which $\lambda_i^\theta > 0$. We will say that an i -tree z_i is negligible relative to state i for small ϵ (for θ near θ^*) if there exists a neighborhood U of θ^* such that

$$\frac{\prod_{(j,k) \in z_i} p_{jk}^{\theta, \epsilon}}{\max_{z' \in Z_i} \left(\prod_{(j,k) \in z'} p_{jk}^{\theta, \epsilon} \right)} \rightarrow 0 \text{ as } \epsilon \rightarrow 0,$$

for all θ in U .

Remark 3 *A weaker sufficient condition for the conclusion of Theorem 3 is that $p_{jk}^{\theta, \epsilon}$ be continuous as $\theta \rightarrow \theta^*$ for all (j, k) , and that*

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

holds for those ordered pairs (j, k) for which there exists a state i and an i -tree z_i such that i is a long-run equilibrium for θ near θ^ and z_i contains (j, k) and is nonnegligible relative to i for small ϵ .*

Proof. The argument for this extension is straightforward. It suffices to show that $\lambda_i^\theta \rightarrow \lambda_i^{\theta^*}$ for states i that are long-run equilibria in the θ^* process. In the proof of Theorem 3 we noted that

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

Let W be the set of all trees that are nonnegligible relative to a state that is a long-run equilibrium for θ near θ^* . To prove the extension of the theorem it suffices to show that each of the four sums on the right side of this expression are asymptotically unchanged if the sum is restricted to trees belonging to W , e.g. to show

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

This is immediate from the definition of negligible. \square

We now discuss a class of smoothed imitation processes as an illustration of when the order of limits doesn't 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.

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

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

with $\mu_{ij}^{(\theta)} \rightarrow \mu_{ij}^{(\theta^*)}$ as $\theta \rightarrow \theta^*$. 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.

We assume that this matrix is irreducible and aperiodic; one case in which that is true is if every strategy has a positive chance of increasing its share, so that $\Lambda^{(\theta)}$ is a strictly positive stochastic matrix. The $\Lambda^{(\theta)}$ matrix is not irreducible for the RV 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. The property would be satisfied if under the no-mutation process, every existing strategy had some probability of being copied, 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].

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

Theorem 4 *Suppose that the no-mutation processes $X_t^{(\theta,0)}$ are imitation processes for every $\theta \in \Theta \cup \{\theta^*\}$. Let $X_t^{(\theta,\epsilon)}$ be standard mutation extensions of the $X_t^{(\theta,0)}$ and suppose that for all $\epsilon > 0$ and every $\theta \in \Theta \cup \{\theta^*\}$, $\{X_t^{(\theta,\epsilon)}\}$ is irreducible with unique ergodic distribution $v^{(\theta,\epsilon)}$ and that $\Lambda^{(\theta)}$ is irreducible and has a unique invariant distribution. Finally, suppose that all the functions $f_{ij}^{(\epsilon)}(z)$ are continuous in z . Then the limit distributions*

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

exist and $\lim_{\theta \rightarrow \theta^} \lambda^\theta = \lambda^*$, 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 1a and the least-cost-trees analysis of Theorem 3 and Remark 3.

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 1a) and c), $\Lambda^{(\theta)} \rightarrow \Lambda^{(\theta^*)}$ as $\theta \rightarrow \theta^*$, and it follows by Theorem 1b) 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 and aperiodic implies that for any θ (including θ^*), there is an s_i -tree with probability that is $O(\epsilon^{K-1})$. 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 Remark 3. \square

4.2 Noisy matching in finite-memory fictitious play

Ellison [1997] discusses a model in which $N \geq 4$ players are repeatedly randomly matched to play the symmetric 3×3 coordination game shown below (with η a small positive number).

$$\begin{pmatrix} 10 + \eta & 0 & 0 \\ 1 & 1 & 1 \\ 11 & -100 & 0 \end{pmatrix}$$

He assumes that players are matched with only one opponent in each period ($m = 1$), that players only observe the outcomes of matches they participate in, and that players follow finite-memory fictitious play learning rules (with a ten period memory). When no mutations are present, play converges in the long run to all players playing B . Ellison notes, however, that there is another behavior that is more relevant in the intermediate run. A single mutation away from the steady-state sets in motion a fairly regular cycle in which most of the population shifts to A , then to C , then to B , and then back to A again. In even moderately large populations (e.g. $N = 100$) this cycle will persist for millions of periods before play returns to the steady state.

We can illustrate and elaborate on this point using a version of the model with ϵ -mutations. Although the invariant distributions do satisfy $\lim_{\epsilon \rightarrow 0} v^{(1,\epsilon)} = v^{(1,0)}$, $v^{(1,\epsilon)}$ will look very different from $v^{(1,0)}$ if ϵ is not extremely small, e.g. if there is one mutation somewhere in the population every thousand periods. A model in which players are matched with every opponent in every period ($m = \infty$) would not share this property, as all players would continue to play B after any moderate number of mutations, and even if a shift to A was set in motion by all N players mutating simultaneously, the system would return to the steady state after a single cycle through A and C .

The $m = 1$ and $m = \infty$ versions of this model can be fit directly into our framework by letting the state space \mathcal{S} be the set of $N \times 10$ matrices describing possible observations in each of the ten previous periods by each player. In order to apply our theorem to examine the behavior of the system with large finite m , however, some modification of the model would be needed. Our theorem assumes a common state space across m , whereas the possible realizations of average play change with m . One way to modify the large finite m model so that our theorem would be applicable would be to have players beliefs always lie on some finite grid (e.g. by discretizing them to the closest value involving multiples of $1/N$).² Theorem 1 would then imply that the behavior of the ϵ -perturbed models is continuous as $m \rightarrow \infty$. Hence, behavior with a large m would look very different from behavior with $m = 1$. This highlights that the assumption of a small number of rounds is important for two reasons: First, it is an amplification mechanism that lets a small number of mutations influence future play (because they are a disproportionate part of the history of players who see them); second, it allows belief heterogeneity to persist in the population, which plays an important role in the recurrence of the cycles.

5 Basins of attraction

In this section we 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])

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

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

Theorem 5 *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.

²One might also need to smooth best responses when players are nearly indifferent between two strategies if any such indifferences can occur in the $m = \infty$ model for the given value of N .

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

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

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

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

This 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

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

This implies that

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

Since $N_{i_j, i_{j+1}}$ is convex, it follows that $P\{Z_{i_j} \in N_{i_j, i_{j+1}}\} < 1$. Therefore, $p_{i_j, i_{j+1}}^{(1)} = Ef_{i_j, i_{j+1}}(Z_{i_j}) > 0$ 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

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

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

Remark 4 *The convexity assumption in part b) of Theorem 5 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 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 5 b) is satisfied in the example.

Example 5 *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 $m < \infty$,

$$\mathcal{D}^{(m)}(\{0\}) = \{0, 1\}, \quad \mathcal{D}^*(\{0\}) \supseteq \{0, 1, 2\}. \quad (7)$$

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 6 Note that when mutations are added (as in RV) 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 6 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 $m \geq m_0$,

$$\mathcal{D}^{(m)}(\Omega) \subseteq \mathcal{D}^*(\Omega) \quad (8)$$

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

$$p_{ij}^* > 0 \quad \implies \quad p_{ij}^{(m)} > 0 \quad \text{for all } m \geq m_0. \quad (9)$$

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 (9), 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 (8). \square

Example 7 In the setting of Example 5, (7) 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 (8) even if m is large.

The next example shows that the inclusion (8) need not hold for m sufficiently 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 8 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 agent 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 everytime 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,

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

References

- [1] J. Bergin, B. L. Lipman, Evolution with state-dependent mutations, *Econometrica* 64 (1996) 943-956.
- [2] K. Binmore, L. Samuelson, Muddling through: noisy equilibrium selection, *J. Econ. Theory* 74 (1997) 235-265.
- [3] G.E. Cho, C.D. Meyer, Comparison of perturbation bounds for the stationary distribution of a Markov chain, *Lin. Alg. Appl.* 335 (2001) 137-150.
- [4] G. Ellison, Learning from personal experience: one rational guy and the justification of myopia, *Games Econ. Behav.* 19 (1997) 180-210.

- [5] G. Ellison, Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution, *Rev. Econom. Studies* 67 (2000) 17-45.
- [6] M. Freidlin, A. Wentzell, *Random Perturbations of Dynamical Systems*, Springer Verlag, New York, 1984.
- [7] D. Fudenberg, L. A. Imhof, Imitation processes with small mutations, *J. Econ. Theory* 131 (2006) 251-262.
- [8] D. Fudenberg, L. A. Imhof, M. Nowak, and C. Taylor, Evolutionary game dynamics in finite populations with strong selection and weak mutations, *Theoretical Population Biology* 70 (2006) 352-363.
- [9] L.A. Imhof, M.A. Nowak, Evolutionary game dynamics in a Wright-Fisher process, *J. Math. Biol.* 52 (2006) 667-681.
- [10] M. Iosifescu, *Finite Markov Processes and Their Applications*, Wiley, New York, 1980.
- [11] M. Kandori, G. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, *Econometrica* 61 (1993) 29-56.
- [12] S. Karlin, H.M. Taylor, *A First Course in Stochastic Processes*, Academic Press, San Diego, 1975.
- [13] A.J. Robson, F. Vega-Redondo, Efficient equilibrium selection in evolutionary games with random matching, *J. Econ. Theory* 70 (1996) 65-92.
- [14] H.P. Young, The evolution of conventions, *Econometrica* 61 (1993) 57-84.