Monotone Imitation Dynamics in Large Populations

Citation

Fudenberg, Drew, and Lorens A. Imhof. 2008. Monotone imitation dynamics in large populations. *Journal of Economic Theory* 140, no. 1: 229-245.

Published Version

<http://dx.doi.org/10.1016/j.jet.2007.08.002>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3196338>

Terms of Use

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

Share Your Story

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

[Accessibility](#)

HIER

Harvard Institute of Economic Research

Discussion Paper Number 2141

Monotone Imitation Dynamics in
Large Populations

by

Drew Fudenberg
and
Lorenz A. Imhof

November 2007

HARVARD UNIVERSITY
Cambridge, Massachusetts

This paper can be downloaded without charge from:
<http://post.economics.harvard.edu/hier/2007papers/2007list.html>

The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1029691>

Monotone Imitation Dynamics in Large Populations

Drew Fudenberg
Department of Economics
Harvard University
Cambridge, MA 02138
dfudenberg@harvard.edu

Lorens A. Imhof
Department of Economics
Bonn University
D-53113 Bonn
Germany
limhof@uni-bonn.de

February 28, 2007

Abstract

We analyze a class of imitation dynamics with mutations for games with any finite number of actions, and give conditions for the selection of a unique equilibrium as the mutation rate becomes small and the population becomes large. Our results cover the multiple-action extensions of the aspiration-and-imitation process of Binmore and Samuelson [3] and the related processes proposed by Benaïm and Weibull [2] and Traulsen et al. [24], as well as the frequency-dependent Moran process studied by Fudenberg et al. [15]. We illustrate our results by considering the effect of the number of periods of repetition on the selected equilibrium in repeated play of the prisoner's dilemma when players are restricted to a small set of simple strategies.

JEL classification: C62; C72; C73.

Keywords: Equilibrium selection; Learning model; Limit distribution; Markov process; Prisoner's dilemma.

1 Introduction

We study a class of imitation dynamics in large populations playing an $n \times n$ game. Our main assumptions are that at every time step, at most one agent changes his strategy, and that this agent may only imitate a strategy that is

currently in use. In addition, we assume that if only two strategies are present in the population, the probabilities of the possible transitions depend only on the current payoffs to these strategies, and that the expected motion is in the direction of the better response; this is the sense in which the dynamics are "monotone."¹ Finally we assume that a small mutation term makes the system ergodic. This class of dynamics encompasses various models that have been studied in the literature, mainly for 2×2 games; e.g. the aspiration-and-imitation process of Binmore and Samuelson [3], the related imitation processes proposed by Benaïm and Weibull [2], Björnerstedt and Weibull [5] and Traulsen et al. [24], and the frequency-dependent Moran process introduced by Nowak et al. [20].

The models we consider have a unique "limit distribution" as the mutation rate becomes small; we are interested in the convergence of this limit distribution as the population size becomes large. Our analysis builds on our work in Fudenberg and Imhof [11], which shows that for every fixed population size, the limit distribution can be obtained as the unique invariant distribution of the $n \times n$ transition matrix of a certain embedded chain. The entries of this matrix are given by absorption probabilities of the original process restricted to the edges of the state space where only two strategies are present. The large-population behavior of the limit distribution is therefore determined by the large population behavior of these absorption probabilities, so we need to develop a fairly precise characterization of their asymptotic behavior. To this end we approximate the probabilities using Riemann sums and apply the Laplace method to analyse the behavior of the resulting integrals. This yields simple criteria that determine whether the probabilities converge to zero and, in the case of convergence, explicit expressions for the rate of convergence. We then show how these rates, if interpreted at transition costs, can be combined with the now-familiar machinery of "least cost

¹Friedman [10] and Samuelson and Zhang [22] consider related monotonicity conditions in deterministic models with a continuum population; our condition is weaker in that it applies only to states where only two strategies are present. Our model allows for stochastic transitions, so it is not a monotone system in their sense even in games with only two strategies, but when only two strategies are present the mean field of the processes we consider satisfies their assumptions. However, our model need not satisfy the "Darwinian" monotonicity assumption in the finite-population model of Kandori et al. [18], which requires that the probability that the no-mutation process moves in the direction of the better response is equal to 1. Our monotonicity condition requires only that this probability is larger than that of a step in the opposite direction; this is sufficient to imply that the expected motion of the system (i.e. its mean field) is in the direction of increasing payoffs.

graphs²" to solve the equilibrium selection problem.

In the 2×2 case there are only two absorption probabilities to consider and the selected strategy can be determined by checking whether their ratio converges to 0 or to infinity. Our selection result for this case coincides with that of Binmore and Samuelson [3]³. Neither the asymptotics for the individual absorptions probabilities nor the least-cost arguments are needed in this case, and Binmore and Samuelson's direct analysis of an explicit expression of the ergodic distribution is simpler than our approach. However, in the general case, the limit distribution cannot be expressed in terms of ratios of absorption probabilities and the individual rates of convergence are required to determine the least-cost graph.

As an application we study strategic forms matrices arising from finitely or infinitely repeated play of the prisoner's dilemma when players are restricted to a small set of simple strategies such as "tit for tat," and the payoffs are evaluated with the time-average criterion. For any finite number m of repetitions, the system converges to the state where all agents play "Always Defect," but with an infinite number of repetitions the "Always Defect" is weakly dominated, so its share of the population goes to 0, and the long-run distribution is some combination of the strategies "Tit for Tat" and "Perfect Tit-for-tat." We show that the conclusion from large-population limit of the case $m = \infty$ is robust in the sense that it also applies to the case of a finite but large number of repetitions and a range of "intermediate" population sizes. This proves a conjecture of Imhof et al. [16]. We state and prove a more general version of this result, relating the large-population limit in one game to the invariant distributions of finite-population versions of "nearby" games.

2 The model

Consider a symmetric two-player game with pure strategies $1, \dots, n$ and pay-off matrix $A = (a_{ij})_{i,j=1}^n$. We consider a population of size $N \geq 2$ and describe its evolution by a homogeneous Markov chain $\{X(t; \epsilon, N) : t = 0, 1, \dots\}$ with

²This way of computing the support of the limit distribution was introduced by Freidlin and Wentzell [9]; its use in evolutionary game theory is due to Kandori et al. [18] and Young [25]. See Fudenberg and Levine [14] or Samuelson [21] for an introduction to the relevant probability theory and survey of some of its applications to game dynamics.

³The selection result of Binmore and Samuelson applies to their more general muddling process for 2×2 games. In [12] we analyze a multi-dimensional muddling analog of generalized muddling, and show that it covers the extension of the Ellison and Fudenberg [7] word-of-mouth learning model to the case of more than two brands.

state space

$$\mathcal{S}_N = \{(x_1, \dots, x_n) \in \{0, 1, \dots, N\}^n : x_1 + \dots + x_n = N\}.$$

The i th component of $X(t; \epsilon, N)$ is the number of individuals that play strategy i during $[t, t + 1)$. The parameter $\epsilon \geq 0$ corresponds to the size of the mutation rates; $\{X(t; 0, N)\}$ is the no-mutation process. We denote the transition probabilities by $p_{\epsilon, N}(x, x')$.

Assumption 1 *For every $N \geq 2$, the no-mutation transition probabilities are such that*

- (i) *at every time t , at most one agent changes his strategy,*
- (ii) *absent strategies will never be re-introduced,*
- (iii) *for every strategy that is currently played, there is a positive probability that its frequency increases in the next step, unless all members of the population use the same strategy.*

Thus, in the absence of mutations, every state is transient except for the pure states $s_1 = (N, 0, \dots, 0), \dots, s_n = (0, \dots, 0, N)$, which are absorbing.⁴ For every pair of different strategies i, j let $s_{i/j}$ denote the state where every agent plays i except for one, who plays j . The next assumption specifies how the mutations modify the no-mutation process.

Assumption 2 (i) *If $\epsilon > 0$, then $\{X(t; \epsilon, N)\}$ is irreducible.*

(ii) *The transition probabilities $p_{\epsilon, N}(x, x')$ depend continuously on ϵ .*

(iii) *For every $i \neq j$, the limit*

$$\lim_{\epsilon \rightarrow 0} \frac{p_{\epsilon, N}(s_i, s_{i/j})}{\epsilon} = \mu_{ij}$$

exists and does not depend on N .

(iv) *The matrix $(\mu_{ij})_{i,j=1}^n$, where $\mu_{11} = \dots = \mu_{nn} = 0$, is irreducible.*

⁴Part (ii) of this assumption rules out processes where adjusting agents play the best response (or even a smoothed best response) to the current state, as in Benaïm and Hirsch [1], Fudenberg and Kreps [13], Sandholm [23], and Young [25]. Theorem 2 of Fudenberg and Imhof [11] relaxed (ii). We do not know to what extent a similar extension is possible here.

(v) For every $N \geq 2$, every strategy i and all $x = (x_1, \dots, x_n) \in \mathcal{S}_N$ with $x_i \leq N - 2$,

$$\lim_{\epsilon \rightarrow 0} \frac{p_{\epsilon, N}(s_i, x)}{\epsilon} = 0.$$

Assumption 2 implies that each transition probability out of a pure state has a well-defined limiting order, with the probability that a single mutant invades a pure population being exactly of order ϵ or of order $o(\epsilon)$, while the probability that two or more mutants invade simultaneously is $o(\epsilon)$. Assumptions 1 and 2 imply that for every $N \geq 2$ and every sufficiently small $\epsilon > 0$, the process $\{X(t; \epsilon, N)\}$ spends most of the time at the pure states and transitions occur mainly along the edges of the state space where only two pure strategies are present in the population.⁵

We now place more restrictions on the behavior of the no-mutation process on these edges. For every pair of different strategies i, j let

$$r_{ij}(k, N) = p_{0, N} \left\{ \frac{k}{N} s_i + \left(1 - \frac{k}{N}\right) s_j, \frac{k+1}{N} s_i + \left(1 - \frac{k+1}{N}\right) s_j \right\},$$

$$k = 0, \dots, N-1,$$

$$\ell_{ij}(k, N) = p_{0, N} \left\{ \frac{k}{N} s_i + \left(1 - \frac{k}{N}\right) s_j, \frac{k-1}{N} s_i + \left(1 - \frac{k-1}{N}\right) s_j \right\},$$

$$k = 1, \dots, N.$$

Thus if k agents play strategy i and the rest play j , then $r_{ij}(k, N)$ is the probability that one agent switches from j to i , and $\ell_{ij}(k, N)$ is the probability that one agent switches from i to j . Clearly, $r_{ij}(k, N) = \ell_{ji}(N - k, N)$.

If the population is in state $x = (x_1, \dots, x_n) \in \mathcal{S}_N$ and $x_i \geq 1$, then, under random matching, the average payoff to agents that use strategy i is

$$u_i(x) = \frac{\sum_{j=1}^n a_{ij} x_j - a_{ii}}{N - 1},$$

where we assume that agents play infinitely⁶ often and do not interact with themselves⁷. The following assumption characterizes our imitation model by specifying how the transition probabilities may depend on the payoff functions $u_i(x)$. Let $u_{\min} = \min_{i,j} a_{ij}$ and $u_{\max} = \max_{i,j} a_{ij}$.

⁵See Young [25, Theorem 4] or Fudenberg and Imhof [11].

⁶We explore the consequences of randomness due to finitely many pairings per period in Ellison et al. [8]. Imhof and Nowak [17] consider a frequency-dependent Wright-Fisher process and compare its behavior when the game is played infinitely often to the behavior when the game is played just once in each period.

⁷All our results carry over to the case where each player may also play against himself, where $u_i(x)$ would be replaced by $\tilde{u}_i(x) = \sum_{j=1}^n a_{ij} x_j / N$. For simplicity, we restrict attention to the functions $u_i(x)$.

Assumption 3 (i) *There is a strictly positive Lipschitz continuous function f on $[u_{\min}, u_{\max}]^2$ such that for each pair of different strategies i, j ,*

$$\frac{r_{ij}(k, N)}{\ell_{ij}(k, N)} = f \left[u_i \left(\frac{k}{N} s_i + \frac{N-k}{N} s_j \right), u_j \left(\frac{k}{N} s_i + \frac{N-k}{N} s_j \right) \right],$$

$k = 1, \dots, N-1$ and $N \geq 2$.

(ii) *For all $u, v \in [u_{\min}, u_{\max}]$,*

$$f(u, v) \begin{matrix} > \\ < \end{matrix} 1 \iff u \begin{matrix} > \\ < \end{matrix} v.$$

Assumption 3 (i) says that along an edge the relative probabilities of upwards and downwards shifts in the number of agents playing i depends only on the current payoffs, and that this dependence has a well-behaved limit as N grows large. Assumption 3 (ii) adds the condition that the expected motion is in the direction of the better response. By "monotone imitation dynamics" we will mean any dynamics that meet Assumptions 1, 2, and 3.

3 Examples

We now present several examples of models that meet our assumptions. In each case, we describe only the no-mutation process and assume that mutations are modeled in agreement with Assumption 2.

Example 1 *In the aspiration and imitation model of Binmore and Samuelson [3], $u_i(x)$ is considered as an expected payoff and the realized payoff to any individual is given by the expected payoff plus a random variable that captures aggregate shocks. These random perturbations are assumed to be identically distributed and independent across all players and rounds. Let F denote the common distribution function. At each time step, a randomly chosen player compares his realized payoff to a deterministic aspiration level Δ . If the realized payoff is above that level, he keeps his strategy. Otherwise he imitates the strategy of a randomly chosen individual. Thus*

$$\begin{aligned} r_{ij}(k, N) &= \frac{k(N-K)}{N(N-1)} F \left[\Delta - u_j \left(\frac{k}{N} s_i + \frac{N-k}{N} s_j \right) \right], \\ \ell_{ij}(k, N) &= \frac{k(N-K)}{N(N-1)} F \left[\Delta - u_i \left(\frac{k}{N} s_i + \frac{N-k}{N} s_j \right) \right]. \end{aligned}$$

This model satisfies our assumptions with

$$f(u, v) = \frac{F(\Delta - v)}{F(\Delta - u)},$$

provided that F is strictly increasing and Lipschitz continuous and that $F(\Delta - u_{\max}) > 0$. The inequality means that there is a positive probability that the shock is large enough that even a strategy whose expected payoff is $u_{\max}$ can lead to a realized payoff below the aspiration level. $\square$

Example 2 Suppose that at each time step, an agent A_1 is chosen at random to re-evaluate his strategy. To this end, he chooses randomly another individual, A_2 . If A_2 uses the same strategy as A_1 , then A_1 keeps his strategy. If they use different strategies, say i_1 and i_2 , then A_1 imitates the strategy of A_2 with a probability that depends on their respective payoffs. Let this probability be given by $g[u_{i_2}(x), u_{i_1}(x)]$. With probability $1 - g[u_{i_2}(x), u_{i_1}(x)]$, A_1 keeps his strategy. This model satisfies our assumptions with $f(u, v) = g(u, v)/g(v, u)$, provided g is positive and Lipschitz continuous and $g(u, v) > g(v, u)$ if and only if $u > v$.

In [2], Benaïm and Weibull consider a model of aspiration and imitation of success, where the individual that reviews his strategy switches to the strategy of another individual if the payoff difference exceeds a random threshold. Denoting the distribution function of the threshold by F , we can cover this scheme by choosing $f(u, v) = F(u - v)/F(v - u)$.

If we choose

$$f(u, v) = \frac{\delta + u - v}{\delta + v - u},$$

where $\delta > u_{\max} - u_{\min}$, we obtain the evolutionary process with the local updating rule introduced by Traulsen et al. [24].

In Björnerstedt and Weibull's [5] success-oriented imitation dynamics, the probability that the reviewing agent A_1 imitates the strategy of A_2 is a strictly increasing function h , say, of A_2 's payoff, but does not depend on A_1 's own payoff. A finite population version of their model can be obtained from ours by setting $f(u, v) = h(u)/h(v)$.

An alternative to selecting the agents with each member of the population having of the same chance of being chosen is to choose a pair of strategies with each strategy currently present having the same probability of being chosen disregarding the number of agents that use it. This does not affect the ratio $r_{ij}(k, N)/\ell_{ij}(k, N)$ and leads to the same function f . $\square$

Example 3 In the frequency-dependent Moran process introduced by Nowak et al. [20], the fitness of an individual using strategy i in a population in state x is given by $\phi_i(x) = 1 - w + wu_i(x)$, where $w \in (0, 1]$ is a parameter that describes the intensity of selection. At each time step, one individual is chosen to reproduce and the probability that an individual using strategy i is chosen is $x_i\phi_i(x)/(\sum_j x_j\phi_j(x))$. The offspring then replaces a randomly

chosen individual. This model satisfies our assumptions with

$$f(u, v) = \frac{1 - w + wu}{1 - w + wv}. \quad \square$$

To see what our assumptions rule out, consider a model where at each time step, an agent is drawn at random and chooses a best reply to the current population. This model would in general not satisfy Assumption 1 (ii) because even without mutations absent strategies could be re-introduced. In addition, even if the agent is only allowed to choose the best reply among the currently present strategies, the functions r_{ij} and ℓ_{ij} would in general fail to be positive, so that Assumption 1 (iii) and Assumption 3 (i) would not be satisfied. Assumption 1 (iii) can be satisfied by using smooth best replies as in Fudenberg and Kreps [13] and Benaïm and Weibull [2], but even with smooth best replies Assumption 3 (i) can still fail. To see this, consider the neutral 2×2 case, where either strategy is a best reply and thus chosen with probability $\frac{1}{2}$. Then $r_{12}(k, N) = (N - k)/(2N)$ and $\ell_{12}(k, N) = k/(2N)$ for $k = 1, \dots, N - 1$, so that the ratio $r_{12}(k, N)/\ell_{12}(k, N)$ would be unbounded.

However, the following variant of a smooth best reply dynamics is consistent with our assumptions.

Example 4 Suppose that at each time step, one of the currently present strategies, $i_1, \dots, i_k$ say, is chosen at random and then one of the agents that use this strategy changes his strategy as follows. He makes noisy observations on the current average payoffs to $i_1, \dots, i_k$ and chooses a strategy with the highest observed payoff. Specifically, we assume there are independent random variables $\epsilon_{i_1}, \dots, \epsilon_{i_k}$ with a common continuous distribution, so that the observed payoffs are $u_{i_1}(x) + \epsilon_{i_1}, \dots, u_{i_k}(x) + \epsilon_{i_k}$. This gives rise to an imitation model with $f(u, v) = G(u - v)/G(v - u)$, where G is the distribution function of $\epsilon_{i_1} - \epsilon_{i_2}$, and we assume that G is Lipschitz continuous and that $G(u_{\min} - u_{\max}) > 0$. $\square$

The musical chairs model of Binmore, Samuelson and Vaughan [4] is similar to ours but does not exactly fit into our framework. In their model, there is a positive probability that more than two agents change their strategies in one time step. In the limit they consider, this probability becomes negligible, and the remaining transition probabilities (more specifically, the relevant ratios) behave as required by our assumptions, provided that one considers the payoff matrix $\{a_{ij}\}$ obtained by transforming the payoff matrix of the given game by the appropriate "anti-fitness function."⁸

⁸Binmore, Samuelson and Vaughan [4, page 12] suppose that the transition probabilities are the result of pairwise competition in a "death game" whose payoffs depend on the payoffs of the original game in a particular way. Thus, the learning process operates on this transformed payoff matrix and not the original one.

4 Equilibrium selection in large populations

Suppose throughout that $\{X(t; \epsilon, N)\}$ satisfies Assumptions 1-3. By Assumption 2 (i), $\{X(t; \epsilon, N)\}$ has a unique invariant distribution $\pi(x; \epsilon, N)$, provided $\epsilon > 0$. As we noted earlier, it is easy to show that for every N ,

$$\lim_{\epsilon \rightarrow 0} \pi(x; \epsilon, N) = 0 \quad \text{for all } x \in \mathcal{S}_N \setminus \{s_1, \dots, s_n\}.$$

To determine the limits for $x = s_1, \dots, s_n$ consider for every pair of distinct strategies i, j the probability that the no-mutation process will be absorbed at s_j if initially $N-1$ agents play i and one agent plays j . Denote this absorption probability by $\rho_{ij}(N)$. Define the $n \times n$ matrix $\Lambda(N) = [\Lambda_{ij}(N)]$ by

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

Lemma 1 *For every $N \geq 2$, the limits*

$$\pi_i^*(N) = \lim_{\epsilon \rightarrow 0} \pi(s_i; \epsilon, N), \quad i = 1, \dots, n,$$

exist and are strictly positive. Moreover, $\boldsymbol{\pi}^(N) = (\pi_1^*(N), \dots, \pi_n^*(N))$ is the unique vector such that*

$$\boldsymbol{\pi}^*(N) \Lambda(N) = \boldsymbol{\pi}^*(N), \quad \pi_1^*(N) + \dots + \pi_n^*(N) = 1. \quad (1)$$

This lemma says that there is a unique "limit distribution" $\boldsymbol{\pi}^*$, and that it is the unique invariant distribution of the matrix Λ . The proof of the lemma is simply to verify that the model satisfies the assumptions of Theorem 1 of Fudenberg and Imhof [11]. See Section 7 for this and all other proofs. Intuitively, the result follows from the fact that the system spends almost all of its time on the edges of the state space: starting from a steady state of the no-mutation process, a single rare mutation puts the process somewhere on an edge; when the mutation rate is small, the process will be absorbed at one of the two relevant vertices before the next mutation occurs.⁹

We now turn to the asymptotic behavior of the fixation probabilities $\rho_{ij}(N)$ and $\rho_{ji}(N)$. The behavior depends on whether, in the subgame with pure strategies i and j , one strategy dominates the other, both pure strategies are equilibria, or there is a mixed-strategy equilibrium. We write $a_N \asymp b_N$ if the ratio a_N/b_N is bounded and bounded away from zero.

⁹Note that it is important here that we first send the mutation probability to 0 and then send the population size to infinity; with the other order of limits, the result need not be concentrated on the vertices; see for example the discussion of the hawk-dove game in [15].

Lemma 2 *Let $i, j \in \{1, \dots, n\}$ be two different pure strategies. Let*

$$\phi_{ij}(x) = \log f[xa_{ij} + (1-x)a_{ii}, xa_{jj} + (1-x)a_{ji}], \quad \psi_{ij}(\xi) = \int_0^\xi \phi_{ij}(x) dx.$$

If i dominates j , that is, if $a_{ii} \geq a_{ji}$ and $a_{ij} > a_{jj}$,

$$\rho_{ij}(N) \asymp e^{-N\psi_{ij}(1)}, \quad \rho_{ji}(N) \asymp 1.$$

If j dominates i , that is, if $a_{ii} < a_{ji}$ and $a_{ij} \leq a_{jj}$,

$$\rho_{ij}(N) \asymp 1, \quad \rho_{ji}(N) \asymp e^{N\psi_{ij}(1)}.$$

In the coordination case, that is, if $a_{ii} \geq a_{ji}$ and $a_{ij} \leq a_{jj}$ with at least one inequality being strict,

$$\rho_{ij}(N) \asymp \frac{\exp\{-N\psi_{ij}(\xi_{ij}^*)\}}{\sqrt{N}}, \quad \rho_{ji}(N) \asymp \frac{\exp\{-N[\psi_{ij}(\xi_{ij}^*) - \psi_{ij}(1)]\}}{\sqrt{N}},$$

where $\xi_{ij}^ = (a_{ii} - a_{ji})/(a_{ii} - a_{ji} + a_{jj} - a_{ij})$.*

In the hawk-dove case, that is, if $a_{ii} < a_{ji}$ and $a_{ij} > a_{jj}$,

$$\rho_{ij}(N) \asymp \exp\{-N \max[0, \psi_{ij}(1)]\}, \quad \rho_{ji}(N) \asymp \exp\{N \min[0, \psi_{ij}(1)]\}.$$

If i and j are neutral, that is, if $a_{ii} = a_{ji}$ and $a_{ij} = a_{jj}$,

$$\rho_{ij}(N) \asymp \frac{1}{N}, \quad \rho_{ji}(N) \asymp \frac{1}{N}. \quad (2)$$

Note that the cases in Lemma 2 are exhaustive. Note also that if $a_{ii} = a_{ji} = a_{jj}$, Assumption 3 implies that $r_{ij}(k, N) = \ell_{ij}(k, N)$ for every k . Thus on the edge from s_i to s_j , the no-mutation process is a martingale, and this yields

$$\rho_{ij}(N) = \rho_{ji}(N) = \frac{1}{N}, \quad (3)$$

improving (2) for this case.

The proof of Lemma 2 rests on the representation

$$\frac{1}{\rho_{ij}(N)} = \sum_{\nu=0}^{N-1} \exp \left\{ \sum_{k=1}^{\nu} \log \frac{r_{ij}(N-k, N)}{\ell_{ij}(N-k, N)} \right\}.$$

We first approximate the interior sum by an integral over ϕ_{ij} , which leads to $N\psi_{ij}$. We then approximate the remaining sum by $N \int_0^1 \exp\{N\psi_{ij}(\xi)\} d\xi$. The Laplace method ([6]) shows that the asymptotic behavior of this integral

is essentially given by $\exp\{N \max_{\xi} \psi_{ij}(\xi)\}$ and the value of the maximum must then be determined for each case.

Lemma 2 shows in particular that for every pair i, j with $j \neq i$, the limit

$$\beta_{ij} = - \lim_{N \rightarrow \infty} \frac{\log \rho_{ij}(N)}{N} \quad (4)$$

exists and is a non-negative number. If $\beta_{ij} > 0$, then β_{ij} is the exponential rate with which $\rho_{ij}(N)$ converges to 0. If $\beta_{ij} = 0$, $\rho_{ij}(N)$ may or may not converge to 0.

Our main result on equilibrium selection, Theorem 1, is based on the least cost graphs of Freidlin and Wentzell [9], where we approximate the costs of the edges of the graphs by the exponential rates. Although these rates contain less accurate information on the asymptotic behavior of the absorption probabilities than available from Lemma 2, they are often sufficient to determine a selected strategy. For example, the β_{ij} contain enough information to find the selected strategy for coordination games, see Corollary 3. For the prisoner's dilemma game considered in Example 5, it is even enough to know which β_{ij} are positive. On the other hand, in Example 6, the more precise asymptotics of Lemma 2 are needed. For convenience, we collect the exponential rates for all 2×2 games in the following corollary.

Corollary 1 *If i dominates j ($a_{ii} \geq a_{ji}$ and $a_{ij} > a_{jj}$), then $\beta_{ij} = \psi_{ij}(1) > 0$. In the coordination case ($a_{ii} > a_{ji}$ and $a_{ij} \leq a_{jj}$), $\beta_{ij} = \psi_{ij}(\xi_{ij}^*) > 0$. In particular, if i is a strict Nash equilibrium ($a_{ii} > a_{ji}$), $\beta_{ij} > 0$.*

In the hawk-dove case ($a_{ii} < a_{ji}$ and $a_{ij} > a_{jj}$), $\beta_{ij} = \max\{0, \psi_{ij}(1)\}$. In every other case, $\beta_{ij} = 0$.

To formulate our main result on equilibrium selection we need the concept of an i -graph, where i is a pure strategy. A graph consisting of arrows $j \rightarrow k$, where $j \in \{1, \dots, n\} \setminus \{i\}$, $k \in \{1, \dots, n\}$ and $j \neq k$, is called an i -graph if it satisfies the following conditions: (a) every $l \in \{1, \dots, n\} \setminus \{i\}$ is the initial point of exactly one arrow in the graph, and (b) for any $l \in \{1, \dots, n\} \setminus \{i\}$ there is a sequence of arrows in the graph that leads from l to i . Let $\mathcal{G}_i$ denote the set of i -graphs and let $\mathcal{G} = \mathcal{G}_1 \cup \dots \cup \mathcal{G}_n$.

Theorem 1 *Consider a monotone imitation dynamic (i.e. a model satisfying Assumptions 1, 2, and 3) with invariant distribution $\pi(x; \epsilon, N)$ and limit distribution $\pi_i^*(N) = \lim_{\epsilon \rightarrow 0} \pi(s_i; \epsilon, N)$. For all $j \neq k$ let β_{jk} denote the exponential rate defined by (4). For every graph $g \in \mathcal{G}$ let*

$$\gamma(g) = \begin{cases} \sum_{(j \rightarrow k) \in g} \beta_{jk} & \text{if } \mu_{jk} > 0 \text{ for all } (j \rightarrow k) \in g, \\ \infty & \text{otherwise,} \end{cases}$$

and let $\gamma^* = \min\{\gamma(g) : g \in \mathcal{G}\}$. If i is a strategy such that the minimum γ^* is not attained by any i -graph, then

$$\lim_{N \rightarrow \infty} \pi_i^*(N) = 0.$$

If there exists a graph $g^* \in \mathcal{G}_i$ such that $\gamma(g) > \gamma(g^*)$ for all $g \in \mathcal{G} \setminus \mathcal{G}_i$, then i will be selected, that is,

$$\lim_{N \rightarrow \infty} \pi_i^*(N) = 1.$$

The proof of these limit results builds on the fact that, by Lemma 1, $\pi^*(N)$ is the unique solution to (1). According to a well-known representation of solutions to systems of equations of this type, see Freidlin and Wentzell [9], $\pi_i^*(N)$ is proportional to a sum of certain products of the entries $\Lambda_{jk}(N)$, where the choice of the factors is given by the set of i -graphs. The limit behavior of $\pi^*(N)$ can thus be obtained from the asymptotic behavior of the $\Lambda_{jk}(N)$ as given in Lemma 2.

We now apply Theorem 1 to obtain explicit selection results under suitable conditions on the underlying game.

Corollary 2 *Suppose strategy i weakly dominates every strategy $j \neq i$ in the subgame with pure strategies i and j . Suppose further that $\mu_{ji} > 0$ for all $j \neq i$. Then strategy i will be selected: $\pi_i^*(N) \rightarrow 1$.*

As another immediate consequence of Theorem 1 we obtain the following selection result for coordination games, where every pure strategy is a strict Nash equilibrium. For Moran processes, the special case of 3×3 coordination games has already been considered by Fudenberg et al. [15].

Corollary 3 *Consider a monotone imitation dynamic for an $n \times n$ coordination game. For every pair of different strategies j, k let $\beta_{jk} = \psi_{jk}(\xi_{jk}^*)$. Let $\gamma_j = \min\{\gamma(g) : g \in \mathcal{G}_j\}$ and $\gamma^* = \min\{\gamma_1, \dots, \gamma_n\}$.*

If $\gamma_i > \gamma^$, then $\pi_i^*(N) \rightarrow 0$. If $\gamma_i = \gamma^*$, then i is not selected against, that is, $\liminf_{N \rightarrow \infty} \pi_i^*(N) > 0$. If i is the unique strategy with $\gamma_i = \gamma^*$, then $\pi_i^*(N) \rightarrow 1$.*

5 Applications to the prisoner's dilemma

Here we show how our general results can be applied to determine selected equilibria for a class of 4×4 games arising from the finitely or infinitely repeated prisoner's dilemma. In Example 5, we consider the finite case and obtain by a straightforward application of Theorem 1 and Corollary 1 that

defectors will be selected in large populations. Using the more precise asymptotic results of Lemma 2, we show in Example 6 that in the case of infinitely many rounds conditional cooperators are selected. Moreover it turns out that even with finitely many rounds, the population consists mainly of conditional cooperators if the population is of moderate size and the number of rounds is sufficiently large. We precisely state this observation in Example 7. A general theorem in the next section will yield a rigorous justification. A heuristic argument for a similar result in a simpler model was given in Imhof et al. [16].

Consider the prisoner's dilemma game with strategies 'cooperate' and 'defect' and payoff matrix

$$\begin{pmatrix} R & S \\ T & P \end{pmatrix}, \quad T > R > P > S > 0.$$

We study the repeated prisoner's dilemma game with an expected number of rounds m and focus on the following $n = 4$ strategies: 1: cooperate in each round (ALLC), 2: defect in each round (ALLD), 3: tit-for-tat (TFT) and 4: perfect-tit-for-tat (PTFT), which is also called win-stay, lose-shift. TFT cooperates in the first round and then does what the opponent did in the previous round, PTFT cooperates in the first round and then cooperates after receiving R or P in the previous round and defects otherwise.

For simplicity, we assume that all types of mutations are equally likely. Denote the small mutation limit of the invariant distribution by $(\pi_1^*(N, m), \dots, \pi_4^*(N, m))$, where $m \leq \infty$.

Example 5 *If the number of rounds m is finite, the payoff matrix is*

$$\begin{pmatrix} R & S & R & R \\ T & P & \frac{1}{m}\{T + (m-1)P\} & \frac{1}{m}\{\lfloor \frac{m}{2} \rfloor P + \lfloor \frac{m+1}{2} \rfloor T\} \\ R & \frac{1}{m}\{S + (m-1)P\} & R & R \\ R & \frac{1}{m}\{\lfloor \frac{m}{2} \rfloor P + \lfloor \frac{m+1}{2} \rfloor S\} & R & R \end{pmatrix},$$

where $\lfloor z \rfloor$ denotes the largest integer less than or equal to z . ALLC, TFT and PTFT are neutral against each other. Hence, by Corollary 1, $\beta_{31} = \beta_{43} = 0$. Also, $\beta_{12} = 0$. Since ALLD is a strict Nash equilibrium, $\beta_{21} > 0$, $\beta_{23} > 0$ and $\beta_{24} > 0$. Consider $g^* = \{(4 \rightarrow 3), (3 \rightarrow 1), (1 \rightarrow 2)\} \in \mathcal{G}_2$. We have $\gamma(g^*) = 0$, and for every $g \in \mathcal{G} \setminus \mathcal{G}_2$, there occurs one of the positive numbers β_{21} , β_{23} , β_{24} in the sum defining $\gamma(g)$, so that $\gamma(g) > 0 = \gamma(g^*)$. It follows from Theorem 1 that $\lim_{N \rightarrow \infty} \pi_2^*(N, m) = 1$. That is, ALLD is selected. $\square$

Example 6 For the prisoner's dilemma game with infinitely many rounds, we obtain the payoff matrix

$$A = \begin{pmatrix} R & S & R & R \\ T & P & P & \frac{1}{2}P + \frac{1}{2}T \\ R & P & R & R \\ R & \frac{1}{2}P + \frac{1}{2}S & R & R \end{pmatrix}.$$

(A complete derivation is given in Subsection 7.6; we report only the selection results here.) If PTFT is a strict Nash equilibrium against ALLD, that is, if $2R > P + T$, then $\pi_3(N, \infty) \rightarrow \frac{2}{3}$ and $\pi_4(N, \infty) \rightarrow \frac{1}{3}$. That is, for large populations, a monotone imitation dynamics spends nearly all the time at the states where everyone plays TFT or where everyone plays PTFT, and TFT will be observed about twice as often as PTFT. In this case the relevant least cost graphs represent the transitions $\{(PTFT \rightarrow ALLC \rightarrow ALLD \rightarrow TFT)\}$, $\{(ALLC \rightarrow ALLD \rightarrow TFT), (PTFT \rightarrow TFT)\}$, which correspond to TFT-graphs, and $\{(ALLC \rightarrow ALLD \rightarrow TFT \rightarrow PTFT)\}$, which corresponds to a PTFT-graph. If PTFT is not a strict Nash equilibrium against ALLD, then $\pi_3(N, \infty) \rightarrow 1$, so that most of the time everyone plays TFT. Here the graph specified in the least cost argument represents the transitions $\{(ALLC \rightarrow ALLD \rightarrow TFT), (PTFT \rightarrow ALLD)\}$. $\square$

The following results for the case of finitely many rounds are a consequence of Theorem 2 in the next section.

Example 7 Consider again the case where $m < \infty$. Suppose first that PTFT is a strict Nash equilibrium against ALLD. Let $\eta > 0$. Then there exists a population size N_0 such that for every $N_1 > N_0$, there is a number of rounds m_0 so that

$$\left| \pi_3^*(N, m) - \frac{2}{3} \right| \leq \eta, \quad \left| \pi_4^*(N, m) - \frac{1}{3} \right| \leq \eta$$

for every $N \in \{N_0, \dots, N_1\}$ provided $m \geq m_0$.

If PTFT is not a strict Nash equilibrium against ALLD, then for every $\eta > 0$, there exists N_0 such that for every $N_1 > N_0$, there is m_0 so that

$$\pi_3^*(N, m) \geq 1 - \eta$$

for every $N \in \{N_0, \dots, N_1\}$ provided $m \geq m_0$. $\square$

For an intuitive explanation consider the dynamics in the case where $P + T \geq 2R$. We denote by ALLC (respectively ALLD, ...) also the state where everyone plays ALLC (respectively ALLD, ...). ALLC, TFT and PTFT are

neutral, and so the evolution of a population where only these strategies occur is determined by random drift. ALLD dominates ALLC and PTFT. Thus as soon as the population is in the state ALLC or PTFT, an ALLD invader will quickly take over the whole population and the population is then unlikely to return directly to ALLC or PTFT. ALLD against TFT is a coordination game, provided m is not too small. Therefore, a population in state ALLD is to some extent resistant to invasion by TFT. However, if m is sufficiently large, the basin of attraction of ALLD (i.e. the part of the edge from TFT to ALLD with drift towards ALLD) is small and if the population is not too large, there will soon be enough TFT players in the population that can take over the population. ALLD invaders then have only a small chance of taking over again. Therefore, the time spent at ALLD is relatively short compared to the time spent at TFT, which explains the result in Example 7. On the other hand, for any fixed $m < \infty$, the basin of attraction of ALLD corresponds to a fixed proportion of TFT players necessary for a reasonable chance to take over. As N gets large, it becomes increasingly unlikely that enough TFT players occur. Thus for fixed m and N sufficiently large, the process spends most of the time at ALLD, as shown in Example 5.

6 Population size relative to other parameters

To motivate the results in this section, we first give a heuristic derivation of Example 7, which is based on a certain continuity argument rather than on payoff considerations as described at the end of the last section. Recall that if m , the number of rounds in the prisoner's dilemma is finite, then ALLD is selected in the sense that $\pi_2^*(N, m) \rightarrow 1$ as $N \rightarrow \infty$. Of course, this does not imply that $\pi_2^*(N, m)$ is close to 1 for any given N . Indeed, a continuity argument shows that for fixed N , $\pi_3^*(N, m)$ is close to $\pi_3^*(N, \infty)$ when m is sufficiently large. Now if $P + T \geq 2R$, then $\pi_3^*(N, \infty)$ has been shown to be close to 1 when $N \geq N_0$ for some N_0 . Thus, for this N_0 and in fact for a whole range of values $N = N_0, \dots, N_1$, say, $\pi_3^*(N, m)$ is close to 1 when m is sufficiently large. That is, the behavior of large populations in the limit case $m = \infty$ is closely related to the behavior of medium-sized populations in the finite case $m < \infty$ when sufficiently many rounds are played.

This type of relationship between the behavior of medium-sized populations playing one game and that of a large population playing another game holds in more general situations. We first describe a suitable setting with quite general imitation processes and will then establish our result on populations of medium size within this framework.

Consider a family of homogeneous Markov chains $\{X(t; \epsilon, N, \delta) : t = 0, 1, \dots\}$ with state space $\mathcal{S}_N$ and transition probabilities $p_{\epsilon, N, \delta}(x, x')$. Again, $\epsilon \geq 0$ determines mutation rates, and $\delta \geq 0$ is a further parameter of the family. For example, in an application to the prisoner's dilemma, $\delta = 1/m$. We are interested in the behavior of the chain for very small mutation rates, moderate population sizes and small but positive values of δ .

Recall that $s_1, \dots, s_n$ denote the pure states in $\mathcal{S}_N$, and for $j \neq i$, $s_{i/j}$ denotes the state where every agent plays i except for one, who plays j . Let $\rho_{ij}(N, \delta)$ denote the probability that the no-mutation process $\{X(t; 0, N, \delta)\}$ will be absorbed at s_j if the initial state is $s_{i/j}$.

Assumption 4 *For every $N \geq 2$ and $\delta \geq 0$, under the no-mutation process,*

- (i) *a strategy that is absent remains absent,*
- (ii) *all states in $\mathcal{S}_N \setminus \{s_1, \dots, s_n\}$ are transient,*
- (iii) *the fixation probabilities $\rho_{ij}(N, \delta)$ are positive.*

Assumption 5 *For every $N \geq 2$,*

- (i) *$\{X(t; \epsilon, N, \delta)\}$ is irreducible if $\epsilon > 0$,*
- (ii) *for every $\delta \geq 0$, the transition probabilities $p_{\epsilon, N, \delta}(x, x')$ depend continuously on ϵ ,*
- (iii) *the transition probabilities of the no-mutation process are continuous in δ at $\delta = 0$,*
- (iv) *for $\delta \geq 0$ and $i \neq j$,*

$$\lim_{\epsilon \rightarrow 0} \frac{p_{\epsilon, N, \delta}(s_i, s_{i/j})}{\epsilon} = \mu_{ij} > 0,$$

- (v) *for $\delta \geq 0$, $i = 1, \dots, n$ and $s \in \mathcal{S}_N \setminus \{s_i, s_{i/1}, \dots, s_{i/i-1}, s_{i/i+1}, \dots, s_{i/n}\}$,*

$$\lim_{\epsilon \rightarrow 0} \frac{p_{\epsilon, N, \delta}(s_i, s)}{\epsilon} = 0.$$

These assumptions ensure that for every $\epsilon > 0$, $N \geq 2$ and $\delta \geq 0$, the process $\{X(t; \epsilon, N, \delta)\}$ has a unique invariant distribution $\pi(x; \epsilon, N, \delta)$ and that the limits

$$\pi_i^*(N, \delta) = \lim_{\epsilon \rightarrow 0} \pi(s_i; \epsilon, N, \delta), \quad i = 1, \dots, n.$$

exist, see Fudenberg and Imhof [11, Theorem 1].

Theorem 2 *Suppose that Assumptions 4 and 5 are satisfied and that for $\delta = 0$ the large population limit of the limit distribution exists:*

$$\lambda_i^* = \lim_{N \rightarrow \infty} \pi_i^*(N, 0), \quad i = 1, \dots, n. \quad (5)$$

Then, given $\eta > 0$, there exists a population size N_0 such that the following holds. For every $N_1 > N_0$, there exists $\delta_0 > 0$ such that

$$|\pi_i^*(N, \delta) - \lambda_i^*| < \eta, \quad i = 1, \dots, n, \quad (6)$$

for all $0 < \delta < \delta_0$ and $N \in \{N_0, \dots, N_1\}$.

The main step in the proof of Theorem 2 consists in showing that, under the present assumptions, for every N , $\pi^*(N, \delta) \rightarrow \pi^*(N, 0)$ as $\delta \rightarrow 0$. This amounts to showing that one may interchange the order of limits for $\epsilon \rightarrow 0$ and $\delta \rightarrow 0$. Note that $\pi^*(N, \delta)$ is not the unique invariant distribution of $\{X(t; 0, N, \delta)\}$. The proof is completed by exploiting (5). The reason that (6) can be guaranteed only for N in a suitable range and not for all N sufficiently large, is that we do not allow δ_0 to depend on N .

To prove the assertion made in Example 7 apply Theorem 2 with $\delta = 1/m$, so that the limit case $\delta = 0$ corresponds to the case of infinitely many rounds. The large population limits for this case were derived in Example 6.

7 Proofs

7.1 Proof of Lemma 1

The present Assumptions 1 and 2 ensure that Assumptions 1-4 of Fudenberg and Imhof [11] are satisfied. Assumption 1 also implies that $\rho_{ij}(N) > 0$ for all $i \neq j$. By Assumption 2 (iv), the matrix (μ_{ij}) is irreducible, and it follows that $\Lambda(N)$ is irreducible as well. Therefore there exists a unique solution $\pi^*(N)$ of (1), and $\pi^*(N)$ is positive. In particular, Assumption 5 of [11] is satisfied. The limit assertion now follows from [11, Theorem 1]. $\square$

7.2 Proof of Lemma 2

Fix any two different strategies i and j . We will only prove the assertions for $\rho_{ij}(N)$. By Karlin and Taylor [19, page 113] and Assumption 3 (i),

$$\frac{1}{\rho_{ij}(N)} = \sum_{\nu=0}^{N-1} \prod_{k=1}^{\nu} \frac{r_{ij}(N-k, N)}{\ell_{ij}(N-k, N)} = \sum_{\nu=0}^{N-1} \prod_{k=1}^{\nu} f[u(k, N), v(k, N)],$$

where

$$u(k, N) = \frac{(N - k - 1)a_{ii} + ka_{ij}}{N - 1}, \quad v(k, N) = \frac{(N - k)a_{ji} + (k - 1)a_{jj}}{N - 1}$$

and the empty product is equal to 1. Assumption 3 (i) implies that $\log f$ is Lipschitz continuous. Since

$$u(k, N) - \frac{(N - k)a_{ii} + ka_{ij}}{N} = O\left(\frac{1}{N}\right), \quad v(k, N) - \frac{(N - k)a_{ji} + ka_{jj}}{N} = O\left(\frac{1}{N}\right),$$

it follows that for some constant c_1 ,

$$\phi_{ij}\left(\frac{k}{N}\right) - \frac{c_1}{N} \leq \log f[u(k, N), v(k, N)] \leq \phi_{ij}\left(\frac{k}{N}\right) + \frac{c_1}{N}$$

for all k and N . Hence for $\nu = 0, \dots, N - 1$,

$$\exp\left\{-c_1 + \sum_{k=1}^{\nu} \phi_{ij}\left(\frac{k}{N}\right)\right\} \leq \prod_{k=1}^{\nu} f[u(k, N), v(k, N)] \leq \exp\left\{c_1 + \sum_{k=1}^{\nu} \phi_{ij}\left(\frac{k}{N}\right)\right\}.$$

Lipschitz continuity of $\log f$ implies that ϕ_{ij} is Lipschitz continuous with Lipschitz constant c_2 , say. Thus

$$\left|\sum_{k=1}^{\nu} \phi_{ij}\left(\frac{k}{N}\right) - N\psi_{ij}\left(\frac{\nu}{N}\right)\right| \leq N \sum_{k=1}^{\nu} \int_{\frac{k-1}{N}}^{\frac{k}{N}} \left|\phi_{ij}\left(\frac{k}{N}\right) - \phi_{ij}(x)\right| dx \leq \frac{c_2\nu}{2N}.$$

Consequently,

$$e^{-c_1 - c_2} \sum_{\nu=0}^{N-1} \exp\left\{N\psi_{ij}\left(\frac{\nu}{N}\right)\right\} \leq \frac{1}{\rho_{ij}(N)} \leq e^{c_1 + c_2} \sum_{\nu=0}^{N-1} \exp\left\{N\psi_{ij}\left(\frac{\nu}{N}\right)\right\}.$$

For every N ,

$$\begin{aligned} & \sum_{\nu=0}^{N-1} \exp\left\{N\psi_{ij}\left(\frac{\nu}{N}\right)\right\} \\ &= N \sum_{\nu=0}^{N-1} \int_{\frac{\nu}{N}}^{\frac{\nu+1}{N}} \exp\left\{N\left[\psi_{ij}\left(\frac{\nu}{N}\right) - \psi_{ij}(\xi)\right]\right\} \exp\{N\psi_{ij}(\xi)\} d\xi. \end{aligned}$$

If $\nu/N \leq \xi \leq (\nu + 1)/N$, then, by the mean value theorem,

$$N\left|\psi_{ij}\left(\frac{\nu}{N}\right) - \psi_{ij}(\xi)\right| \leq N\left|\frac{\nu}{N} - \xi\right| \|\psi'_{ij}\| \leq \|\psi'_{ij}\| = \|\phi_{ij}\| < \infty,$$

where $\|\cdot\|$ denotes the sup-norm on $[0, 1]$. It follows that

$$\frac{1}{\rho_{ij}(N)} \asymp N \int_0^1 \exp\{N\psi_{ij}(\xi)\} d\xi.$$

To complete the proof the asymptotic behavior of $\int_0^1 \exp\{N\psi_{ij}(\xi)\} d\xi$ has be determined. If $a_{ii} = a_{ji}$ and $a_{ij} = a_{jj}$, then by Assumption 3 (ii), $\psi_{ij}(\xi) = 0$ for all $\xi \in [0, 1]$, and so $\rho_{ij}(N) \asymp 1/N$. If $a_{ii} \geq a_{ji}$ and $a_{ij} > a_{jj}$, then $\phi_{ij}(x) > 0$ for all $x \in (0, 1]$. Thus $\psi_{ij}(\xi)$ attains its unique maximum over $[0, 1]$ at $\xi = 1$ and $\psi'_{ij}(1) = \phi_{ij}(1) > 0$. The Laplace method for integrals (see e.g. [6, Chapter 4]) now yields that

$$\int_0^1 \exp\{N\psi_{ij}(\xi)\} d\xi \asymp \frac{\exp\{N\psi_{ij}(1)\}}{N}$$

and it follows that $\rho_{ij}(N) \asymp \exp\{-N\psi_{ij}(1)\}$.

In the coordination case, $\psi_{ij}(\xi)$ attains its unique maximum at $x = \xi_{ij}^*$ and $\psi'_{ij}(\xi_{ij}^*) = 0$. In this case, the Laplace method yields

$$\int_0^1 \exp\{N\psi_{ij}(\xi)\} d\xi \asymp \frac{\exp\{N\psi_{ij}(\xi_{ij}^*)\}}{\sqrt{N}},$$

so that $\rho_{ij}(N) \asymp \exp\{-N\psi_{ij}(\xi_{ij}^*)\}/\sqrt{N}$.

The arguments for the remaining cases are similar and are therefore omitted. $\square$

7.3 Proof of Theorem 1

The limit distribution $(\pi_1^*(N), \dots, \pi_n^*(N))$ is determined by (1) in Lemma 1 and the matrix $\Lambda(N)$ is irreducible. It follows from Freidlin and Wentzell [9, Lemma 3.1, page 177] that the limit distribution can be expressed as

$$\pi_i^*(N) = \frac{\sum_{g \in \mathcal{G}'_i} \prod_{(j \rightarrow k) \in g} \Lambda_{jk}(N)}{\sum_{g \in \mathcal{G}'} \prod_{(j \rightarrow k) \in g} \Lambda_{jk}(N)}, \quad i = 1, \dots, n,$$

where $\mathcal{G}'_i = \{g \in \mathcal{G}_i : \gamma(g) < \infty\}$ and $\mathcal{G}' = \mathcal{G}'_1 \cup \dots \cup \mathcal{G}'_n$. If $g \in \mathcal{G}'_i$, then by Lemma 2,

$$\prod_{(j \rightarrow k) \in g} \Lambda_{jk}(N) \asymp \prod_{(j \rightarrow k) \in g} \rho_{jk}(N) \asymp \frac{\exp\left\{-N \sum_{(j \rightarrow k) \in g} \beta_{jk}\right\}}{N^{d(g)}} = \frac{\exp\{-N\gamma(g)\}}{N^{d(g)}},$$

where $d(g) \in \{0, \frac{1}{2}, 1, \frac{3}{2}, \dots, n-1\}$. The assertion is now easily verified. $\square$

7.4 Proof of Corollary 2

For every $j \neq i$, $a_{ii} \geq a_{ji}$ and $a_{ij} \geq a_{jj}$ with at least one of the inequalities being strict. Therefore, by Corollary 1, $\beta_{ji} = 0$ and $\beta_{ij} > 0$. Let g^* denote the i -graph that consists of all the arrows $j \rightarrow i$ with $j \neq i$. Then $\gamma(g^*) = 0$. Every graph $g \in \mathcal{G} \setminus \mathcal{G}_i$ contains one of the arrows $i \rightarrow j$ with $j \neq i$, and so $\gamma(g) > 0$. The assertion now follows from Theorem 1. $\square$

7.5 Proof of Corollary 3

For coordination games, $\rho_{jk}(N) \asymp e^{-N\beta_{jk}}/\sqrt{N}$ by Lemma 2. Thus, in the proof of Theorem 1, $d(g) = (n-1)/2$ for all $g \in \mathcal{G}'$, and it follows that

$$\pi_i^*(N) \asymp \frac{\sum_{g \in \mathcal{G}'_i} \exp\{N[\gamma^* - \gamma(g)]\}}{\sum_{g \in \mathcal{G}'} \exp\{N[\gamma^* - \gamma(g)]\}} \asymp \exp\{N[\gamma^* - \gamma_i]\}.$$

The limit assertions are now obvious. $\square$

7.6 A selection result for a class of 4×4 games

The following lemma yields the selection results indicated in Example 6. The class of games considered here is slightly more general and the μ_{ij} need not coincide.

Lemma 3 *Consider a 4×4 game that satisfies the following four conditions.*
a) Strategies 1, 3 and 4 are neutral against each other:

$$a_{11} = a_{13} = a_{31} = a_{33} = a_{14} = a_{41} = a_{44} = a_{34} = a_{43}. \quad (7)$$

b) In the subgame with strategies 1 and 2, strategy 1 is strictly dominated by strategy 2:

$$a_{11} < a_{21}, \quad a_{12} < a_{22}. \quad (8)$$

c) In the subgame with strategies 2 and 3, strategy 3 is a strict Nash equilibrium that weakly dominates strategy 2:

$$a_{33} > a_{23}, \quad a_{32} \geq a_{22}. \quad (9)$$

d) In the subgame with strategies 2 and 4, strategy 2 is a strict Nash equilibrium:

$$a_{22} > a_{42}. \quad (10)$$

Then we have the following equilibrium selection result for monotone imitation dynamics with $\mu_{jk} > 0$ for all $j \neq k$.

If in the subgame with strategies 2 and 4, strategy 4 is a strict Nash equilibrium, that is, if

$$a_{44} > a_{24}, \quad (11)$$

then

$$\lim_{N \rightarrow \infty} \boldsymbol{\pi}^*(N) = \left(0, 0, \frac{\mu_{41} + \mu_{43}}{\mu_{34} + \mu_{41} + \mu_{43}}, \frac{\mu_{34}}{\mu_{34} + \mu_{41} + \mu_{43}}\right). \quad (12)$$

If in that subgame strategy 4 is weakly dominated by strategy 2, that is, if

$$a_{44} \leq a_{24}, \quad (13)$$

then

$$\lim_{N \rightarrow \infty} \boldsymbol{\pi}^*(N) = (0, 0, 1, 0). \quad (14)$$

Proof. We rely on the representation (see Freidlin and Wentzell [9, Lemma 3.1, page 177] and Lemma 1)

$$\pi_i^*(N) = \frac{\sum_{g \in \mathcal{G}_i} w(g, N)}{\sum_{g \in \mathcal{G}} w(g, N)}, \quad i = 1, \dots, 4, \quad (15)$$

where for every graph g , $w(g, N) = \prod_{(j \rightarrow k) \in g} \mu_{jk} \rho_{jk}(N)$. To derive the asymptotic behavior of the functions $w(g, N)$ for $N \rightarrow \infty$, we first gather asymptotic results for the $\rho_{jk}(N)$. It follows from (7) and (3) that

$$\rho_{13}(N) = \rho_{31}(N) = \rho_{14}(N) = \rho_{41}(N) = \rho_{34}(N) = \rho_{43}(N) = \frac{1}{N}. \quad (16)$$

By (8) and Lemma 2,

$$\rho_{12}(N) \asymp 1. \quad (17)$$

By (8) through (10) and Corollary 1, there exists $\beta > 0$ such that

$$\rho_{21}(N) = O(e^{-\beta N}), \quad \rho_{32}(N) = O(e^{-\beta N}), \quad \rho_{24}(N) = O(e^{-\beta N}). \quad (18)$$

By (9) and Lemma 2, $\rho_{23}(N) \asymp 1$ or $\rho_{23}(N) \asymp 1/\sqrt{N}$. Thus,

$$\lim_{N \rightarrow \infty} N \rho_{23}(N) = \infty. \quad (19)$$

Assume now that (11) holds. Then, by Corollary 1, there exists $\beta_1 > 0$ such that

$$\rho_{42}(N) = O(e^{-\beta_1 N}). \quad (20)$$

Consider the three graphs

$$\begin{aligned} g_3^* &= \{(1 \rightarrow 2), (2 \rightarrow 3), (4 \rightarrow 1)\} \in \mathcal{G}_3, \\ g_3^{**} &= \{(1 \rightarrow 2), (2 \rightarrow 3), (4 \rightarrow 3)\} \in \mathcal{G}_3, \\ g_4^* &= \{(1 \rightarrow 2), (2 \rightarrow 3), (3 \rightarrow 4)\} \in \mathcal{G}_4. \end{aligned}$$

Setting $h(N) = \mu_{12}\mu_{23}\rho_{12}(N)\rho_{23}(N)/N$, we have by (16),

$$w(g_3^*, N) = \mu_{41}h(N), \quad w(g_3^{**}, N) = \mu_{43}h(N), \quad w(g_4^*, N) = \mu_{34}h(N), \quad (21)$$

and by (17),

$$\frac{1}{h(N)} = O\left(\frac{N}{\rho_{23}(N)}\right).$$

Every graph $g \in \mathcal{G}_1$ is of the form $g = \{(2 \rightarrow i), (3 \rightarrow j), (4 \rightarrow k)\}$ for some $i, j, k \in \{1, \dots, 4\}$. Thus, by (16), (18), (19) and (20),

$$w(g, N) = O\left(\frac{\rho_{23}(N)}{N^2}\right),$$

and it follows that

$$\lim_{N \rightarrow \infty} \frac{w(g, N)}{h(N)} = 0 \quad \text{for all } g \in \mathcal{G}_1.$$

A similar argument shows that $w(g, N) = O(N^{-2})$ for every $g \in \mathcal{G}_2$, so that $w(g, N)/h(N) = O[(N\rho_{23}(N))^{-1}]$. Hence, by (19),

$$\lim_{N \rightarrow \infty} \frac{w(g, N)}{h(N)} = 0 \quad \text{for all } g \in \mathcal{G}_2.$$

If $g \in \mathcal{G}_3 \cup \mathcal{G}_4 \setminus \{g_3^*, g_3^{**}, g_4^*\}$, then $w(g, N)/h(N) = O(N^{-1})$. It has thus been shown that

$$\lim_{N \rightarrow \infty} \frac{w(g, N)}{h(N)} = 0 \quad \text{for all } g \in \mathcal{G} \setminus \{g_3^*, g_3^{**}, g_4^*\}.$$

The claimed limit assertion (12) is now a consequence of the representation (15) and (21).

If (13) holds, one may prove limit assertion (14) by showing that $\lim_{N \rightarrow \infty} w(g, N)/w(\bar{g}_3, N) = 0$ for all $g \in \mathcal{G}_1 \cup \mathcal{G}_2 \cup \mathcal{G}_4$, where $\bar{g}_3 = \{(1 \rightarrow 2), (2 \rightarrow 3), (4 \rightarrow 2)\} \in \mathcal{G}_3$. The details are omitted. $\square$

7.7 Proof of Theorem 2

We will first show that for the fixation probabilities of the no-mutation processes we have

$$\lim_{\delta \rightarrow 0} \rho_{ij}(N, \delta) = \rho_{ij}(N, 0). \quad (22)$$

Fix $N \geq 2$ and a strategy j . Let $h_\delta(x)$ denote the probability that $\{X(t; 0, N, \delta)\}$ will be absorbed at s_j if the initial state is x . Thus $\rho_{ij}(N, \delta) = h_\delta((1 - 1/N)s_i + (1/N)s_j)$. Write $\mathcal{S}' = \mathcal{S}_N \setminus \{s_1, \dots, s_n\}$. By Assumption 4, $s_1, \dots, s_n$

are absorbing states and the elements of $\mathcal{S}'$ are transient states of the no-mutation processes. Hence, for every $\delta \geq 0$, h_δ is the unique solution of

$$h_\delta(x) = p_{0,N,\delta}(x, s_j) + \sum_{y \in \mathcal{S}'} p_{0,N,\delta}(x, y) h_\delta(y), \quad x \in \mathcal{S}',$$

see e.g. Karlin and Taylor [19, page 90]. It therefore follows from Assumption 5 (iii) and Cramer's rule that $h_\delta \rightarrow h_0$ as $\delta \rightarrow 0$. This proves (22).

Assumptions 4 and 5 ensure that for every N and δ , Theorem 1 in Fudenberg and Imhof [11] can be applied to compute $\pi_i^*(N, \delta)$. Accordingly, $(\pi_1^*(N, \delta), \dots, \pi_n^*(N, \delta))$ is the unique invariant distribution of a transition matrix whose off-diagonal elements are $\mu_{ij}\rho_{ij}(N, \delta)$. Hence, using (22), we obtain

$$\lim_{\delta \rightarrow 0} \pi_i^*(N, \delta) = \pi_i^*(N, 0), \quad i = 1, \dots, n \text{ and } N \geq 2. \quad (23)$$

Let $\eta > 0$. By (5), there exists N_0 such that

$$|\lambda_i^* - \pi_i^*(N, 0)| < \frac{\eta}{2} \quad \text{for all } N \geq N_0 \text{ and } i = 1, \dots, n.$$

Choose any $N_1 > N_0$. In view of (23), there exists $\delta_0 = \delta_0(N_0, N_1) > 0$ such that

$$|\pi_i^*(N, \delta) - \pi_i^*(N, 0)| < \frac{\eta}{2} \quad \text{for all } \delta \in (0, \delta_0), \ N = N_0, \dots, N_1, \ i = 1, \dots, n.$$

This completes the proof. $\square$

Acknowledgments We thank an Associate Editor and a referee whose remarks led to a substantial improvement of the presentation of our ideas. We also thank Larry Samuelson for helpful comments, and Harvard's Program for Evolutionary Dynamics for its hospitality. NSF grant SES-0426199 provided financial support.

References

- [1] M. Benaïm, M. Hirsch, Mixed equilibria and dynamical systems arising from fictitious play in perturbed games, *Games Econ. Behav.* 29 (1999) 36-72.
- [2] M. Benaïm, J. W. Weibull, Deterministic approximation of stochastic evolution in games, *Econometrica* 71 (2003) 873-903.
- [3] K. Binmore, L. Samuelson, Muddling through: noisy equilibrium selection, *J. Econ. Theory* 74 (1997) 235-265.

- [4] K. Binmore, L. Samuelson, R. Vaughan, Musical chairs: modelling noisy evolution, *Games Econ. Behav.* 11 (1995) 1-35.
- [5] J. Björnerstedt and J. Weibull, Nash equilibrium and evolution by imitation. In: *The Rational Foundations of Economic Analysis*, ed. K. Arrow et al., McMillan, London 1995.
- [6] N. G. de Bruijn, *Asymptotic Methods in Analysis*, North-Holland, Amsterdam, 1958.
- [7] G. Ellison, D. Fudenberg, Word of mouth communication and social learning, *Quart. J. Econ.* 110 (1995) 93-126.
- [8] G. Ellison, D. Fudenberg, L. A. Imhof, Random matching in adaptive dynamics, *Mimeo*, 2006.
- [9] M. I. Freidlin, A. D. Wentzell, *Random Perturbations of Dynamical Systems*, second edition, Springer, New York, 1998.
- [10] D. Friedman, Evolutionary Games in Economics, *Econometrica* 59 (1991) 637-666.
- [11] D. Fudenberg, L. A. Imhof, Imitation processes with small mutations, *J. Econ. Theory* 131 (2006) 251-262.
- [12] D. Fudenberg, L. A. Imhof, Multi-dimensional muddling in large populations, *Mimeo*, 2006.
- [13] D. Fudenberg, D. Kreps, Learning mixed equilibria, *Games Econ. Behav.* 5 (1993) 320-367.
- [14] D. Fudenberg, D. Levine, *The Theory of Learning in Games*, MIT Press, Cambridge, MA, 1998.
- [15] D. Fudenberg, M. A. Nowak, C. Taylor, L. A. Imhof, Evolutionary game dynamics in finite populations with strong selection and weak mutation, *Mimeo*, Harvard University, 2005.
- [16] L. A. Imhof, D. Fudenberg, M. A. Nowak, Evolutionary cycles of cooperation and defection, *Proc. Natl. Acad. Sci. USA* 102 (2005) 10797-10800.
- [17] L. A. Imhof, M. A. Nowak, Evolutionary game dynamics in a Wright-Fisher process, *J. Math. Biol.* 52 (2006) 667-681.
- [18] M. Kandori, G. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, *Econometrica* 61 (1993) 29-56.

- [19] S. Karlin, H. M. Taylor, A First Course in Stochastic Processes, second edition, Academic Press, New York, 1975.
- [20] M. A. Nowak, A. Sasaki, C. Taylor, D. Fudenberg, Emergence of co-operation and evolutionary stability in finite populations, *Nature* 428 (2004) 646-650.
- [21] L. Samuelson, Evolutionary Games and Equilibrium Selection, MIT Press, Cambridge, MA, 1997.
- [22] L. Samuelson, J. Zhang, Evolutionary stability in asymmetric games, *J. Econ. Theory* 57 (1992) 363-391.
- [23] W. Sandholm, Simple and clever decision rules for a model of evolution, *Econ. Letters* 61 (1998) 165-170.
- [24] A. Traulsen, J. C. Claussen, C. Hauert, Coevolutionary dynamics: from finite to infinite populations, *Phys. Rev. Lett.* 95 (2005) 238701.
- [25] P. Young, The evolution of conventions, *Econometrica* 61 (1993) 57-84.