

Stochastic Imitation in Finite Games*

SSE/EFI Working Paper Series in Economics and Finance No 363

Jens Josephson

Department of Economics, Stockholm School of Economics, P.O. Box 6501,
S-113 83 Stockholm, SWEDEN. E-mail: jens.josephson@hhs.se.

Alexander Matros

ELSE, Department of Economics University College London, Gower Street,
London WC1E 6BT, UK. E-mail: a.matros@ucl.ac.uk.

Abstract

In this paper we model an evolutionary process with perpetual random shocks, where individuals sample population-specific strategy and payoff realizations and imitate the most successful behavior. For finite n -player games we prove that in the limit, as the perturbations tend to zero, only strategy-tuples in minimal sets closed under single better replies will be played with positive probability. If the strategy-tuples in one such minimal set have strictly higher payoffs than all outside strategy-tuples, then the strategy-tuples in this set will be played with probability one in the limit, provided the minimal set is a product set and the sample is sufficiently large.

*The authors are grateful for helpful comments and suggestions from Jörgen W. Weibull and an anonymous referee. We have also benefited from comments by Bo Becker, Ken Binmore, Eric van Damme, Maria Sáez-Martí, Philippe Solal, H. Peyton Young, and seminar participants at the Stockholm School of Economics 1999, the Young Economist Meeting 2000, and the First World Congress of the Game Theory Society 2000. We gratefully acknowledge financial support from the Jan Wallander and Tom Hedelius Foundation.

1 Introduction

In most game-theoretical models of learning, the individuals are assumed to know a great deal about the structure of the game, such as their own payoff function and all players' available strategies. However, for many applications, this assumption is neither reasonable nor necessary; in many cases, individuals may not even be aware that they are playing a game. Moreover, equilibrium play may be achieved even with individuals who have very little knowledge of the game, an observation made already in 1950 by John F. Nash. In his unpublished Ph.D. thesis (1950), he referred to it as "the 'mass-action' interpretation of equilibrium points." Under this interpretation:

"It is unnecessary that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal."

In the present paper, we develop a model in this spirit, where individuals are only required to know their own available pure strategies and a sample of the payoffs that a subset of these strategies have earned in the past. We use an evolutionary framework with perpetual random shocks similar to Young (1993), but our assumption of individual behavior is different. Whereas the individuals in his model play a myopic best reply to a sample distribution of their opponents' strategies, the individuals in our model imitate other individuals in their own population. Imitation is a behavior with both experimental, empirical, and theoretical support.¹

¹For experimental support of imitation, see for example, Huck *et al.* (1999, 2000) and Duffy and Feltovich (1999), for empirical support see Graham (1999), Wermers

More specifically, we assume that in every period, individuals are drawn at random from each of n arbitrary-size populations to play a finite n -player game. Each of these individuals observes a sample from a finite history of her population's past strategy and payoff realizations. Thereafter, she imitates by choosing the most attractive strategy in her sample. This could, for instance, be the strategy with the highest average payoff, or that with the highest maximum payoff. In the special case when each population consists of only one individual, this behavior can be interpreted as a special kind of reinforcement learning.² In this case, each individual plays the most successful strategy in a sample of her own previous strategy choices. With some small probability, the individuals also make errors or experiment, and instead choose any strategy at random from their set of strategies.³ Altogether, this results in an ergodic Markov process, which we denote *imitation play*, on the space of histories. We study the stationary distribution of this process as the experimentation probability tends to zero.

Imitation in a stochastic setting has previously been studied by Robson and Vega-Redondo (1996), who modify the framework of Kandori *et al.* (1993) to allow for random matching. More precisely, they assume that in each period, individuals are randomly matched for a finite

 (1999), and Griffiths *et al.* (1998), and for theoretical support, see Björnerstedt and Weibull (1996) and Schlag (1998, 1999).

²This behavior is related to one of the interpretations of individual behavior in Osborne and Rubinstein (1998), where each individual first samples each of her available strategies once and then chooses the strategy with the highest payoff realization.

³An alternative interpretation, which provides a plausible rationale for experimentation and is consistent with the knowledge of individuals in the model is the following: if and only if the sample does not contain all available strategies, then with a small probability, the individual instead picks a strategy not included in the sample at random.

number of rounds and tend to adopt the strategy with the highest average payoff across the population. Robson and Vega-Redondo (1996) assume either single- or two-population structures and obtain results for symmetric 2×2 games and two-player games of common interest.

Our model differs from these and other stochastic learning models, and has several advantages. First, we are able to prove general results, applicable to any finite n -player game, about the limiting distribution of imitation play. We are thus not restricted to the two classes of games in Robson and Vega-Redondo (1996), or even to a generic class of games, as in Young (1998). Second, we find that this distribution has some interesting properties. For instance, it puts probability one on an efficient set of outcomes in a large class of n -player games. Third, the speed of convergence of our process is relatively high. We show that in 2×2 Coordination Games, the expected first passage time may be considerably shorter than in Young (1993), Kandori *et al.* (1993), and Robson and Vega-Redondo (1996), for small experimentation probabilities.

The perturbed version of imitation play is a regular perturbed Markov process. This implies that the methods employed by Young (1993) can be used to calculate the states that will be played with positive probability by the stationary distribution of the process as the experimentation probability tends to zero, i.e. the *stochastically stable states*.

We prove three results which facilitate this calculation and enable us to characterize the set of such states. First, we show that from any initial state, the unperturbed version of imitation play converges to a state which is a repetition of a single pure-strategy profile, a *monomorphic state*. Hence, the stochastically stable states of the process belong to the set of monomorphic states.

Second, we prove for the perturbed process that in the limit, as the experimentation probability tends to zero, only pure-strategy profiles in particular subsets of the strategy-space are played with positive probability. These sets, which we denote *minimal sets closed under single better replies* (*minimal cusber sets*), are minimal sets of strategy profiles such that no player can obtain a weakly better payoff by deviating unilaterally and playing a strategy outside the set.

Minimal cusber sets are similar to Sobel's (1993) definition of *non-equilibrium evolutionary stable* (*NES*) sets for two-player games and to what Nöldeke and Samuelson (1993) call *locally stable components* in their analysis of extensive form games. They are also closely related to *minimal sets closed under better replies* (Ritzberger and Weibull, 1995). We show that every minimal set closed under better replies contains a minimal cusber set and that if a minimal cusber set is a product set, then it is also a minimal set closed under better replies. The relationship between minimal cusber set and the limiting distribution of imitation play should be contrasted with Hurkens's (1995) and Young's (1998) findings that adaptive best-reply processes for generic games selects pure-strategy profiles in minimal sets closed under best replies.

Finally, we show that in a certain class of games, imitation play selects efficient outcomes. If the pure-strategy profiles in a minimal cusber set have strictly higher payoffs than all other pure-strategy profiles, then the pure-strategy profiles in this set will be played with probability one in the limit, provided that the minimal cusber set is a product set. This is a generalization of previous results for games of common interest. Robson and Vega-Redondo (1996) prove that in their model a Pareto-dominant pure-strategy profile is selected in two-player games of common interest.

Applied to 2×2 games, our three results give clear predictions. In Coordination Games, imitation play selects the strictly Pareto-superior Nash equilibrium. This result differs sharply from the predictions in Young's (1993) and Kandori *et al.*'s (1993) models, where the stochastically stable states correspond to the risk-dominant equilibria, but it is consistent with the predictions of Robson and Vega-Redondo's (1996) model for symmetric Coordination Games. However, if neither equilibria Pareto dominates the other, the latter model may select the risk-dominant equilibrium, whereas both equilibria are played with positive probability in our model. In games without pure Nash equilibria, all four monomorphic states are stochastically stable.

The paper is organized as follows. In Section 2, we define the unperturbed and perturbed versions of imitation play. In Section 3, we derive general results for the limiting distribution of the process. In Section 4, we apply our results to 2×2 games and compare our findings to those in previous literature. In Section 5, we discuss an extension of the model and in Section 6, we conclude. Omitted proofs can be found in the Appendix.

2 Model

The model described below is similar to Young (1993), but the sampling procedure is modified and individuals employ a different decision-rule. Let Γ be a n -player game in strategic form. Let X_i be the finite set of pure strategies x_i available to player $i \in \{1, \dots, n\} = N$ and let $\Delta(X_i) = \{p \in \mathbb{R}_+^{|X_i|} \mid p' \mathbf{1} = 1\}$ be player i 's set of mixed strategies p_i . Define the product sets $X = \prod_i X_i$ and $\square(X) = \prod_i \Delta(X_i)$.

To each player role i in the game Γ corresponds a finite and non-empty population of individuals. These populations need not be of the

same size, nor need they be large. The population of individuals corresponding to player role i have payoffs represented by the utility function $\pi_i : X \rightarrow \Pi_i$, where $\Pi_i \subset \mathbb{R}$.⁴ Expected payoffs are represented by the function $u_i : \square(X) \rightarrow \mathbb{R}$. Note that we write “players” when referring to the game Γ and “individuals” when referring to the members of the populations.

The play proceeds as follows. Let $t = 1, 2, \dots$ denote successive time periods. The game Γ is played once every period. In period t , one individual is drawn at random from each of the n populations and assigned to play the corresponding role. An individual in role i chooses a pure strategy x_i^t from a subset of her strategy space X_i , according to a rule that will be defined below. The pure-strategy profile $x^t = (x_1^t, \dots, x_n^t)$ is recorded and referred to as *play* at time t . The *history* of plays up to time t is the sequence $h^t = (x^{t-m+1}, \dots, x^t)$, where m ($4 \leq m$) is a given positive integer, the *memory size* of all individuals.

Let h be an arbitrary history. Denote a sample of s ($1 < s \leq m/2$) elements from the m most recent strategy choices by individuals in population i by $w_i \in X_i^s$, and the corresponding payoff realizations by $v_i \in \Pi_i^s$. For any history h , the *maximum average correspondence*, $\alpha_i : X_i^s \times \Pi_i^s \rightarrow X_i$, maps each pair of strategy sample w_i and payoff sample v_i to the pure strategy (or the set of strategies) with the highest average payoff in the sample. Following Young (1993), we can think of the sampling process as beginning in period $t = m + 1$ from some arbitrary initial sequence of m plays h^m . In this period and every period thereafter, each individual in player role i inspects a pair (w_i, v_i) and plays a pure strategy $x_i \in \alpha_i(w_i, v_i)$. This defines a finite Markov

⁴Actually, utility functions need not be identical within each population for any of the results in this paper. It is sufficient if each individual’s utility function is a positive affine transformation of a population-specific utility function.

process $P^{\alpha, m, s, 0}$ on the finite state space $H = X^m$ of histories. Given a history $h^t = (x^{t-m+1}, \dots, x^t)$ at time t , the process moves to a state of the form $h^{t+1} = (x^{t-m+2}, \dots, x^t, x^{t+1})$ in the next period. Such a state is called a *successor* of h^t . We call the process $P^{\alpha, m, s, 0}$ *imitation play with memory m and sample size s* .

As an example, consider imitation play with memory $m = 6$ and sample size $s = 3$ in the 2×3 game in Figure 1.

	a	b	c
A	2, 2	0, 0	3, 3
B	0, 0	1, 1	0, 0

FIGURE 1

Let $h = ((A, a), (B, a), (A, b), (B, b), (A, c), (B, c))$ be the initial history. Assume that the individual in the role of the row player (player 1) draws the last three plays from this history, such that $w_1 = (B, A, B)$ and $v_1 = (1, 3, 0)$. This gives an average payoff of 3 to strategy A and $1/2$ to strategy B . Hence, the individual in role of the row player will choose strategy A in the next period. Further, assume that the individual in the role of the column role (player 2) draws the first three plays, such that $w_2 = (a, a, b)$ and $v_2 = (2, 0, 0)$. This gives an average payoff of 1 to strategy a and 0 to strategy b . Strategy c cannot be chosen since it is not included in the sample. Hence, the individual in the column role will choose strategy a in the next period. Altogether, this implies that the unperturbed process will move to state $h' = ((B, a), (A, b), (B, b), (A, c), (B, c), (A, a))$ in the next period.

The perturbed process can be defined as follows. In each period, there is some small probability $\varepsilon > 0$ that each individual i drawn to play chooses a pure strategy at random from X_i , instead of according to the imitation rule. The event that i experiments is assumed to be

independent of the event that j experiments for every $j \neq i$ and across time periods. The process defined in this way is denoted $P^{\alpha, m, s, \varepsilon}$ and it is referred to as *imitation play with memory m , sample size s , and experimentation probabilities ε* .

3 Stochastic Stability

In this section, we turn our attention to the limiting distribution of imitation play as the experimentation probability tends to zero. We first show that we can apply some of the tools in Young (1993) to calculate this distribution. Thereafter, we prove that it puts positive probability only on the pure-strategy profiles in particular subsets of the strategy-space.

3.1 Preliminaries

In what follows, we will make use of the following definitions. A *recurrent class* of the process $P^{\alpha, m, s, 0}$ is a set of states such that there is zero probability of moving from any state in the class to any state outside, and there is a positive probability of moving from any state in the class to any other state in the class. We call a state h' *absorbing* if it constitutes a singleton recurrent class and we refer to a state $h_x = (x, x, \dots, x)$, where x is any pure-strategy profile from X , as a *monomorphic state*. In other words, a monomorphic state is a state where the individuals in each player role played the same pure strategy in the last m periods. If each player i has $|X_i| \geq 1$ strategies in the game Γ , then there are $|X| = \prod_i |X_i|$ monomorphic states in this game. The following result shows that the monomorphic states correspond one to one with the recurrent classes of imitation play.

Theorem 1 *All monomorphic states are absorbing states of the unper-*

turbed process $P^{\alpha,m,s,0}$. From any state, the process converges almost surely to a monomorphic state.

PROOF: It is evident that any monomorphic state is an absorbing state, since any sample from a monomorphic state will contain one strategy only. We shall prove that if $s/m \leq 1/2$, then the monomorphic states are the only recurrent classes of the unperturbed process. Consider an arbitrary initial state $h^t = (x^{t-m+1}, \dots, x^t)$. If $s/m \leq 1/2$, there is a positive probability that all individuals drawn to play sample from x^{t-s+1}, \dots, x^t (i.e. i , for $i \in N$, sample $x_i^{t-s+1}, \dots, x_i^t$ and $\pi_i^{t-s+1}, \dots, \pi_i^t$) in every period from $t+1$ to $t+s$ inclusive. All of them play the pure strategy with the highest average payoff in their sample. Without loss of generality, assume that this is a unique pure strategy x_i^* for each of the player roles (if there is more than one pure strategy, all of them have positive probability according to the assumptions). With positive probability, all the individuals drawn to play thereafter sample only from plays more recent than x^t in every period from $t+s+1$ to $t+m$ inclusive. Since all of these samples have the form $w_i^* = (x_i^*, \dots, x_i^*)$ and $v_i^* = (\pi_i^*, \dots, \pi_i^*)$, the unique pure strategy with the highest payoff in the sample is x_i^* . Hence, there is a positive probability of at time $t+m$ obtaining a history $h^{t+m} = (x^*, \dots, x^*)$, a monomorphic state. It follows that for $s/m \leq 1/2$, the only recurrent classes of the unperturbed process are the monomorphic states. ■

By the same logic as in Young (1993), the perturbed process $P^{\alpha,m,s,\varepsilon}$ is a *regular perturbed Markov process*, and hence it has a unique stationary distribution μ^ε satisfying the equation $\mu^\varepsilon P^{\alpha,m,s,\varepsilon} = \mu^\varepsilon$. Moreover, by Theorem 4 in Young (1993) $\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon = \mu^0$ exists, and μ^0 is a stationary distribution of $P^{\alpha,m,s,0}$.

The following concepts are due to Freidlin and Wentzell (1984), Foster and Young (1990), and Young (1993). A state $h \in H$ is *stochastically stable* relative to the process $P^{\alpha, m, s, \varepsilon}$ if $\lim_{\varepsilon \rightarrow 0} \mu^\varepsilon(h) > 0$. Let h' be a successor of h and let x be the right-most element of h' . A *mistake* in the transition from h to h' is a component x_i of x that does not have the maximum average payoff in any sample of strategies and payoffs from h . For any two states h, h' , the *resistance*, r , is the total number of mistakes involved in the transition $h \rightarrow h'$ if h' is a successor of h , otherwise $r = \infty$. For each pair of distinct monomorphic states, an *xy-path* is a sequence of states $\zeta = (h_x, \dots, h_y)$ beginning in h_x and ending in h_y . The *resistance* of this path is the sum of the resistances on the edges that compose it. Let r_{xy} be the least resistance over all *xy*-paths. Construct a complete directed graph with $|X|$ vertices, one for each recurrent class. The weight on the directed edge $h_x \rightarrow h_y$ is r_{xy} . A *tree* rooted at h_x is a set of $|X| - 1$ directed edges such that, from every vertex different from h_x , there is a unique directed path in the tree to h_x . The *resistance of such a rooted tree* $\mathfrak{S}(x)$ is the sum of the resistances $r_{x'x''}$ on the $|X| - 1$ edges that compose it. The *stochastic potential* $\rho(x)$ of a monomorphic state h_x is the minimum resistance over all trees rooted at h_x .

The following theorem describes the long-run behavior of the perturbed process as the experimentation probability tends to zero.

Theorem 2 *The stochastically stable states of $P^{\alpha, m, s, \varepsilon}$ are the monomorphic states with minimum stochastic potential.*

PROOF: This follows from Theorem 1 above and Theorem 4 in Young (1993). ■

In order to characterize the sets of monomorphic states with minimum stochastic potential, the *better-reply correspondence* $\gamma = \prod_{i \in N} \gamma_i :$

$\square(X) \rightarrow X$ turns out to be useful. This concept, due to Ritzberger and Weibull (1995), is defined as follows:

$$\gamma_i(p) = \{x_i \in X_i \mid u_i(x_i, p_{-i}) \geq u_i(p)\}, \forall i \in N. \quad (1)$$

If $p \in \square(X)$ has support only on a single pure-strategy profile $x \in X$, we write $\gamma_i(x)$ instead of $\gamma_i(p)$ in order to simplify the notation.

The following lemma is helpful in calculating the stochastic potential of a monomorphic state.

Lemma 1 *The resistance from h_x to h_y is positive for any $y \neq x$. It is equal to one if and only if $y \in (\gamma_i(x), x_{-i})$ for some player i .*

PROOF: The first statement follows since a monomorphic state consists of a repetition of a single strategy profile and since only strategies included in the sample can be selected in the absence of perturbations. The resistance is equal to one if $y \in (\gamma_i(x), x_{-i})$ for some player i , since then h_y can be reached if an individual in player role i plays y_i by mistake, and $m - 1$ consecutive individuals in player role i thereafter draw a sample including strategy y_i . Moreover, it is clear that if $y \notin (\gamma_i(x), x_{-i})$ for all i , then a single mistake is insufficient to make any individual change strategy. Hence, the resistance of the transition from h_x to h_y is equal to one only if $y \in (\gamma_i(x), x_{-i})$ for some player i . ■

In order to illustrate how to calculate the stochastic potential under imitation play, we present an example of a two-player game. In the game in Figure 2, every player has three strategies, labeled A, B and C for the first player and a, b and c for the second player. The game has one strict Nash equilibrium (A, a) , where both players gain less than in a mixed equilibrium with the probability mixture $1/2$ on $B(b)$ and $1/2$ on $C(c)$ for the first (second) player.⁵

⁵There is also a third equilibrium, $((\frac{5}{7}, \frac{1}{7}, \frac{1}{7}), (\frac{3}{5}, \frac{1}{5}, \frac{1}{5}))$.

	a	b	c
A	1, 1	0, 0	0, 0
B	0, 0	3, 2	0, 3
C	0, 0	0, 3	3, 2

FIGURE 2

Denote by $x_1 \in \{A, B, C\}$ some strategy choice by player 1 and $x_2 \in \{a, b, c\}$ some strategy choice by player 2. To find the stochastically stable monomorphic states, construct directed graphs with nine vertices, one for each monomorphic state. In Figure 3, we illustrate two such trees. The numbers in the squares correspond to the resistances of the directed edges and the numbers in the circles represent the payoffs associated with the monomorphic states. It is easy to check that for $s > 2$, the stochastic potential $\rho(A, a) = 8$, whereas all other monomorphic states have a stochastic potential of 9. Hence, the monomorphic state $h_{(A,a)}$ is stochastically stable.

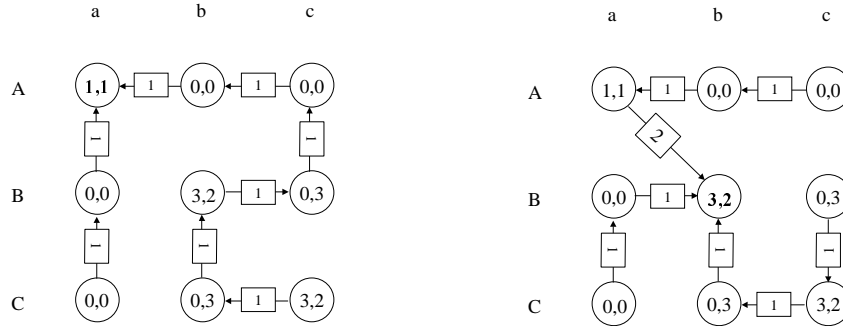


FIGURE 3—Minimum-resistance trees rooted at $h_{(A,a)}$ and $h_{(B,b)}$, respectively.

3.2 Stochastically Stable Sets

In this subsection, we will show that the stochastically stable states of imitation play correspond to pure-strategy profiles in particular subsets

of the strategy-space. In order to do this, we will introduce some new concepts and prove four lemmata.

Lemma 2 *If $x' \in (\gamma_i(x), x_{-i})$ for some player $i \in N$, then the following inequality holds for the stochastic potentials $\rho(x)$ and $\rho(x')$ of the monomorphic states h_x and $h_{x'}$, respectively:*

$$\rho(x) \geq \rho(x'). \quad (2)$$

PROOF: By definition, $\rho(x)$ is the minimum resistance over all trees rooted at state h_x . Construct a tree rooted at $h_{x'}$ by taking (one of) the tree(s) with minimum resistance rooted at h_x , adding the directed edge from x to x' and deleting the directed edge from x' . By Lemma 1 above, the resistance of the added edge is exactly one and that of the deleted edge at least one, so the total resistance of the new tree is $\rho(x)$ at most.

■

Definition 1 *A non-empty set of strategy profiles V is **closed under single better replies**, or a **cusber** set, if, for each $x \in V$ and $i \in N$, $(\gamma_i(x), x_{-i}) \subseteq V$. Such a set is called a **minimal cusber** set if it does not properly contain another cusber set.*

From the definition it follows that every game contains a minimal cusber set and that any *minimal closed sets under better replies* (Ritzberger and Weibull, 1995) contains a minimal cusber set. However, unlike minimal closed sets under better replies, minimal cusber sets are not necessarily product sets (see for example the game in Figure 4).

Let a *better-reply path* be a sequence of pure-strategy profiles x^1, x^2, \dots, x^l such that for every $k \in \{1, \dots, l-1\}$, there exists a unique player, say player i , such that $x^{k+1} \in (\gamma_i(x^k), x_{-i}^k)$ and $x_i^{k+1} \neq x_i^k$.

Lemma 3 *For any two strategy profiles x, x' in a minimal cusber set V , there exist better-reply paths x, \dots, x' and x', \dots, x , which connect these strategy profiles.*

PROOF: Let V be a minimal cusber set. Suppose that the claim is false and there exist two pure strategy profiles $x, x' \in V$ such that there is no better-reply path from x to x' . Consider all better-reply paths, starting at strategy profile x . There are a finite number of paths and a finite number of pure strategy profiles along all these paths. Collect all these strategy profiles in a set. By construction this set is a cusber set and by assumption it does not contain the vertex x' , contradicting the assumption of V being a minimal cusber set. ■

To every strategy profile $x \in X$ of a minimal cusber set V corresponds a monomorphic state h_x which is a repetition of the associated pure-strategy profile. We call the set of these monomorphic states an α -set and denote it by \mathbb{V} . In other words, $\mathbb{V} = \{h_x \in H : x \in V\}$.

Lemma 4 *All monomorphic states in an α -set have equal stochastic potential.*

PROOF: On the one hand, by Lemma 3, there exists a better-reply path from an arbitrary strategy profile x in a minimal cusber set to any other strategy profile x' in the same minimal cusber set. Let the sequence x, \dots, x' be such a path. By Lemma 2, the following inequalities hold for the stochastic potential of the corresponding monomorphic states $h_x, \dots, h_{x'}$:

$$\rho(x) \geq \dots \geq \rho(x'). \quad (3)$$

On the other hand, by applying Lemma 3 once more, there exist a better-reply path from the strategy tuple x' to the strategy tuple x . Using

Lemma 2, gives:

$$\rho(x') \geq \dots \geq \rho(x). \quad (4)$$

From the inequalities in (3) and (4) it follows that $\rho(x) = \rho(x')$ for the monomorphic states h_x and $h_{x'}$. ■

Based on Lemma 4, we can define the *stochastic potential* of an α -set as the stochastic potential of any monomorphic state contained in it.

Lemma 5 *For every monomorphic state which does not belong to any α -set, there exists a monomorphic state with lower stochastic potential.*

PROOF: See the Appendix.

We are now in a position to state the following main theorem.

Theorem 3 *The stochastically stable states of $P^{\alpha, m, s, \varepsilon}$ are the monomorphic states contained in α -set(s) with minimum stochastic potential.*

PROOF: The theorem follows immediately from Lemma 4, Lemma 5, and Theorem 2. ■

Theorem 3 establishes a relation between the stochastically stable states of imitation play and minimal cusber sets, which is similar to the relationship between the stochastically stable states and minimal curb sets first proved for a particular dynamic by Hurkens (1995), and later modified for a different dynamic by Young (1998, p. 111).

We say that a finite set Y of pure-strategy profiles *strictly Pareto dominates* a pure-strategy profile x if for any pure-strategy profile $y \in Y$, $\pi_i(y) > \pi_i(x)$, for all i . The following theorem shows that if the sample size is sufficiently large, imitation play selects sets of efficient outcomes in a large class of games.

Theorem 4 *If a minimal cusber set V is a product set and strictly Pareto dominates all pure-strategy profiles outside V , then there exists an $s^* \geq 1$, such that for every sample size $s > s^*$, the stochastically stable states of $P^{\alpha, m, s, \varepsilon}$ are the monomorphic states contained in the α -set \mathbb{V} .*

PROOF: See the Appendix.

The intuition behind this result is that for a sufficiently large sample size, the transition from a state inside \mathbb{V} to any state outside \mathbb{V} requires more mistakes than the number of player roles, while the opposite transition requires one mistake per player role at most. The following corollary follows immediately from Theorem 4.

Corollary 1 *If a strict Nash Equilibrium x strictly Pareto dominates all other pure-strategy profiles, then there exists an $s^* \geq 1$, such that for every sample size $s > s^*$, h_x is a unique stochastically stable state of $P^{\alpha, m, s, \varepsilon}$.*

The requirement in Theorem 4 that V be a product set is necessary, as shown by the game in Figure 4.

	a	b	c
A	1, 1	0, 0	0, 0
B	0, 0	0, 0	5, 3
C	0, 0	2, 4	5, 4

FIGURE 4

In this game, the minimal cusber set $V = \{(C, b), (C, c), (B, c)\}$ strictly Pareto dominates all pure-strategy profiles outside V . Let \mathbb{V} be the α -set associated with V . It is evident that two mistakes are enough to move from monomorphic state $h_{(A, a)}$ to any monomorphic state in \mathbb{V} . We will show that two mistakes are also enough to move from \mathbb{V} to

a monomorphic state outside of \mathbb{V} . Suppose the process is in the state $h_{(C,c)}$ at time t . Further, suppose that the individual in player role 1 plays B instead of C at time $t + 1$ by mistake. This results in play (B, c) at time $t + 1$. Assume that the individual in player role 2 makes a mistake and plays b instead of c , and that the individual in player role 1 plays C in period $t + 2$. Hence, the play at time $t + 2$ is (C, b) . Assume that the individuals in both player roles sample from period $t - s + 2$ to period $t + 2$ for the next s periods. This means that the individuals in player role 1 choose to play B and the individuals in player role 2 choose to play b from period $t + 3$ to period $t + s + 2$. There is a positive probability that from period $t + s + 3$ through period $t + m + 2$, the individuals in both player roles will sample from periods later than $t + 2$. Hence, by the end of period $t + m + 2$, there is a positive probability that the process will have reached the monomorphic state $h_{(B,b)}$ outside \mathbb{V} . It is now straightforward to show that all the monomorphic states $h_{(A,a)}$, $h_{(C,b)}$, $h_{(C,c)}$, and $h_{(B,c)}$ have equal stochastic potential.

4 Applications to 2×2 Games

In this section, we apply the results from Section 3 to 2×2 games. First, we find the stochastically stable states in three special classes of games and second, we study the speed of convergence in general 2×2 games. In all of the following games, we denote player 1's strategies A and B , and player 2's strategies a and b , respectively.

4.1 Stochastically Stable States

In this subsection, we analyze the stochastic stability in games with two strict Nash equilibria, games with one strict Nash equilibrium and games without Nash equilibria in pure strategies. We start with the class of games with a unique strict Nash equilibrium, which includes, for

example, *Prisoners' Dilemma Games*.

Corollary 2 *In 2×2 games with a unique strict Nash equilibrium, the corresponding monomorphic state is a unique stochastically stable state.*

PROOF: Games in this class contain exactly one minimal cusber set, consisting of the unique strict Nash equilibrium. By Theorem 3, the corresponding monomorphic state must be the unique stochastically stable state. ■

We now proceed with the class of games with two strict Nash equilibria. *Coordination Games* and *Hawk-Dove Games* are two examples of games in this class.

Corollary 3 *In 2×2 games with two strict Nash equilibria, where one Nash Equilibrium strictly Pareto dominates the other, there exists an $s^* \geq 1$, such that for every sample size $s > s^*$, the unique stochastically stable state corresponds one to one with the monomorphic state of the Pareto dominant equilibrium.*

PROOF: Games in this class contain two minimal cusber sets, either $\{(A, a)\}$ and $\{(B, b)\}$ or $\{(A, b)\}$ and $\{(B, a)\}$. Without loss of generality, assume that $\{(A, a)\}$ and $\{(B, b)\}$ are the minimal cusber sets. If (A, a) strictly Pareto dominates (B, b) , it must also strictly Pareto dominate the two other pure-strategy profiles, (A, b) and (B, a) . Hence, by Corollary 1, the monomorphic state $h_{(A, a)}$ is the unique stochastically stable state. ■

This implies that, unlike Young's (1993) process of adaptive play, imitation play does not generally converge to the risk-dominant equilibrium in Coordination Games. Our result is consistent with Robson

and Vega-Redondo's (1996) result for generic symmetric Coordination Games. However, for the non-generic case when $\pi_i(A, a) = \pi_i(B, b)$, the stochastically stable states in their model depend on the details of the adjustment process, whereas imitation play always selects both equilibria.

Proposition 1 *In 2×2 games with two strict Nash equilibria, where neither strictly Pareto dominates the other, the stochastically stable states correspond one to one with the monomorphic states of the two equilibria.*

PROOF: See the Appendix.

Finally, we consider games which do not have any Nash equilibrium in pure strategies. One of the games in this class is the *Matching Pennies Game*.

Corollary 4 *In 2×2 games without Nash equilibria in pure strategies, the stochastically stable states correspond one to one with the four monomorphic states of the game.*

PROOF: Games in this class contain exactly one minimal cusber set $\{(A, a), (A, b), (B, a), (B, b)\}$. By Theorem 3, the four corresponding monomorphic states are all stochastically stable. ■

4.2 Speed of Convergence

In this subsection, we analyze the speed of convergence of imitation play.

Proposition 2 *In 2×2 games, the maximum expected first passage time for the perturbed process $P^{\alpha, m, s, \varepsilon}$ from any state to a stochastically stable state is at most $\delta \varepsilon^{-2}$ units of time, for some positive constant δ .*

PROOF: The claim follows from the observation that in any 2×2 game, the transition from an arbitrary state to a stochastically stable monomorphic state requires two mistakes at most. ■

This result should be contrasted with the speed of convergence in Young (1993), Kandori *et al.* (1993), and Robson and Vega-Redondo (1996). In Young's (1993) model, the maximum expected first passage for a 2×2 Coordination Game is at least $\delta_Y \varepsilon^{-v}$ where v depends on the sample size and both players' payoffs. In Kandori *et al.* (1993) the maximum expected first passage time is of the order $\delta_{KMR} \varepsilon^{-Nu}$, where N is the size of the population and u is determined by the game's payoff structure. In Robson and Vega-Redondo (1996), the corresponding figure is $\delta_{RV} \varepsilon^{-q}$, where q is a positive integer independent of the payoffs and the current state. Thus, when v , Nu , and q are greater than two and ε is sufficiently small, then imitation play converges considerably faster than the processes in these three models.

5 Extensions

All results in this paper hold for a more general class of imitation dynamics. Let the *maximum correspondence* be a correspondence which maps a strategy sample w_i and the associated payoff sample v_i to a strategy with the highest payoff in the sample. This correspondence defines a new Markov process on the space of histories with the same set of absorbing states and (for a sufficiently large sample size) stochastically stable states as imitation play. Moreover, if each population consists of arbitrary shares of individuals who make choices based on the maximum correspondence and the maximum average correspondence, respectively, then the results of this paper still hold. Hence, the model allows for a certain kind of population heterogeneity, where individuals make their

choices based on different rules.

6 Conclusion

In this paper we develop an evolutionary model with perpetual random shocks where individuals, in every period, choose the strategy with the highest average payoff in a finite sample of past play. We denote the resulting Markov process imitation play and prove that, provided information is sufficiently incomplete and the sample size sufficiently large, the stochastically stable states of imitation play are repetitions of the pure-strategy profiles in minimal closed sets under single better-replies. We call such sets minimal cusber sets. These sets are related to minimal closed sets under better replies. We also prove that if the pure-strategy profiles in a minimal cusber set have strictly higher payoffs than all outside pure-strategy profiles, then, provided that the minimal cusber set is a product set and the sample is sufficiently large, the pure-strategy profiles in this set will be played with probability one in the limit as the experimentation probability tends to zero. Our results give clear predictions in 2×2 games. In Coordination Games, where one equilibrium strictly Pareto dominates the other, imitation play selects the strictly Pareto superior Nash equilibrium. If neither equilibria strictly Pareto dominates the other, then both are stochastically stable. Finally, we show that the speed of convergence for imitation play in many cases is higher than in other known models.

The objective of this paper is to derive predictions for general finite games in a world of truly boundedly rational individuals. The assumption underlying the model, that individuals do not make decisions based on the predictions of their opponents' future strategies, but rather based on which strategies have been successful in the past, is maybe most ap-

peeling in the class of games where it is costly to obtain information about the opponents. A high cost may be due to the size or the complexity of the game or to institutional factors preventing the release of information about the opponents. It would be particularly interesting to test the implications of our model against empirical or experimental evidence in this class of games.

7 Appendix

PROOF OF LEMMA 5: For every strategy profile not included in any minimal cusber set, there exists a finite better-reply path which ends in some minimal cusber set. Let this path be $x^1, x^2, \dots, x^{T-1}, x^T$, where x^1 is an arbitrary strategy profile that does not belong to any minimal cusber set and x^T the first monomorphic state on the path belonging to some minimal cusber set, V . By Lemma 2, it follows that the following inequalities hold for the stochastic potential of the corresponding monomorphic states:

$$\rho(x^1) \geq \rho(x^2) \geq \dots \geq \rho(x^{T-1}) \geq \rho(x^T). \quad (5)$$

We will show that in fact, $\rho(x^{T-1}) > \rho(x^T)$. Note that $\rho(x^{T-1})$ is the minimum resistance over all trees rooted at the state $h_{x^{T-1}}$. Denote (one of) the tree(s) that minimizes resistance by $\mathfrak{S}(x^{T-1})$. Find in the tree $\mathfrak{S}(x^{T-1})$ a directed edge from some vertex h_y such that y is in the minimal cusber set V , to some other vertex $h_{y'}$ such that y' is outside this minimal cusber set. It will be shown later that there is only one such directed edge in the minimal resistance tree $\mathfrak{S}(x^{T-1})$. Delete in the tree $\mathfrak{S}(x^{T-1})$ the directed edge $h_y \rightarrow h_{y'}$ and add the directed edge $h_{x^{T-1}} \rightarrow h_{x^T}$. As a result, we obtain a tree $\mathfrak{S}(y)$ rooted at the state h_y . By Lemma 1, the resistance of the deleted edge is greater than one, and the resistance of the added edge is one. Therefore,

the total resistance of the new tree $\mathfrak{S}(y)$ is less than the stochastic potential $\rho(x^{T-1})$. Moreover, by Lemma 4, the monomorphic state h_{x^T} has the same stochastic potential as the monomorphic state h_y . Hence, $\rho(x^{T-1}) > \rho(x^T)$.

We will now consider the tree $\mathfrak{S}(x^{T-1})$ and show that there is only one directed edge from the monomorphic states which consists of a repetition of a strategy profile in a minimal cusber set to a state which consists of a repetition of a strategy profile outside the cusber set. Suppose there is a finite number of such directed edges $h_{y^j} \rightarrow h_{z^j}$, $j = 1, 2, \dots, l$, where y^1, \dots, y^l are strategy profiles in the minimal cusber set and z^1, \dots, z^l strategy profiles outside the cusber set. It is clear that there cannot be an infinite number of outgoing edges since the game Γ is finite. Recall that a tree rooted at vertex h_{y^j} is a set of $|X| - 1$ directed edges such that, from every vertex different from h_{y^j} , there is a unique directed path in the tree to h_{y^j} . The resistance of any directed edge $h_{y^j} \rightarrow h_{z^j}$, $j = 1, 2, \dots, l$ is at least two. By Lemma 3, there exists a finite better-reply path from vertex y^1 to vertex y^2 in the minimal cusber set. Let $y^1, f^1, \dots, f^k, y^2$ be such a path.

Consider the vertex h_{f^1} . There are two mutually exclusive cases:

- 1.a) there exists a directed path from h_{f^1} to one of the vertices h_{y^2}, \dots, h_{y^l} in the initial tree $\mathfrak{S}(x^{T-1})$, or
- 1.b) there exists a directed path from h_{f^1} to h_{y^1} .

In case 1.a) by deleting the directed edge $h_{y^1} \rightarrow h_{z^1}$ and adding the directed edge $h_{y^1} \rightarrow h_{f^1}$ to the tree $\mathfrak{S}(x^{T-1})$, we obtain a new tree $\mathfrak{S}^1(x^{T-1})$ with lower stochastic potential than $\mathfrak{S}(x^{T-1})$, because the resistance of the directed edge $y^1 \rightarrow f^1$ is one. This means that we are done, since it contradicts the assumption of $\mathfrak{S}(x^{T-1})$ being a minimal resistance tree.

In case 1.b), we will use the following procedure for vertex h_{f^1} : delete the initial directed edge from h_{f^1} and add the directed edge $h_{f^1} \rightarrow h_{f^2}$. As above, there are two cases:

2.a) there exists a directed path from h_{f^2} to one of the vertices h_{y^2}, \dots, h_{y^l} in the initial tree $\mathfrak{S}(x^{T-1})$, or

2.b) there exists a directed path from h_{f^2} to h_{y^1} .

In case 2.a), we obtain a new tree $\mathfrak{S}^2(x^{T-1})$ with lower stochastic potential than $\mathfrak{S}(x^{T-1})$, because the resistance of the directed edge $h_{f^1} \rightarrow h_{f^2}$ is one. This means that we are done, since it contradicts the assumption of $\mathfrak{S}(x^{T-1})$ being a minimal resistance tree.

In case 2.b), we repeat the procedure for vertices h_{f^2}, h_{f^3}, \dots . The better-reply path $y^1, f^1, \dots, f^k, y^2$ from y^1 to y^2 is finite. Hence, after $k + 1$ steps at most, we have constructed a tree $\mathfrak{S}''(x^{T-1})$ rooted at the state $h_{x^{T-1}}$ with lower stochastic potential than $\mathfrak{S}(x^{T-1})$. ■

PROOF OF THEOREM 4: Suppose $V \subset X$ is a minimal cusber set which strictly Pareto dominates all pure-strategy profiles outside V . We will prove Theorem 4 using two lemmata.

Lemma 6 *The transition from any monomorphic state h_z to a monomorphic state h_x such that $x \in V$ requires at most n mistakes, independently of the sample size.*

PROOF: Assume that the process is in state h_z and that the individuals in all player roles simultaneously mistakes, so that x is played instead of z . Since, by assumption, $\pi_i(x) > \pi_i(z)$ for all player roles i , if the individuals in all player roles sample the most recent plays for the next $m - 1$ periods, this will take the process to the state h_x . ■

Lemma 7 *There exist an $s^* \geq 1$, such that for sample size $s > s^*$, the transition from a monomorphic state h_x such that $x \in V$ to a monomorphic state h_z such that $z \notin V$ requires at least $n + 1$ mistakes.*

PROOF: The transition from h_x to a monomorphic state h_z such that $z \notin V$ can be made if individuals in one of the player roles, say i , make (at least) s consecutive mistakes and play a strategy $z_i \notin V_i$ every time. If the individuals in player role i , thereafter sample from the most recent plays for $m - s$ periods, the Markov chain will end up in the state h_z where $z = (z_i, x_{-i})$. Hence, if $s > n$, this kind of transition will require more than n mistakes.

Alternatively, the transition from h_x to a monomorphic state h_z such that $z \notin V$ can be made if a single individual in one of the player roles, say i , makes a mistake and play z_i and individuals in other player roles thereafter (or before this) make sufficiently many mistakes to make the average payoff of x_i lower than that of z_i in a sample of the most recent plays. The number of mistakes required for this kind of transition will be as low as possible if $\pi_i(x)$ is as *low* as possible, $\pi_i(z)$ is as *high* as possible, and the minimum expected payoff is achieved when i plays z_i , an individual in a different player role j plays $q_j \notin V_j$, and all other individuals drawn to play x_{-i} .

Let $\underline{\pi}$ be the *minimum* payoff to any player role for any pure-strategy profile *in* V , let $\underline{\pi}$ be the *minimum* payoff to any player role for any pure-strategy profile *outside* V , and let $\bar{\pi}$ be the *maximum* payoff to any player role for any pure-strategy profile *outside* V . By assumption, $\underline{\pi} > \bar{\pi} \geq \underline{\pi}$. By the above logic, the transition from h_x to h_z will require

at least $n + 1$ mistakes if

$$\frac{(s - n)\check{\pi} + (n - 1)\underline{\pi}}{s - 1} > \bar{\pi} \quad (6)$$

\Leftrightarrow

$$s > \frac{n(\check{\pi} - \underline{\pi}) + \underline{\pi} - \bar{\pi}}{\check{\pi} - \bar{\pi}}. \quad (7)$$

Note that the right-hand side in the last expression is greater than or equal to n . Hence, if $s > s^* = \frac{n(\check{\pi} - \underline{\pi}) + \underline{\pi} - \bar{\pi}}{\check{\pi} - \bar{\pi}}$, then the transition from a monomorphic state h_x , such that $x \in V$, to a monomorphic state h_z , such that $z \notin V$, requires more than n mistakes. ■

Let h_y be an arbitrary monomorphic state such that $y \notin V$ and consider a minimal resistance tree $\mathfrak{S}(y)$ rooted at h_y . Let h_x be a monomorphic state such that $x \in V$ and such that there is a directed edge from h_x to a state h_z , with $z \notin V$, in the tree $\mathfrak{S}(y)$. Create a new tree $\mathfrak{S}(x)$ rooted at h_x by adding a directed edge from h_y to h_x and deleting the directed edge from h_x to h_z in the tree $\mathfrak{S}(y)$. By Lemma 7, the deleted edge has a resistance greater than n provided that $s > s^*$, and by Lemma 6, the added edge has a resistance of at most n . Hence, for $s > s^*$ the total resistance of the new tree $\mathfrak{S}(x)$ is less than that of $\mathfrak{S}(y)$. Theorem 4 now follows by Lemma 4, according to which that all monomorphic states of an α -set have equal stochastic potential. ■

PROOF OF PROPOSITION 3: As in the proof of Corollary 3, without loss of generality, assume that $\{(A, a)\}$ and $\{(B, b)\}$ are the minimal cusber sets of the game. By Theorem 3, it follows that monomorphic states $h_{(A, a)}$ and $h_{(B, b)}$ are the only two candidates for the stochastically stable states. Suppose that only one of these monomorphic states is stochastically stable, say $h_{(A, a)}$. Let $\mathfrak{S}(A, a)$ be a minimum resistance tree with resistance $\rho(A, a)$ rooted at $h_{(A, a)}$. In this tree, there is an outgoing edge from the monomorphic state $h_{(B, b)}$.

First, note that the resistance of this edge is at least two, such that at least two mistakes are needed to move from the monomorphic state $h_{(B,b)}$. This follows since $\pi_1(B, b) > \pi_1(A, b)$ and $\pi_2(B, b) > \pi_2(B, a)$.

Second, note that two mistakes are sufficient to move the process from the monomorphic state $h_{(A,a)}$ to the monomorphic state $h_{(B,b)}$. Suppose the process is in state $h_{(A,a)}$. Since neither of the Nash equilibria is strictly Pareto superior, either $\pi_1(B, b) \geq \pi_1(A, a)$ and/or $\pi_2(B, b) \geq \pi_2(A, a)$. Without loss of generality, assume that the first of these inequalities holds. Then, there is a positive probability that the individuals in both player roles simultaneously make mistakes at time t . There is also a positive probability that the individuals in player role 1 draw the fixed sample $(x_1^{t-s+1}, \dots, x_1^t)$ with corresponding payoffs $(\pi_1^{t-s+1}, \dots, \pi_1^t)$ and that the individuals in player role 2 sample from plays earlier than x_2^t from period $t + 1$ to, and including, period $t + s - 1$. With positive probability individuals in player role 1 play B and individuals in player role 2 play a in all of these periods. This implies that if the individuals in both player roles sample from plays more recent than x_i^{t-1} from period $t + s$ to, and including, period $t + s - 1 + m$, 1's sample will only contain strategy B and 2's sample will always contain strategy b , and possibly, strategy a . Furthermore, the average payoff of strategy b will be $\pi_2(B, b)$ as compared to an average payoff of $\pi_2(B, a)$ for strategy a , when the latter strategy is included in the sample. Hence, with positive probability the process will be in state $h_{(B,b)}$ at time $t + s - 1 + m$.

Finally, create a new tree rooted at $h_{(B,b)}$ by deleting the outgoing edge from the monomorphic state $h_{(B,b)}$ in the tree $\mathfrak{S}(A, a)$ and adding an edge from $h_{(A,a)}$ to $h_{(B,b)}$. The resistance of the deleted edge is at least two and that of the added edge two. Hence, the total resistance of the new tree is at most $\rho(A, a)$, thereby contradicting the assumption

that only $h_{(A,a)}$ is stochastically stable. ■

References

- [1] Basu, K., and Weibull J. (1991). “Strategy Subsets Closed under Rational Behavior,” *Economics Letters* 36, 141-146.
- [2] Björnerstedt, J., and Weibull J. (1996). “Nash Equilibrium and Evolution by Imitation,” in *The Rational Foundations of Economic Behaviour* (K. Arrow et al. Eds.), 155-171. MacMillan.
- [3] Duffy, J., and Feltovich N. (1999). “Does Observation of Others Affect Learning in Strategic Environments? An Experimental Study,” *International Journal of Game Theory* 28, 131-152.
- [4] Graham, J. R. (1999). “Herding among Investment Newsletters: Theory and Evidence,” *Journal of Finance* 54, 237-268.
- [5] Griffiths, M. D., Smith, B. F., Turnbull, D., Alasdair, S., and White R. W. (1998). “Information flows and open outcry: evidence of imitation trading,” *Journal of International Markets, Institutions & Money* 8, 101-116.
- [6] Huck, S., Normann, H.-T., and Oechssler J. (1999). “Learning in Cournot Oligopoly: An experiment,” *Economic Journal* 109, 80-95.
- [7] ——— (2000). “Does information about competitors’ actions increase or decrease competition in experimental oligopoly markets?,” *International Journal of Industrial Organization* 18, 39-58.
- [8] Hurkens, S., (1995). “Learning by forgetful players,” *Games and Economic Behavior* 11, 304-329.
- [9] Foster, D., and P. Young (1990). “Stochastic Evolutionary Game Dynamics,” *Theoretical Population Biology* 38, 219-32.
- [10] Freidlin, M., and A. Wentzell (1984). *Random Perturbations of Dy-*

namical Systems. Springer-Verlag.

- [11] Kandori, M., Mailath, G., and Rob R. (1993). "Learning, Mutation and Long-Run Equilibria in Games," *Econometrica* 61, 29-56.
- [12] Nash, J. (1950). "Non Cooperative Games," Ph.D. thesis, Princeton University.
- [13] Nöldeke, G., and Samuelson L. (1993). "An Evolutionary Analysis of Backward and Forward Induction", *Games and Economic Behavior* 5, 425-454.
- [14] Osborne, M. J., and Rubinstein A. (1998). "Procedurally Rational Players," *American Economic Review* 88, 834-847.
- [15] Ritzberger, K., and Weibull J. (1995). "Evolutionary Selection in Normal-Form Games" *Econometrica* 63, 1371-1400.
- [16] Robson, A. J., and Vega-Redondo F. (1996). "Efficient Equilibrium Selection in Evolutionary Games with Random Matching", *Journal of Economic Theory* 70, 65-92.
- [17] Schlag, K. (1998). "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits," *Journal of Economic Theory* 78, 130-56.
- [18] ——— (1999). "Which One Should I Imitate?," *Journal of Mathematical Economics* 31, 493-522.
- [19] Sobel, J. (1993). "Evolutionary Stability and Efficiency," *Economic Letters* 42, 301-312.
- [20] Wermers, R. (1999). "Mutual Fund Herding and the Impact on Stock Prices," , *Journal of Finance* 54, 237-268.
- [21] Young, P. H. (1993). "The Evolution of Conventions," *Econometrica* 61, 57-84.
- [22] ——— (1998). *Individual Strategy and Social Structure*. Princeton University Press.