

Advances in Theoretical Economics

Volume 2, Issue 1

2002

Article 1

Local Conventions

Jeffrey C. Ely*

*Economics Department, Northwestern University, junk@wilamette.econ.northwestern.edu

Copyright ©2002 by the authors. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher, bepress, which has been given certain exclusive rights by the author. *Advances in Theoretical Economics* is produced by The Berkeley Electronic Press (bepress). <http://www.bepress.com/bejte>

Local Conventions

Jeffrey C. Ely

Abstract

It is shown that player mobility has important consequences for the long-run equilibrium distribution in dynamic evolutionary models of strategy adjustment, when updating is prone to small probability perturbations, i.e. “mistakes” or “mutations.” Ellison (1993) concluded that the effect on the matching process of localized “neighborhoods” was to strengthen the stability of risk-dominant outcomes, originally demonstrated by Kandori, Mailath, and Rob (1993) (KMR) and Young (1993). I consider a model in which players can choose the neighborhoods to which they belong. When strategies and locations are updated simultaneously, only efficient strategies survive. The robustness of this conclusion is emphasized in a general locational model in which strategy revision opportunities are allowed to arrive at a faster rate than opportunities to change locations. The efficient strategy persists in all cases in which the locational structure is non-trivial. Moreover, even as the relative frequency of player mobility approaches zero, the efficient strategy occurs with boundedly positive relative frequency. This result is in stark contrast to the conclusions of the previous models.

KEYWORDS: stochastic evolution, local interaction, mobility

1 Introduction

Much literature in evolutionary game theory has focused on equilibrium selection in coordination games. A great deal of attention has been focused on the result that the equilibrium satisfying the *risk-dominance* criterion of Harsanyi and Selten (1988) is the unique stable outcome in a wide variety of dynamic models. The models of Kandori, Mailath, and Rob (1993), and Young (1993) are the best-known versions of this result, and it is supported in related models by Fudenberg and Harris (1992), Foster and Young (1990), Ellison (1993), and Blume (1993), among others.

The common feature of each of these models is an evolutionary process consisting of *random matching* and *stochastic adjustment*. Players are randomly paired in each of an infinite sequence of rounds. Following each round of play, strategies are updated in favor of best replies, but this process is subject to stochastic “mistakes.”

The original papers assumed uniform matching: each player was equally likely to be paired with every other. In a subsequent paper, Ellison (1993) focused on the role of this type of matching rule in the dynamics. He compared models with uniform matching with one in which players were arranged within a network of small, overlapping “neighborhoods” and were paired only within their own neighborhood. He found that this type of local matching accelerates the transition to risk-dominant outcomes and thus can only strengthen the stability results of previous papers.

Local matching is a natural point of extension to evolutionary models of equilibrium selection. In most environments, an agent’s set of potential opponents will be a strict subset of the total population, confined to a specific location, place of work, social group, etc. As Ellison showed, this feature plays an important role in the dynamic process. Once we recognize this role, however, it seems natural to question the assumption that the neighborhood structure is exogenous to the evolutionary process. In many environments, agents have some control over the set of individuals with whom they may interact. People choose where and with whom to live, where to work, and which bowling leagues to join.

This paper begins with the assumption that through choices such as these, players have some control over their set of opponents. The results show that the logic underlying the selection of risk-dominant outcomes no longer applies in such an environment. Rather, the ability to seek out “good” outcomes favors the emergence of Pareto-efficient equilibria. Moreover, deviators to “bad” strategies can be avoided, offering efficient outcomes a further edge in stability.

An example and a brief review of the literature on evolution and coordination games help to illustrate the idea. Consider the following 2x2 coordination game.

	k_1	k_2
k_1	4, 4	-2, 3
k_2	3, -2	0, 0

Each strategy is a strict symmetric Nash equilibrium. The outcome (k_2, k_2) is risk-dominant but Pareto dominated by (k_1, k_1) . I will refer to k_2 as the risk-dominant strategy and k_1 as the efficient strategy.

This game has been analyzed by Kandori, Mailath, and Rob (1993), and Young (1993) using a model of stochastic adjustment. Each period, players in a population are randomly matched to play the stage game. Following each round of play, some subset of players are given the opportunity to revise their behavior. In the version of the model I will use, this opportunity arises with some positive probability, independent across players. When the opportunity arises, players update *myopically*. That is, they choose a best reply to the current population configuration. This process of updating, in which decisions are motivated by myopic payoff maximization is called a *Darwinian adjustment process*.

In the game above, regardless of the matching rule, the Darwinian adjustment process leads with probability one to a configuration in which all players play the same pure strategy.¹ Following Young (1993), such profiles will be called *conventions*. Which convention will be reached depends on the initial conditions, hence the Darwinian adjustment process alone does not yield a sufficiently strong prediction to be a model of equilibrium selection.

Adding small probability “mistakes” in the updating process yields a unique prediction in the form of an invariant probability distribution over states of the system. The *limit distribution*, i.e. the limit of this frequency distribution as mistake probabilities go to zero, will assign positive probability only to those states that are sufficiently robust to these mistakes. KMR and Young show that when the matching process is uniform, the limit distribution puts probability 1 on the risk-dominant convention, k_2 . They conclude that the risk-dominant equilibrium, by virtue of being less susceptible to deviation, is the more stable outcome.

Ellison (1993) observed that the k_1 convention, despite being less stable than k_2 , is nevertheless quite stable in absolute terms. Many near-simultaneous mistakes are required to overturn a convention when every

¹The extension of this result to arbitrary local interaction structures is proved in Proposition 3.

player in the population is a potential opponent. Thus the k_1 convention may persist for a long time before the occurrence of a number of simultaneous mistakes sufficient to lead to k_2 , a fact which raises doubts about the predictive power of this type of analysis.

With this motivation, Ellison (1993) compared the uniform matching process to one in which players occupy distinct locations on a circular grid, with the assumption that each player is paired only with a subset of players occupying nearby locations (i.e. the “neighborhood”).²

While this modification does not affect the character of the limit distribution, Ellison showed that localized matching introduced two features which accelerate the transition to the risk-dominant convention. First, with a small set of potential opponents, fewer mistakes are required to establish k_2 as a best-reply, and secondly, any localized outbreak of k_2 quickly sweeps through the population through the fixed network of overlapping neighborhoods.

Furthermore, his analysis demonstrates that this phenomenon relies on the risk-dominance property of k_2 . Thus, the same features do not accelerate the undoing of a k_2 convention. The result is that in Ellison’s local matching model, transition to the risk-dominant outcome takes place much sooner, and transitions away from it are much less frequent than in the uniform matching of KMR and Young.

In this paper, I examine these conclusions in a model in which in addition to being matched within localized neighborhoods, players can choose the neighborhoods to which they belong. In the simplest version of the model, when a revision opportunity arises, a player simultaneously chooses a strategy and location in order to maximize his expected payoff. In such a model, risk-dominance ceases to play an important role in determining long-run equilibria. In fact, the limit distribution places probability 1 on the efficient convention.

Consider Ellison’s circle model starting with a profile in which all players are playing the efficient strategy k_1 and earning the maximal per-period payoff of 4. Suppose a few neighbors of player i mistakenly switch strategies to k_2 . If i ’s location were exogenously fixed as in Ellison’s model, i will switch to k_2 in order to play a best-response to his (permanent) neighbors. If instead, i had the option to move elsewhere on the circle, he would prefer to avoid the deviators and find a location where he could continue to play k_1 and receive the efficient payoff.

²See also Blume (1993), who considers a related spatial environment, and obtains similar results.

But this possibility prevents the type of contagion process which is central to the risk-dominance selection. When all of the deviators' closest neighbors have moved away, their link to the rest of the population is cut. In fact, as soon as they have the chance, the deviators themselves will move away and revert to k_1 in order to maximize their payoff.

The argument above demonstrates the difficulty of upsetting an efficient convention when players can choose their location. Equally importantly, endogenous location facilitates the *emergence* of efficient conventions. Suppose all players are playing k_2 . If a single player mistakenly moves to an unoccupied location and begins playing k_1 , all players will quickly follow in order to play k_1 against him and receive the efficient payoff. This type of event can be triggered by a single mistake under any (non-degenerate) locational structure.

The reader is cautioned in interpreting these results that while there are many natural ways of extending a model of local interaction to account for mobility, I have taken a specific approach. Among the modeling choices to be made is the specification of payoffs. I have chosen to calculate the payoff to a player by averaging over the payoffs to the interactions with each of his neighbors. Alternatively, one could consider the sum of payoffs. Note that with a fixed interaction structure, the number of neighbors is fixed and these measures are equivalent. On the other hand, when a player can control the number of neighbors by relocating, the specification clearly matters. A more general model would allow a more general specification of payoffs, allowing for trade-offs between average payoffs and the number of interactions. The effects of such an extension are discussed at the end of Section 3.

These strategic features are closely related to the conclusions in the literature on evolution and pre-game communication. For example, Kim and Sobel (1995) and Matsui (1991) found that when players are able to make costless announcements prior to each round of play, evolutionary pressures select efficient outcomes. In these models, efficient strategies are protected from bad outcomes through the use of a "secret handshake." The efficient strategy would be used only upon receiving a special signal identifying the opponent's intention to do the same. In the models of this paper, locations play a role analogous to messages. Players of the efficient identify one another through their choice of location.

In the communication literature, the main obstacle to proving that efficient outcomes emerge is to ensure there is at least one message unused by the current population which can become a secret handshake. The analogue to this difficulty in the context of location choice is to guarantee an empty neighborhood. In this paper, no assumption is made to guarantee

that there will be an empty neighborhood in every state. Rather, mobility itself ensures this eventuality: by a sequence of isolated mistakes, players will drift among locations until a neighborhood is emptied, paving the way for the emergence of the efficient strategy.

These arguments are used to show that when location and strategy decisions are made simultaneously, efficient conventions survive. This is first established in Section 3 in the context of a simple interaction structure in which neighborhoods consist of mutually exclusive locations within each of which matching is uniform. In Section 4, this result is shown to hold under a general class of interaction structures (part 2 of Theorem 2).

Section 4 further generalizes the simple model by allowing location and strategy decision opportunities to arise independently. This provides a framework for a type of robustness test for models without mobility. When strategy revision opportunities arise infinitely more frequently than the opportunities to change locations, the model approximates one in which migration is impossible. Yet, as is shown in part 3 of Theorem 2, efficient conventions occur with a frequency bounded away from zero. Thus, adding a little migration to the model results in a discontinuity in the limit distribution.

2 Related Literature

Some previous authors have studied mobility and its implications for equilibrium selection in random matching models. Most closely related to this paper is Oechssler (1999), who considers the simple interaction structure analyzed in Section 1. He assumes that in the initial state of the system, each equilibrium is represented as a convention in some location. While in general the limit point of his dynamic process depends on the initial conditions, he demonstrates that under his assumption, the system will converge to the efficient outcome. Section 1 builds on this analysis by showing how these initial conditions can arise via a sequence of single mistakes, and that overturning efficient conventions requires events of infinitely smaller probability. Section 2 extends these results to general interaction structures.

Mailath, Samuelson, and Shaked (2001) study endogenous matching with general interaction structures. In a model with simultaneous strategy/location revision, and an interaction structure similar to the types considered here, they also obtain an efficiency result. They also provide examples of alternative interaction structures in which efficiency needn't obtain. For example, if players are unable to segregate themselves from those playing

the inefficient strategy, then the presence of a few may not provide enough incentive for others to switch. These examples illustrate that an important assumption about interaction structures in this paper is that they always provide players with the ability to isolate themselves from others.

Robson (1993) analyzes a model motivated by biological evolution. The population is partitioned into a finite set of subpopulations which grow independently at rates which reflect the payoffs earned by the players within them. Populations in which the efficient strategy is played therefore grow faster than others. At exogenously fixed intervals, all populations become extinct and are re-populated by small groups randomly drawn from the preceding generation. When the time between extinction events is sufficiently long, the populations in which the efficient strategy is played grow arbitrarily large relative to other populations.

3 A Simple Model

This section presents a simple model intended to highlight the effect of mobility in a standard random matching model. A finite population \mathcal{I} of players, and a finite set of neighborhoods \mathcal{J} , are specified. I assume $|\mathcal{J}| \geq 2$. Each period, players are matched within neighborhoods to play a symmetric 2x2 coordination game \mathcal{G} , with strategy set $\mathcal{K} = \{k_1, k_2\}$ and payoff function $u : \mathcal{K} \times \mathcal{K} \rightarrow \mathbf{R}$ represented by the following matrix.

	k_1	k_2
k_1	a,a	e,f
k_2	f,e	b,b

I assume the generic case that the two pure strategy equilibria are strict and Pareto rank-able. Assume $a > b$, and call k_1 the efficient strategy. Let τ be the probability assigned to k_1 by the (symmetric) mixed equilibrium strategy. That is, τ solves $\tau a + (1-\tau)e = \tau f + (1-\tau)b$. Only if the opponent plays k_1 with at least probability τ is k_1 a best-reply. I shall assume the interesting case that $\tau > 1/2$ so that k_2 is a best reply to an equal mixing, and hence risk-dominant.

A *demographic* is an assignment $d : \mathcal{I} \rightarrow \mathcal{J} \times \mathcal{K}$ of a strategy and location to each player. The set of all demographics is \mathcal{D} . A demographic is called a *convention* if for some $k \in \mathcal{K}$, $d_{\mathcal{K}}(i) = k$ for every player i , where $d_{\mathcal{K}}(i)$ represents the strategy component of $d(i)$. The set of demographics which are k_n -conventions, and for which there are no loners will be denoted Z_n , for $n = 1, 2$.

A location is an exclusive environment within which matching can take place. Given a demographic d , let $N_d(i)$ be the set of i 's neighbors: $N_d(i) = \{\hat{i} \neq i : d_{\mathcal{J}}(i) = d_{\mathcal{J}}(\hat{i})\}$. This is the set of potential opponents player i faces in demographic d . This type of locational structure allows the cleanest illustration of the role of player mobility in that it provides a simple way for players to avoid groups in which inferior equilibria are played. In the next section I show that this assumption is not necessary for the result. In fact the basic conclusion holds under any non-trivial neighborhood structure.

Each player assumes he is equally likely to be matched with each of his neighbors. Therefore, when d is the demographic, we can write the expected payoff to player i as follows.

$$\pi_i(d) = \frac{1}{|N_d(i)|} \sum_{\hat{i} \in N_d(i)} u(d_{\mathcal{K}}(i), d_{\mathcal{K}}(\hat{i})) \quad \text{if } N_d(i) \neq \emptyset$$

When $N_d(i) = \emptyset$, we need some way to evaluate the payoff to player i . The choice of assumption is potentially important as it will determine whether singleton neighborhoods can persist, or whether a *loner* will return to an occupied neighborhood when given the chance. I assume that a player with no opponent receives a fixed reservation utility u_r which is smaller than either equilibrium payoff (i.e. $u_r < b$). This ensures that both conventions are preferred to non-participation.

I adapt the stochastic adjustment process used by Samuelson (1994) and Nöldeke and Samuelson (1993) to the context of endogenous location. Following each round of random matching, some subset of the population is given the opportunity to adjust their behavior. This opportunity arises with independent probability β for each player.

When a revision opportunity occurs, a player chooses a strategy/location pair to maximize his per-period payoff given the current demographic. I assume that a player does not adjust to an alternative which is not a strict improvement over his current payoff. Formally, define the best response correspondence for player i as $b_i(d) = \operatorname{argmax}_{j,k} \pi_i(d^{-i}, (j, k))$, where the notation $(d^{-i}, (j, k))$ represents the demographic obtained from d by setting the location and strategy of player i to (j, k) . When a player is selected for revision, he selects an element at random from the *strict-improvement* correspondence $b_i^*(d)$, defined as follows:

$$b_i^*(d) = \begin{cases} \{d(i)\} & \text{when } d(i) \in b_i(d) \\ b_i(d) & \text{otherwise} \end{cases}$$

This rule characterizes the Darwinian component of the stochastic adjustment process, i.e. those adjustments that occur as the result of payoff-

driven decision making. The first result is that with probability 1, these adjustments lead to a stationary demographic. A stationary demographic must have the property that each player plays a best reply to the strategy profile of his current set of neighbors. Thus we have a version of the standard “stability implies Nash” result of evolutionary equilibrium selection. The proof of this and later results are in the appendix.

Proposition 1 *The Darwinian adjustment process converges with probability 1 to a convention.*

The Darwinian dynamics alone cannot select among the set of conventions. The rest point of best-response adjustment processes depends on initial conditions. Equilibrium selection is accomplished by introducing “noise” into the Darwinian process. That is, I assume that on occasion players will make mistakes, experiment, or for some other reason make their decision without regard to stage game payoff. Specifically, when a player is selected for revision, with some small probability ε , the player selects his location and strategy at random, each alternative being equally likely.

Under this assumption, the process can be analyzed using the techniques employed by KMR and Young. For any $\varepsilon > 0$, the system can be described by a regular Markov chain, i.e. one with a unique invariant distribution μ_ε . We can then analyze the behavior of the stationary distribution as the likelihood of mistakes vanishes. The advantage of this technique, as shown by KMR and Young, is that the support limiting distribution is a subset of the recurrent classes of the Darwinian dynamics, i.e. conventions. When the support is a proper subset, the perturbations have selected from the “equilibria” of the unperturbed process.

Definition 1 *Write $\mu^* := \lim_{\varepsilon \rightarrow 0} \mu_\varepsilon$. When it exists, μ^* is called the **limit distribution**. A demographic d is **stochastically stable** iff $\mu^*(d) > 0$.*

Stochastically stable demographics are those which are visited with non-vanishing relative frequency when mistakes are possible but decreasingly likely. Although ergodicity implies every demographic is visited infinitely often, those that are not stochastically stable arise infinitely less often than those that are. Naturally, we conclude that these demographics are not “robust” to the perturbation.

In the simple model of this section, a demographic is stochastically stable only if it is an efficient convention.³

³The condition on the population size is mild. There must be enough players so that if all players occupied the same location and played the same strategy, a single player’s deviation would not change the incentives of any other player.

Theorem 1 *In the simple model, $\mu^*(Z_1) = 1$ for sufficiently large $|\mathcal{I}|$.*

The proof is in the appendix, here I present the basic intuition. The support of the limit distribution μ^* can be characterized using the following techniques from Nöldeke and Samuelson (1993) and Samuelson (1994) (hereafter NS). A *single mutation path* between conventions d_0, d_n is a set of conventions $\{d_k\}_{k=0}^n$ and a sequence of single mutation transitions $\{(d_k, \tilde{d}_{k+1})\}_{k=0}^{n-1}$ such that \tilde{d}_{k+1} is in the Darwinian basin of attraction of d_{k+1} . For any set of conventions Z , let $M(Z)$ denote the set of conventions that are reachable from within Z by single mutation paths. A *connected component* is a minimal set of conventions Z such that $Z = M(Z)$. NS showed that the support of μ^* is a union of some collection of connected components.

In the simple model, no connected component can include elements of Z_2 . To see this, consider any inefficient convention $d \in Z_2$ in which there is an unoccupied neighborhood. Consider the single mutation transition in which a player mistakenly moves to an unoccupied neighborhood and adopts the efficient strategy. By a sequence of Darwinian best-replies to the resulting demographic, all other players will follow and also play the efficient strategy, since the efficient strategy is a best-reply to itself and achieves the maximum possible payoff. This is therefore a single mutation path to Z_1 .

Now suppose there is no unoccupied neighborhood in d . By a sequence of single mutation transitions, players can move from some location j to j' until eventually j is empty. From the resulting demographic, a single mutation path to Z_1 can be constructed as above. Thus, if Z^* is the subset of Z_1 in which all players are in a single location, we have $Z^* \subset M(Z)$ for any Z such that $Z \cap Z_2 \neq \emptyset$. If we can show that $Z^* = M(Z^*)$, it will follow that no such Z can be a connected component.

Consider a single mutation in a demographic $d \in Z^*$. There are two cases to consider: either the mutant moved to an unoccupied location or the mutant stayed in the original location but switched to k_2 . In the first instance, the only player with any incentive to change behavior is the newly-created loner, who will want to return to the original location and play k_1 . In the second situation, provided there are enough players in the original location, this change in the strategy of a single player cannot change the incentives of the others. The mutant of course would like to switch back to k_1 . Thus, in either case, the only basin of attraction the post-mutation demographic can belong to is d . Thus $d = M(d)$ for all $d \in Z^*$, and the result follows.

This simple model is built on a number of assumptions which provide the cleanest illustration of this theme. However, these assumptions are not

necessary for the result. The next section shows that under very general assumptions about the adjustment process and the locational structure, efficient conventions persist in the limit.

On the other hand, as discussed in the introduction, the assumption that players maximize expected payoff without regard to the number of interactions is necessary to obtain the results at the level of generality of Theorems 1 and 2. Here I shall discuss briefly a more general specification of payoffs and the effects this modification would have on the results. Define

$$\hat{\pi}_i(d) = w(|N_d(i)|)\pi_i(d)$$

where $w(\cdot)$ is some strictly increasing concave function. These payoffs would capture a situation in which players are able to interact with more than one neighbor in a given period, with a declining marginal value of each additional interaction. If $w(n) \equiv \bar{w}$, for some constant \bar{w} , then we have a model that is equivalent to the text. On the other hand, if $w(n) \equiv n$, then players are maximizing the sum of their payoffs.

In the latter case, we can see that the stochastic stability of the efficient convention depends on the payoffs in the game. In fact, we can identify necessary and sufficient conditions in this case for the efficient convention to be selected. Define λ by the equation $(1 - \lambda)a - \lambda b = 0$. A player will be indifferent between a location in which a fraction λ of the total population is playing k_2 and another location in which the remaining fraction is playing k_1 . We can extend the previous mutation counting arguments to the present model to show that the efficient convention is selected if and only if⁴ $\lambda > \tau$. To move from the efficient to the risk-dominant convention, $1 - \tau$ fraction of the population should switch actions by mutation, and the remainder will switch by best-reply. On the other hand, to move from the risk-dominant to the efficient convention, a fraction $1 - \lambda$ should, by mutation, move to an empty location and play the efficient action and the remainder will follow by best-reply.

In the general case, the more concave is $w(\cdot)$ the more favorable are the conditions for efficiency. In particular, I conjecture that if w is bounded above by w^* , then for sufficiently large populations, the efficient convention will be selected. This is because as the population size grows, the total number of mutations $(1 - \tau)|\mathcal{I}|$ within a given location necessary to make k_2 a best-reply grows without bound. On the other hand, to move from the risk-dominant to the efficient convention, it is enough to have n^* players move to an empty location and play k_1 where n^* is defined by $w(n^*)a = w^*b$.

⁴This statement is correct up to integer-sized population issues.

This condition is independent of the population size. To establish this result formally one would have to check whether there are quicker routes to the risk-dominant convention involving mutations which change the location as well as the action of some players.

A full exploration of these issues is an interesting direction for further research.

4 A More General Model

In this section I explore the robustness of the efficient convention in a general model of which the basic model in the previous section is a special case. The model is extended in two directions. First, I weaken considerably the assumptions about the structure of neighborhoods. In the previous section, each individual location was identical to the traditional uniform matching environment found in KMR. In this section, I consider a generalization of the type of neighborhood structure considered in Ellison (1993). I show that this generalization does not affect the results of the previous section.

Players again select from a finite set of locations \mathcal{J} . Now however, we associate with each location a set of neighboring locations. Each player in a given location can be matched with any other player in any other neighboring location.

Definition 2 *A neighborhood network is a symmetric and reflexive binary relation \leftrightarrow on \mathcal{J} . Location j is said to neighbor location j' iff $j \leftrightarrow j'$. A neighborhood network \leftrightarrow is **non-degenerate** if there exists $j, j' \in \mathcal{J}$ such that $j \not\leftrightarrow j'$.*

Given a neighborhood network, each demographic determines a neighborhood relation on the set of players. Write $i \leftrightarrow_d i'$ whenever $d_{\mathcal{J}}(i) \leftrightarrow d_{\mathcal{J}}(i')$. Player i 's set of potential opponents is then $N_d(i) := \{\hat{i} \neq i : i \leftrightarrow_d \hat{i}\}$. I make the assumption that each player believes that in any given period he is equally likely to be paired with each of his neighbors, and maximizes his associated expected payoff.⁵ This gives the following payoff function

$$\pi_i(d) = \frac{1}{|N_d(i)|} \sum_{\hat{i} \in N_d(i)} u(d_{\mathcal{K}}(i), d_{\mathcal{K}}(\hat{i}))$$

I continue to assume $N_d(i) = \emptyset \Rightarrow \pi_i(d) = u_r < b$.

⁵In general there may be no matching distribution in which each player is equally likely to be matched with each of his neighbors. This assumption is made for convenience.

The second modification is motivated by the following potential objection. The model of Ellison should be interpreted as a model in which while mobility may be possible, strategy revision occurs with far greater frequency than opportunities to move. This justifies assuming locations fixed while analyzing the process governed solely by strategy adjustment.

The model of this section addresses this argument by assuming a two-tiered adjustment process. As with the simple model, each period, and for each player, there is a probability $\beta > 0$ (independent across players) that the player dies and is replaced by a newborn. The newborn chooses a strategy and location to maximize his current payoff.

Once chosen, the location is fixed for the duration of the player's lifetime. However, opportunities for strategy revision may still arise. When they do, the player selects a strategy which is a best reply to his current set of neighbors. This setup is intended to capture the features of an environment in which location choices are "sticky" relative to strategy decisions.

Formally, the sequence of play is as follows. At the start of each period t , the players observe the current demographic d^t . Matching takes place and payoffs are received. Next, some players will die and be replaced by newborns who choose new locations and strategies. Among the players who survive, some will update strategies. The probability of a death is β and the probability of a revision opportunity is α , conditional on survival. As before, these probabilities are independent of history and across players.

A newborn chooses $(j, k) \in b_i^*(d^t)$. A player with a strategy revision opportunity selects a strategy $k \in a_i := \operatorname{argmax}_{\hat{k} \in \mathcal{K}} \pi_i(d^{-i}, (d_{\mathcal{J}}(i), \hat{k}))$. We assume that a player will only switch strategies to one which is a strict improvement under $\pi_i(\cdot)$. As before, I assume players' decisions are subject to "mistakes." Having settled on his intended choice, with probability ε , the player discards it and chooses at random: In the case of a death/birth this is called a *mutation*, in the case of strategy revision, a *mistake*.. Each alternative is equally likely.

These assumptions characterize the *general model*.

Theorem 2 *Let \leftrightarrow be any non-degenerate neighborhood network. For every $\alpha, \beta \in [0, 1)$, let $\mu_{\alpha\beta}^*$ be the limit distribution of the general model when the revision rate and death rate are α , and β , respectively. Then for sufficiently large $|\mathcal{I}|$, and for every $\bar{\alpha}, \bar{\beta} \in (0, 1)$*

1. $\mu_{\alpha\beta}^*(Z_1) > 0, \mu_{\alpha\beta}^*(Z_1 \cup Z_2) = 1$
2. $\lim_{\alpha \rightarrow 0} \mu_{\alpha\bar{\beta}}^*(Z_1) = \mu_{0\bar{\beta}}^*(Z_1) = 1$

3. $\lim_{\beta \rightarrow 0} \mu_{\alpha\beta}^*(Z_1) > 0$

This result shows the robustness of the results in the simple model, as well as emphasizing the special nature of the models of KMR, Young, and Ellison. First of all $\mu_{0\beta}^*(Z_1) = 1$ means that the analysis of the previous section did not depend on the exclusivity of the neighborhoods. Under any non-degenerate locational structure, efficient conventions are the only stochastically stable outcomes whenever strategies and locations are chosen simultaneously. Furthermore, this conclusion is robust. As long as the difference between revision rates and mobility rates are not too great, the simple model is a reasonable approximation.⁶

The last part of the theorem demonstrates that models without mobility are not good approximations to an environment in which players have any control, however limited, over the set of opponents they may face. When β approaches zero, so that mobility is limited, the probability of the Pareto-dominated convention approaches 1 only when the neighborhood network is degenerate.

The proof of this theorem (in the appendix) is lengthy; a sketch of the approach follows. For all $\alpha, \beta \in (0, 1)$, the limit distribution $\mu_{\alpha\beta}^*$ can be characterized using the NS techniques. In the general model, there is a unique connected component $\mathcal{R} \subset Z_1 \cup Z_2$. It follows from NS that \mathcal{R} is the support of $\mu_{\alpha\beta}^*$.

I extend the NS result to fully characterize $\mu_{\alpha\beta}^*$ using single mutation transitions. Consider all transitions between elements of \mathcal{R} that involve a single mistake/mutation. For each such $d, d' \in \mathcal{R}$, let $\tilde{P}(d, d')$ be the coefficient on ε in the (polynomial) probability of the transition ($d \rightarrow d'$). Then (after a normalization), (\mathcal{R}, \tilde{P}) forms a regular Markov chain, call it $\mathcal{M}(\alpha, \beta)$. It is regular because by definition of \mathcal{R} , between any two demographics $d, d' \in \mathcal{R}$ there exists a path consisting of single mistake/mutation events.

I show that the unique invariant distribution of $\mathcal{M}(\alpha, \beta)$ is $\mu_{\alpha\beta}^*$.⁷ The probabilities in \tilde{P} are polynomials in α and β . The limits $\lim_{\alpha \rightarrow 0} \mu_{\alpha\beta}^*$ and $\lim_{\beta \rightarrow 0} \mu_{\alpha\beta}^*$ can then be analyzed viewing $\mathcal{M}(\alpha, \beta)$ as a perturbation of $\mathcal{M}(0, \beta)$ (respectively of $\mathcal{M}(\alpha, 0)$) and applying the NS techniques.

⁶The condition on the population size is in general stronger than in the simple model. We require $[(N-2)a + e]/N - 1 > b$. In the case of the example in the introduction, any population of at least 8 is sufficient. The stronger condition is only necessary under certain special types of neighborhood networks. For example, in Ellison's circle model, we require only the weaker condition from the previous section.

⁷I am slightly oversimplifying the argument for the purposes of exposition. See proposition 7 in the appendix for a more careful statement

Part 2 of the theorem can then be established as a consequence of the standard result that the support of a limit distribution (here $\lim_{\alpha \rightarrow 0} \mu_{\alpha\beta}^*$) is a subset of the recurrent states of the unperturbed process $\mathcal{M}(0, \beta)$. The transitions in $\mathcal{M}(0, \beta)$ are those that involve only birth/death events with a single mutation. Starting from any Z_2 demographic, there is a sequence of such transitions that lead to (the basin of attraction of) Z_1 : players one-by-one drift around the neighborhood network until there is an empty neighborhood, at which point a single player moves into that neighborhood and plays k_1 . Now there exist demographics in Z_1 from which there is no such path to Z_2 ; namely those in which every player has sufficiently many neighbors. In such a demographic, a single mutation cannot induce any player to choose k_1 . Thus, the recurrent states of $\mathcal{M}(0, \beta)$ must be contained in Z_1 .

The recurrent states of the system $\mathcal{M}(\alpha, 0)$, on the other hand include elements of both Z_1 and Z_2 . Demographics in Z_1 can come about only when a neighborhood becomes empty. Demographics in Z_2 come about when the population becomes vulnerable to contagion effects. The key result is that the latter relies just as strongly on mobility as the former. In order for contagion to be effective, the population must be sufficiently dispersed throughout the network: each player must have relatively few neighbors. Because there is no mobility in $\mathcal{M}(\alpha, 0)$, Z_2 conventions with a concentrated population are recurrent, as well as the Z_1 conventions with a dispersed population.

Part 3 of the theorem is therefore shown by applying the NS techniques to the recurrent classes of $\mathcal{M}(\alpha, 0)$. In particular, it is shown that every component that is connected by single *birth/death* events includes elements of Z_1 . Again, starting from any convention, by a sequence of single mutations, a neighborhood can be emptied. Subsequently, a single player can move into the emptied neighborhood and play k_1 . NS's result now implies part 3.

It is worth noting what is required of a neighborhood network for the purposes of this result. It is clearly necessary, and by the Theorem 2 sufficient, that the network allow demographics in which there are empty neighborhoods as well as demographics in which there is substantial population density in at least some region of the network. The former is guaranteed by the assumption of non-degeneracy. The latter was necessary to conclude that there are Z_1 conventions which are immune to contagion effects, hence recurrent in $\mathcal{M}(\alpha, 0)$. This feature is built into the definition of a neighborhood network: there is no limit to the number of players that can occupy a single location.

This feature makes the definition and the arguments simpler, but it is not the only way to ensure the conclusions of Theorem 2. For example, a natural modification would be to place a limit on the number of players (say, at most 1) that could occupy a given location. Theorem 2 would continue to hold in such a model under the additional assumption that there was at least some region in which locations were tightly networked (i. e. each location had many neighboring locations).

Finally, notice that Theorem 2 says only that $\mu_{\alpha\beta}^*(Z_1) > 0$, whereas in the simple model we had $\mu_{\alpha\beta}^*(Z_1) = 1$. It may be possible to obtain the stronger conclusion in a wider set of neighborhood networks. In particular, the proof of Theorem 2 makes it clear that all stochastically stable demographics would be in Z_1 if it were not possible for a single k_2 -mistake to infect the entire population. In particular, $\mu_{\alpha\beta}^*(Z_1) = 1$ in any neighborhood network in which it is impossible to arrange the players in such a way that this kind of infection can occur. In fact, this conclusion could be strengthened further. In order for Z_2 to be stochastically stable, we it is necessary that starting from *any* Z_1 demographic, such an infection-vulnerable demographic can arise by a sequence of single-mutation events. This condition will not be satisfied if there is at least one isolated region of the neighborhood network which is invulnerable to contagion effects that result from a single mutation. The k_1 convention in which all players are located in this region in particular would be invulnerable to single-mutation infection.

5 Conclusion

I have shown that incorporating local interaction and mobility in the standard framework of random matching and stochastic adjustment has important consequences for the determination of long-run equilibria.

When compared with the original evolution literature, in which variants of ESS were applied to the equilibrium selection problem, these results are not surprising. These solution concepts essentially assumed that the population was arbitrarily large and that “mutations” affected an arbitrarily small component of the population. With such a hypothesis, every strict equilibrium is evolutionarily stable. In order to overturn a strict equilibrium, some discrete-sized deviation is necessary.

Kim and Sobel (1995), and Matsui (1991), however, showed that when a round of pre-play communication is added, even the very weak test of stability inherent in ESS and its variants is sufficient to generate strong predictions. In fact, since in the communication game there are no strict

equilibria, it was necessary to weaken ESS in order to guarantee existence. An interpretation of these results is that when players have a mechanism by which to identify their preferred opponents, catalytic events occur arbitrarily more frequently than is possible in the unmodified game.

This is the logic underlying the results in this paper. I have shown, in parallel with the communication literature, that location and mobility give players a mechanism by which they can identify their preferred opponents. This creates a force for equilibrium selection based on in-equilibrium payoffs in contrast to the criterion of risk-dominance which places equal emphasis on out-of-equilibrium outcomes. For this reason, locational choice tends to favor efficient equilibria.

Acknowledgments

This research was supported by an Alfred P. Sloan Doctoral Dissertation Fellowship. Thanks to my adviser Matthew Rabin for helpful conversations, suggestions, and encouragement. Bob Anderson, Bob Powell, Suzanne Scotchmer, and Chris Shannon also provided helpful comments. All mistakes (in this paper) are mine.

A Proofs

Throughout, let $\mathcal{M}(\varepsilon)$ denote the Markov chain corresponding to the process with mutation/mistake probability ε . In the terminology of Young (1993), the family of chains $\{\mathcal{M}(\varepsilon) : \varepsilon > 0\}$ is a *regular perturbation* of the unperturbed process $\mathcal{M}(0)$. (For $\mathcal{M}(\varepsilon)$ to be a regular perturbation, it is sufficient that the transitions be polynomials in ε .) Young has shown that a regular perturbation has a limit distribution $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu_\varepsilon$, and that μ^* is an invariant distribution of $\mathcal{M}(0)$.

All transition probabilities in the models in this paper are polynomials in ε . For any transition (d, d') , let $C(d, d')$ be the degree of the associated polynomial transition probability. I will refer to it as the *cost* of the transition. Let $c(d, d')$ be the leading coefficient, i.e. the coefficient on $\varepsilon^{C(d, d')}$.

To characterize the limit distribution of a regular perturbation, we use the graph-theoretic technique introduced by Freidlin and Wentzell (1984), and further developed by The approach used here borrows most heavily from the techniques used in Samuelson (1994) and Nöldeke and Samuelson (1993).

A *graph* G is a pair (\mathcal{D}, E) where \mathcal{D} is a finite set, called the *vertices* of G , and E is a collection of ordered pairs of \mathcal{D} , called the *edges* of G .

A *path* is a finite sequence of edges $\{e_i\}, i = (1, \dots, n)$ whose origins are distinct and such that for every $1 \leq i < n$, the endpoint of e_i is the origin of e_{i+1} . The origin and endpoint of a path are defined in the obvious way. Write $d_2 \succeq_G d_1$ to indicate that there is a path in G from d_1 to d_2 . When S and T are subsets of \mathcal{D} , write $T \succeq_G S$ to indicate that for every $s \in S$, there is a $t \in T$ such that $t \succeq_G s$.

A spanning tree h of a graph G is a subset of edges satisfying the following conditions for some fixed vertex r (called the *root*):

- Each $d \in \mathcal{D} \setminus r$ is the origin of exactly one edge in h .
- In the graph $t = (V, h)$, $r \succeq_t d$ for every $d \in \mathcal{D} \setminus r$.

For each $d \in \mathcal{D}$, let H_d be the set of spanning trees in G whose root is d .

Suppose \mathcal{M} is some regular Markov chain on a state space \mathcal{D} , with transition matrix P . The graph derived from \mathcal{M} is $G = (\mathcal{D}, E)$ where the edges are taken to be the set of all positive probability transitions in \mathcal{M} , i.e. $E = \{e = (d_1, d_2) : P(e) > 0\}$.

Freidlin and Wentzell (1984) characterized the invariant distribution μ of a regular Markov chain in terms of its graph. I state below versions of these results for *conditional* distributions $\mu(\cdot | X)$ for $X \subset \mathcal{D}$.

Theorem 3 (Freidlin and Wentzell, 1984) For every $X \subset \mathcal{D}$, $d \in X$, if $\mu(X) > 0$ then

$$\mu(d|X) = \frac{q_d}{\sum_{d' \in X} q_{d'}}$$

where $q_d = \sum_{h \in H_d} \prod_{e \in h} P(e)$

For any spanning tree h , write $C(h) = \sum_{e \in h} C(e)$. For each $X \subset \mathcal{D}$, $d \in \mathcal{D}$ define the sets $H_d^X = \{h' \in H_d : C(h') = \min_{\substack{h \in H_d \\ d \in X}} C(h)\}$.

Corollary 1 Let $\mathcal{M}(\varepsilon)$ be a regular perturbed Markov chain on \mathcal{D} , and G its associated graph. Then if $d \in X$,⁸

$$\mu^*(d|X) = \frac{\sum_{h \in H_d^X} \prod_{e \in h} c(e)}{\sum_{\hat{d} \in X} \sum_{h \in H_{\hat{d}}^X} \prod_{e \in h} c(e)}$$

Remark: When X is a proper subset of \mathcal{D} , $\mu^*(\cdot|X)$ should be understood to be the limit of the conditional distributions of the regular perturbation, as opposed to the conditional of the limit distribution. The latter would be undefined in the case that $\mu^*(X) = 0$. In all of the models of this paper, μ_ε has full support for every $\varepsilon > 0$, and hence the limiting conditional distributions are always well-defined.

Corollary 2 Let $d \in X$. Then $\mu^*(d|X) > 0$ only if $H_d^X \neq \emptyset$.

Let G be any graph. Viewing \succeq_G as a binary relation, we can define $\mathcal{A}(G)$ the collection of \succeq_G -maximal subsets of \mathcal{D} , called the *absorbing sets* of G . Formally $A \in \mathcal{A}(G)$ iff for each $d \in A$, $d' \succeq_G v \iff d' \in A$. For $A \in \mathcal{A}(G)$, the set $D(A) := \{d : A \succeq d\}$ is called the *domain* of A .

Now suppose G is derived from a regular perturbed Markov chain $\mathcal{M}(\varepsilon)$ on the state space \mathcal{D} . We will derive two graphs from G which will be used in characterizing the limit distribution of $\mathcal{M}(\varepsilon)$. Consider the graph $G_0 = (\mathcal{D}_0, E_0)$, where $\mathcal{D}_0 = \mathcal{D}$ and $E_0 := \{e \in E : C(e) = 0\}$. This graph represents all transitions which occur with positive probability in the underlying, unperturbed chain $\mathcal{M}(0)$. In the models in this paper, these are all of the Darwinian transitions: those that involve only changes in behavior motivated by best-response.

We now define a second graph G_1 , taking the vertex set \mathcal{D}_1 to be the collection $\mathcal{A}_0 = \mathcal{A}(G_0)$ of absorbing sets of G_0 . Define the edges E_1 as follows: $(A_1, A_2) \in E_1$ iff $\exists d \in A_1, d' \in \mathcal{D}$ such that

⁸Adopt the convention that the sum of the empty set of numbers is zero.

- i* . $d' \in D_0(A_2)$, and
- ii* . $C(d, d') = 1$.

where $D_0(A_2)$ is the domain of A_2 in graph G_0 . Throughout subscripts will be used to distinguish objects belonging to G_0 from those in G_1 . I will now state one of the main results to be used in characterizing the support of the limit distribution.

Proposition 2 *Let $d_1, d_2 \in \mathcal{D}_l$. Suppose $X \cap (d_1 \cup d_2) \neq \emptyset$. Then*

1. *If $d_2 \succeq_l d_1$, then $\mu^*(d_1|X) > 0 \Rightarrow \mu^*(d_2|X) > 0$.*
2. *If in addition $d_1 \not\preceq_l d_2$ then $\mu^*(d_1|X) = 0$.*

The version for unconditional distributions, i.e. the case of $X = \mathcal{D}$ is a straightforward implication of the results in NS. The logic for conditional distributions is identical and hence the proof is omitted here.

We now turn to the proofs of Theorems 1 and 2.

Proposition 3 $\mathcal{A}_0 = \cup_n Z_n$.

Remark: This proposition implies proposition 1.

Proof: In order to subsume the analogous result for the general model, we will prove this proposition for the case of an arbitrary network neighborhood on \mathcal{J} .

I first show that if $d \in Z_n$, then $d \in \mathcal{A}_0$. For this it is sufficient to show that $b_i^*(d) = d(i)$ for every i . Using the definition of Z_n , we have

$$\begin{aligned} \max_{j,k} \pi_i(d^{-i}, (j, k)) &= \max[u_r, \max_k u(k, k_n)] \\ &= \max[u_r, u(k_n, k_n)] \\ &= u(k_n, k_n) \\ &= \pi_i(d) \end{aligned}$$

Thus, $b_i^*(d) = d(i) \forall i$ implying $d \in \mathcal{A}_0$. Now to establish the proposition, it suffices to show that $d \notin Z_1 \cup Z_2$ implies $\exists n$ and $\hat{d} \in Z_n$ such that $\hat{d} \succeq_0 d$.

If for every i , there exists j such that $(j, k_2) \in b_i^*(d)$, then we can construct a path to Z_2 as follows: Have all players update in one round. They will all choose strategy k_2 in various locations. In the resulting state, move any newly created loners to a single occupied location. The result is a demographic in Z_2 .

Now assume there exists a player i and a location j such that $(j, k_1) \in b_i^*(d)$. Then letting $\hat{d} := (d^{-i}, (j, k_1))$, we have $\hat{d} \succeq_0 d$ and

$$\pi_i(\hat{d}) \geq \pi_i(d^{-i}, (\hat{j}, k_2)) \quad \text{for all } \hat{j} \in \mathcal{J} \quad (1)$$

Now, for every player l and location m , let z_{lm} represent the fraction of k_1 players l would face were he to move to location m starting from demographic \hat{d} . Since i occupies j in \hat{d} , the value z_{ij} represents the actual fraction of k_1 opponents i faces, hence the fact that i chose k_1 implies $z_{ij} \geq \tau$.

If it is the case that all players who play k_2 in \hat{d} wish to switch to k_1 , then we could construct a path to Z_1 by allowing only those players to switch to k_1 , and then relocating loners. Thus, the proposition follows if we can show that if there is any player l who plays k_2 in \hat{d} , then there is no location m such that $(m, k_2) \in b_l^*(\hat{d})$. This will be proved by contradiction.

If conditional on locating in m , l prefers strategy k_2 , it must be that $z_{lm} \leq \tau$. We can draw an inference from l 's revealed preference of (m, k_2) over (j, k_2) .

$$\begin{aligned} \pi_l(\hat{d}^{-l}, (m, k_2)) - \pi_l(\hat{d}^{-l}, (j, k_2)) &\geq 0 \Rightarrow \\ [z_{lj} - z_{lm}] (b - f) &\geq 0 \end{aligned} \quad (2)$$

Were player l to move to location j , he would face a strictly larger fraction of k_1 playing opponents than does i in \hat{d} . (This is because l is playing k_2 whereas i is playing k_1 .) Thus, $z_{lj} > z_{ij} \geq \tau \geq z_{lm}$. Together with equation (2), this implies $b \geq f$. (A k_2 player has a weak preference for neighbors who play k_2 .)

We have assumed that $a > b$ and that both equilibria of the underlying game are strict. Thus $a > b > e$. Now $(m, k_2) \in b_l^*(\hat{d})$ implies $\pi_l(\hat{d}^{-l}, (m, k_2)) \geq \pi_l(\hat{d}^{-l}, (j, k_1))$. Combining these facts,

$$\begin{aligned} \pi_l(\hat{d}^{-l}, (m, k_2)) &\geq z_{lj}a + (1 - z_{lj})e \\ &> z_{ij}a + (1 - z_{ij})e \\ &= \pi_i(\hat{d}) \end{aligned} \quad (3)$$

where the inequality follows from $z_{lj} > z_{ij}$ and $a > e$.

Finally, since in \hat{d} player l plays k_2 and player i plays k_1 in \hat{d} , player l will face no smaller a fraction of k_1 -playing opponents in m than would player i , i.e. $z_{lm} \geq z_{im}$. Noting that \hat{d} differs from d only in the behavior of player i , this yields

$$\begin{aligned}
 \pi_i(d^{-i}, (m, k_2)) &= \pi_i(\hat{d}^{-i}, (m, k_2)) \\
 &= z_{im}f + (1 - z_{im})b \\
 &\geq z_{lm}f + (1 - z_{lm})b \\
 &= \pi_l(\hat{d}^{-l}, (m, k_2))
 \end{aligned} \tag{4}$$

But this last inequality is in contradiction with inequalities (1) and (3).

■

Proof of Theorem 1: In view of Proposition 1, it suffices by Proposition 2 to show that $Z_1 \succeq_1 Z_2$ and that $Z_2 \not\succeq_1 d$ for every $d \in Z_1$.

Let $d \in Z_2$, and suppose there are no empty neighborhoods in d . Then since $|\mathcal{J}| > 1$, there are at least 2 neighborhoods j_1, j_2 which are occupied. I will show that there exists a $d' \in Z_2$ such that j_1 is empty in d' and $d' \succeq_1 d$.

A movement of a single player from j_1 to j_2 , holding the player's strategy fixed, involves a single mutation. If in the resulting state d_1 there remain at least 2 players in j_1 , then $d_1 \in Z_2$ and $(d, d_1) \in E_1$. We can then repeat this process, constructing a sequence of edges $(d_t, d_{t+1}) \in E_1$ resulting in a state d_n in which there is a single player occupying j_1 . Because this player is a loner, his best-reply is to move to an occupied location and play k_2 . Thus, if d' is the state that results from the loner's movement to j_2 , we have $d_n \in D_0(d')$ and hence $(d_{n-1}, d') \in E_1$.

Now starting from d' , as a result of a single mutation, a player can move to j_1 and begin playing k_1 . From the resulting state d'' , the remaining players' best replies are to move to j_1 and play k_1 . Letting d''' be the state which results when all players do so, we have $d''' \in Z_1$, $d'' \in D_0(d''')$ and hence $(d', d''') \in E_1$. Since d was an arbitrary element of Z_2 , we have shown $Z_1 \succeq_1 Z_2$.

To complete the proof I will show that for every $d \in Z_1$, if d' is any state that results from a single mutation in state d , then d' is not in the domain of any state in Z_2 , and hence there is no edge $(d, \tilde{d}) \in E_1$ where $\tilde{d} \in Z_2$.

Suppose d' results from a mutation which moves a player to a previously unoccupied neighborhood. Then regardless of the mutant's strategy, no other player has any incentive to change his strategy. The only Darwinian adjustments that can take place after such a mutation involve loners (the mutant and any loner he might have left behind) finding new neighborhoods in which to play k_1 .

Suppose instead that the mutation involves a player changing strategies to k_2 within an occupied neighborhood j . (Without a change of strategies, d' would be another element of Z_1 .) If there is some other neighborhood j' ,

and there are no loners, then all players in j will wish to move to j' and play k_1 against a population of k_1 players. Thus all Darwinian adjustments involve players choosing k_1 and these can only lead to a new element of Z_1 .

On the other hand, there may be no other occupied neighborhood, or the only other occupied neighborhood could contain a loner left behind by the mutant. In that case a sufficient condition for all players to wish to update to k_1 is that the fraction of k_1 players in j be sufficiently large so that the best reply in j is also k_1 . Since there are at least $|\mathcal{I}| - 1$ players in j and only one is playing k_2 , this requirement is that $1/(|\mathcal{I}| - 2) < 1 - \tau$. This is satisfied provided the population is sufficiently large. ■

We now turn to the general model.

Proposition 4 $\cup_n Z_n = \mathcal{A}_0$

Proof: The edges in G_0 derived from the general model compose a superset of those derived from the simple model. Thus, since Proposition 1 was proved for an arbitrary neighborhood network, the result that $\cup Z_n \succeq_0 [\mathcal{D} \setminus \cup Z_n]$ is subsumed. Therefore we need only show that elements of $\cup Z_n$ are singleton maximal sets under \succeq_0 in the general model. Proposition 1 showed that Darwinian deaths could not alter an element d of $\cup Z_n$. Clearly neither can a Darwinian revision since all players play the same strategy in d . ■

Proof of Theorem 2 To show $\beta > 0 \Rightarrow \mu_{\alpha\beta}^*(Z_1) > 0$ it is sufficient to show that $Z_1 \succeq_1 Z_2$ whenever $\beta > 0$. Let $d \in Z_2$. By non-degeneracy, there exist two locations j_1, j_2 such that $j_1 \not\sim j_2$. Let $M(j)$ represent the set of locations which neighbor j , and $N_d(j)$ the set of players occupying locations in $M(j)$. If $N_d(j_1), N_d(j_2) \neq \emptyset$, then we will first construct a path in G_1 from d to an element of Z_2 in which j_1 has an empty neighborhood.

By a single mutation in d , a player can move from any $j \in M(j_1)$ to j_2 . Following this move, any loners will be indifferent among all locations available. With positive probability they choose a location outside of $M(j_1)$ and hence if d_1 is any state in which all loners move outside of $M(j_1)$, we have $d_1 \in Z_2$ and $(d, d_1) \in E_1$, and there are strictly fewer players occupying $M(j_1)$ in d_1 than in d . Continuing in this way we can construct a sequence of edges $(d_t, d_{t+1}) \in E_1$ resulting in a state d_n in which $N(d_n) = \emptyset$.

Now starting from d_n , following a mutation, a player can move to j_1 and play k_1 , leading to a state d' . Each remaining player wishes to follow suit, and hence if d'' is the state that results when all players do so, we have $d' \in D_0(d'')$ implying $(d_n, d'') \in E_1$. The state d'' is an element of Z_1 and we have thus shown that $Z_1 \succeq_1 Z_2$.

We now turn to part 2 of the Theorem. Because the proof that $Z_1 \succeq_1 Z_2$ holds for any value of α , to prove that $\mu_{0\beta}^*(Z_1) = 1$ it suffices to show that if $\alpha = 0$, then $Z_2 \not\succeq_1 d$ for any $d \in Z_1$.

Consider any d' which occurs as a result of a single mutation starting at d . Suppose in d' that for every player there is some location where that player can move and earn a payoff of a by playing k_1 . Then in any Darwinian successor of d' , every player's best reply involves playing k_1 . This can only lead to an element of Z_1 .

Now suppose in d' there is some player who cannot achieve a payoff of a in any location. This can only happen if the mutant is playing k_2 and he cannot be avoided: he is a neighbor of every occupied location. If a player were to move to the location occupied by the mutant, he would also neighbor every player. His payoff to playing k_1 would then be $[(|\mathcal{I}| - 2)a + e]/(|\mathcal{I}| - 1)$. For $|\mathcal{I}|$ sufficiently large, this is greater than b , the maximum payoff attainable by playing k_2 . Again, we conclude that k_2 is never a best-response in any Darwinian successor to d' . Thus $Z_2 \not\succeq_1 d$.

To prove the limiting results of the theorem, we make use of some additional notation and prove some intermediate results. For any edge $e = (d_1, d_2)$ of G_1 , a *projection* of e onto G is a path in G from d_1 to d_2 with total cost equal to 1. Let \tilde{E} be the set of all edges of G_0 together with all projections of edges of G_1 . Write $\mathcal{A}_1 = \mathcal{A}(G_1)$, and denote by \mathcal{R} the set of states which are stochastically stable when $\alpha, \beta > 0$.

Proposition 5 *If \mathcal{A}_1 is a singleton, then for every $d \in \mathcal{R}$, $h \in H_d^D$ iff $h \subset \tilde{E}$.*

Proof: When \mathcal{A}_1 is a singleton A , then $\mathcal{R} = A$. We first show that for every $d \in \mathcal{R}$, the set of spanning trees h rooted at d such that $h \subset \tilde{E}$ is not empty. First of all, since it has a unique absorbing set, the graph G_1 contains a spanning tree \tilde{h} rooted at any $\tilde{d} \in \mathcal{A}_1$. If $(d, d') \in \tilde{h}$, then there is a $(d, d'') \in \tilde{E}$ such that $d' \succeq_0 d''$. From this we can construct a “skeleton” tree of G consisting of the projections of all the edges in \tilde{h} . The root is \tilde{d} , the tree is a subset of \tilde{E} , and it contains a path to \tilde{d} from every element of \mathcal{A}_0 . For each remaining element, there is a path in G_0 to some element of \mathcal{A}_0 , and so we can complete the spanning tree by adding in these paths. Such paths consist of edges in $E_0 \subset \tilde{E}$, hence the constructed tree is a subset of \tilde{E} .

Now consider the set of all spanning trees h rooted at some element of \mathcal{R} . Since every element of \mathcal{A}_0 is maximal with respect to \succeq_0 , any edge originating in such an element must have cost at least 1. Thus, $|\mathcal{A}_0| - 1$ is

the minimum total cost of any such tree. If $h \subset \tilde{E}$ each element of \mathcal{A}_0 is the origin of a single-cost edge, while all other edges cost zero. Thus h achieves this minimum. Now suppose h has cost $|\mathcal{A}_0| - 1$. Then since the cost of any path between two elements of \mathcal{A}_0 is at least 1, the cost must be exactly 1 since there are $|\mathcal{A}_0| - 1$ many of them. But this implies that these paths are projections of edges in E_1 and hence that $h \subset \tilde{E}$.

I have shown that a lower-bound on tree costs is achieved if and only if the tree is a subset of \tilde{E} . Since such trees exist, the result is proved. ■

For every $d \in \mathcal{R}$, let \tilde{H}_d be the nonempty set of spanning trees h rooted at d such that $h \subset \tilde{E}$. Proposition 5 states that this set is exactly the set of cost minimizing spanning trees whenever \mathcal{A}_1 is a singleton. The following result implies that in the general model we can assume wlog that \mathcal{A}_1 is a singleton and hence by Proposition 2, that $\mathcal{R} = \mathcal{A}_1$.

Proposition 6 *If $\mu_{\alpha\beta}^*(Z_1) \neq 1$ for some $\beta > 0$, then \mathcal{A}_1 is a singleton.*

Suppose there is a unique *component* of neighborhoods so that for every $j_0, j_n \in \mathcal{J}$, there is a sequence j_1, \dots, j_{n-1} such that $j_l \leftrightarrow j_{l+1}$ for $l = 1, \dots, n - 1$. Then starting from any $d \in Z_n$, by a sequence of mutations in which, one player at a time, players change locations but not strategies, we can construct a path to any $d' \in Z_n$.

Since we have already shown that $Z_1 \succeq_1 d'$ for every $d' \in Z_2$, it now follows that any $d \in Z_1$ is reachable in G_1 from any $d' \in Z_1 \cup Z_2$. The proposition then follows immediately.

Generalization to the case of more than one locational component is straightforward but notationally tedious and is omitted here. ■

Henceforth we will assume $\mathcal{A}_1 = \mathcal{R}$. Consider the graph $\tilde{G} := (\mathcal{D}, \tilde{E})$. Define “transition probabilities” \tilde{P} of edges $e \in \tilde{E}$ as follows

$$\tilde{P}(e) = c_e \varepsilon^{C_e}$$

For ε sufficiently small, $\tilde{P}(e) \leq 1$. If we then set $\tilde{P}(d, d) = 1 - \sum_{\hat{d}: (d, \hat{d}) \in \tilde{E}} \tilde{P}(d, \hat{d})$ and $\tilde{P}(d, d') = 0$ for $d \neq d', (d, d') \notin \tilde{E}$ we obtain a well-defined transition matrix on $\mathcal{D} \times \mathcal{D}$.

For every $\alpha, \beta \in (0, 1)$, denote by $\tilde{\mathcal{M}}_{\alpha\beta}$ the Markov chain with transition matrix \tilde{P} when the revision rate and death rate are α and β , respectively.

Proposition 7 *For every $\alpha, \beta \in (0, 1)$, $\tilde{\mathcal{M}}_{\alpha\beta}$ is regular. Let $\tilde{\mu}_{\alpha\beta}$ be its invariant distribution.*

- For every $d \in \mathcal{R}$, and for every $\alpha, \beta > 0$,

$$\mu_{\alpha\beta}^*(d) = \tilde{\mu}_{\alpha\beta}(d|\mathcal{R})$$

- For fixed α (resp. β), $\tilde{\mathcal{M}}_{\alpha\beta}$ is a regular perturbation of $\tilde{\mathcal{M}}_{\alpha 0}$ (resp. $\tilde{\mathcal{M}}_{0\beta}$) with noise parameter β (resp. α).

Proof: First notice that by construction, whenever α, β are greater than 0, the graph of $\mathcal{M}_{\alpha\beta}$ is (\mathcal{D}, \tilde{E}) . Sufficient for regularity is that there exist some state d such that the graph (\mathcal{D}, \tilde{E}) contains a spanning tree rooted at d (irreducibility), and that $\tilde{P}(d, d) > 0$ (aperiodicity). (See Romanovsky (1970, pg 53, Theorem 14.I). These are guaranteed by construction.

By Theorem 3, and the definition of $\tilde{\mathcal{M}}_{\alpha\beta}$,

$$\tilde{\mu}_{\alpha\beta}(d|\mathcal{R}) = \frac{\sum_{h \in \tilde{H}_d} \prod_{e \in h} c_e \varepsilon^{C_e}}{\sum_{\hat{d} \in \mathcal{R}} \sum_{h \in \tilde{H}_{\hat{d}}} \prod_{e \in h} c_e \varepsilon^{C_e}}$$

If t is the minimum cost among spanning trees of G , then by Proposition 5 t is achieved by all $h \in \tilde{H}_{\hat{d}}$ for every $\hat{d} \in \mathcal{R}$. Since $\prod_{e \in h} \varepsilon_e^C = C(h)$, we can factor the right-hand side to obtain

$$\frac{\varepsilon^t \sum_{h \in \tilde{H}_d} \prod_{e \in h} c_e}{\varepsilon^t \sum_{\hat{d} \in \mathcal{R}} \sum_{h \in \tilde{H}_{\hat{d}}} \prod_{e \in h} c_e}$$

and by Corollary 1, this is just $\mu_{\alpha\beta}^*(d)$.

The second claim follows from the fact that the transition probabilities are all polynomials in α and β . ■

We can thus define cost functions C^α , and C^β on \tilde{E} in the same way as $C(\cdot, \cdot)$ was defined. These cost functions simply count the number of revisions and deaths, respectively, involved in transitions in \tilde{E} . In order to calculate, say, $\lim_{\alpha \rightarrow 0} \mu_{\alpha\beta}^*$, we can evaluate the limit conditional distribution $\lim_{\alpha \rightarrow 0} \tilde{\mu}_{\alpha\beta}(\cdot|\mathcal{R})$. For this we can apply Proposition 2. We take the graph of $\tilde{\mathcal{M}}_{\alpha\beta}$ which by construction is \tilde{G} and decompose it into graphs \tilde{G}_0 and \tilde{G}_1 analogously to G_0 and G_1 for G . We use the cost function C^α to define the edges of these graphs. The successor relations will be denoted, e.g. $\tilde{\succeq}_0$, when derived from \tilde{G}_0 .

The paths constructed in the proof that $Z_1 \succeq_1 Z_2$ (which are thus subsets of \tilde{E}) consisted only of deaths and therefore have C^α -cost zero. Hence, that argument implies $Z_1 \tilde{\succeq}_0 Z_2$. Furthermore, the proof that $\mu_{0\beta}^*(Z_1) = 1$ involved showing that there was no path in \tilde{E} from Z_2 to Z_1 involving

only deaths. Therefore $Z_2 \not\stackrel{\sim}{\subseteq}_0 Z_1$ and by Proposition 2, the limit distribution of $\tilde{\mathcal{M}}_{\alpha\beta}$ assigns probability 1 to Z_1 . By Proposition 7, we conclude $\lim_{\alpha \rightarrow 0} \mu_{\alpha\beta}^*(Z_1) = 1$.⁹

What remains is part 3 of the Theorem. To prove it, we analyze the graphs \tilde{G}_0 , and \tilde{G}_1 now together with the cost function C^β .

We begin by characterizing $\tilde{\mathcal{A}}_0 = \mathcal{A}(\tilde{G}_0)$. Unlike the previous cases, it is possible that an element of $\tilde{\mathcal{A}}_0$ be non-singleton, i.e. an absorbing *set*. A typical absorbing set will be represented by A . When A is a singleton $\{d\}$, we will identify A with its element d .

The edges in \tilde{G}_0 will consist of all edges in \tilde{E} which have zero C^β -cost. For those originating in states not in $Z_1 \cup Z_2$, these will be transitions involving Darwinian strategy revision. Therefore, for $d \notin Z_1 \cup Z_2$, if all players are playing best-responses against their neighborhoods, i.e. $d_{\mathcal{K}}(i) \in a_i(d)$, then $d \in \tilde{\mathcal{A}}_0$.

A transition away from some d in $Z_1 \cup Z_2$ will be in \tilde{G}_0 only if it consists of a single mistake by a single player. Consider the set of all demographics which occur as the result of a single mistake in d . If in each of these, the only player not playing a best reply to his neighborhood is the player who made the mistake, then d is the unique successor in \tilde{G}_0 of each of these, and collectively they form an absorbing set in $\tilde{\mathcal{A}}_0$ containing d .

Consider a demographic d in which all players occupy the same location j and play strategy k_1 . This demographic is in Z_1 , and for a sufficiently large population, no single strategy deviation can alter the incentives of any player. Thus, there is an $A \in \tilde{\mathcal{A}}_0$ containing d . Let Z_1^* be the collection of sets in $\tilde{\mathcal{A}}_0$ which intersect Z_1 .

Lemma 1 *Suppose $A \in \tilde{\mathcal{A}}_0$ contains some $d \in Z_2$ such that there is a location j with an empty neighborhood. Then $Z_1^* \stackrel{\sim}{\subseteq}_1 A$.*

Proof: A mutation by which a player moves to j and adopts strategy k_1 has both C-cost and C^β -cost equal to 1. Such a transition is in \tilde{E} because $d \in Z_2 \subset \mathcal{A}_0$. It is therefore also an edge of \tilde{G}_1 . The resulting state is in $\tilde{\mathcal{A}}_0$ because all players play a best-response to their neighborhoods. Following such a state we can find a path consisting of single C^β -cost transitions in which players move one at a time to j and play k_1 . In each state along the

⁹Alternatively we could show that $\mu_{0\beta}^*$ is the invariant distribution of $\tilde{\mathcal{M}}_{0\beta}$. It is well known that the limit distribution of a regular perturbation is a subset of the invariant distributions of the underlying chain. In particular, since $\tilde{\mathcal{M}}_{0\beta}$ is regular, its invariant distribution is unique and hence the limit distribution $\lim_{\alpha \rightarrow 0} \mu_{\alpha\beta}^*$ of the regular perturbation $\tilde{\mathcal{M}}_{\alpha\beta}$ must be $\mu_{0\beta}^*$.

way, all players play best responses to their neighborhoods, hence each state is in $\tilde{\mathcal{A}}_0$. Therefore, this is a path in \tilde{G}_1 . In the endpoint state, all players play k_1 , and each neighbors every other, hence the endpoint is an element of Z_1^* . ■

A demographic d is a *star* if there is exactly one location j (called the *hub*) which is occupied by more than one player, and if the players in j are the only neighbors of all players not in j .

Lemma 2 *Let d be a star in which all players play k_1 , and d' the star in which players occupy the same locations as in d but all play k_2 . Then either $d \in A$ for some $A \in \tilde{\mathcal{A}}_0$ or $d' \in \tilde{\mathcal{A}}_0$ and $d' \tilde{\succeq}_0 d$.*

Proof: Let j be the hub. The demographic d' must be in $\tilde{\mathcal{A}}_0$ because it is in Z_2 and every player has at least 2 opponents.

If d is not in A for any $A \in \tilde{\mathcal{A}}_0$, then following a single mistaken switch to k_2 , there is a sequence of strategy revisions leading to a state \hat{d} from which there is no path in \tilde{G}_0 back to d . Assume wlog that \hat{d} has the fewest k_1 players among successors of d which have this property. Suppose no occupant of j plays k_2 in \hat{d} . Then every opponent of every player is playing k_1 , hence each wishes to play k_1 when the opportunity arises. This implies $d \succeq \hat{d}$, a contradiction. Thus, some player i in j is playing k_2 . Moreover, i must be playing k_2 as a best-reply, else i would revert to k_1 if given the opportunity leading to a state with fewer k_2 players, contradicting the definition of \hat{d} . All other players in j face at least as many k_2 -playing opponents as i does (since i plays k_2), hence must also wish to switch to k_2 . After they do so, all other player will follow suit as their only opponents are the players in j . The resulting state is d' . ■

Define the following subset of demographics.

$$Z_3 := \{d \in \tilde{\mathcal{A}}_0 : \exists d' \in Z_2, d' \neq d \text{ s.t. } d \tilde{\succeq}_0 d'\}$$

Lemma 3 $Z_1^* \tilde{\succeq}_1 Z_3$

Proof: If $d \in Z_3$ then there is a $d' \in Z_2$ such that $d \tilde{\succeq}_0 d'$. Let I be the set of players playing k_1 in d . Since paths in \tilde{G}_0 represent only strategy revisions, d' does not differ from d in terms of player locations. Thus, $I \neq \emptyset$ else $d = d'$, a contradiction.

The path from d' to d is in \tilde{E} and therefore is initiated by a single mistaken adoption of k_1 by some player i . Since k_1 is not risk-dominant, all players who were induced to switch to k_1 in response to i 's mistake must have i as their unique opponent. (More than half of a player's opponents

must play a risk-dominated strategy in order to induce that player to adopt it.) We must also have $i \in I$, else each member of I would face a unique opponent who plays k_2 , and thus would switch to k_2 when given the chance, and this would lead back to d' , a contradiction. Finally, there must be no player $i' \notin I$ who has i as a unique opponent, else i' would eventually switch to k_1 , contradicting $d \in \tilde{\mathcal{A}}_0$. Thus, I consists of i together with all players whose unique opponent is i .

Consider first the case that $I = \mathcal{I}$. Then every player is playing k_1 , hence $d \in Z_1$ and all except i have i as a unique opponent. In this case, a mistaken choice of k_2 by player i would induce all other players to adopt k_2 , returning to demographic d' . This is a single C-cost path leading from $d_0 \in \mathcal{A}_0$ to $d' \in \mathcal{A}_0$, hence it is in \tilde{E} . Furthermore, it has zero C^β -cost, hence it is a path in \tilde{G}_0 . But this is a contradiction since there is no path in \tilde{G}_0 from d to d' . Thus $I \neq \mathcal{I}$.

Suppose that some player in I is earning a payoff less than a . That player must be player i , since all other members of I face a unique opponent who plays k_1 . Player i must have a neighbor who is playing k_2 . He would like to move to any location j occupied by some other member of I , where he could obtain a payoff of a . Such a transition has C^β -cost 1.

If, on the other hand, all players in I are earning a payoff of a in d , then i 's neighborhood is just the set I . This implies that all players not in I can obtain a payoff of a by moving to the location j occupied by player i .

In either case, we reach a demographic in which a proper subset of the population earns a payoff of a . Among the players earning a , there is a player i who is the unique neighbor of all other players earning a . All players not earning a would like to move to j , the location of i , and play k_1 .

The path in which each player does so, one at a time, consists of transitions which have zero C-cost and C^β -cost 1. The endpoint d_1 is an element of Z_1 in which the set of opponents of every player not occupying j is just the set of players occupying j , hence a star. Moreover, it has an empty neighborhood: any neighborhood previously occupied by a player who moved to j . Finally, by Lemma 2 either $d_1 \in A$ for some $A \in \tilde{\mathcal{A}}_0$, in which case we have shown $Z_1^* \succeq_1 d_1$, or the corresponding k_2 star is in $\tilde{\mathcal{A}}_0$, succeeds d_1 (and hence d) in \tilde{G}_1 and has an empty neighborhood. In the latter case, Lemma 1 implies $Z_1^* \tilde{\succeq}_1 d_1$. ■

We can now conclude the proof of part 3 of the theorem by combining these lemmas.

Let $X = Z_1 \cup Z_2 \cup Z_3$. Since $Z_1 \subset \mathcal{R} \subset X$, it follows that if $\tilde{\mu}_{\alpha\beta}(Z_1|X) > 0$ then $\lim_{\beta \rightarrow 0} \tilde{\mu}_{\alpha\beta}(Z_1|\mathcal{R}) > 0$ which by Proposition 7 is sufficient for the result. We will show that $Z_1^* \tilde{\succeq}_1 A$ for all $A \in \tilde{\mathcal{A}}_0$ such that $A \cap X \neq \emptyset$. By

Proposition 2 it will follow that $\tilde{\mu}_{\alpha\beta}(Z_1^*|X) > 0$.

In view of Lemma 3, we need only show that $Z_1^* \tilde{\succeq}_1 A$ for every A that meets Z_2 . Consider such an A , and let $d \in A \cap Z_2$.

In proving that $\mu_{\alpha\beta}^*(Z_1) > 0$, we constructed a path $\{d, d_1, \dots, d_n\}$ from d through Z_2 to a state $d_n \in Z_2$ in which there was an empty neighborhood. This path was in G_1 , each of its edges had C-cost 1, hence each edge is its own projection, and thus the path is in \tilde{E} .

Consider any edge (d_t, d_{t+1}) of this path for which $d_t \in A_t$ for some $A_t \in \tilde{\mathcal{A}}_0$. Since the associated transition involved a single mutation, it has C^β -cost 1. Therefore if $d_{t+1} \in A_{t+1}$ for some $A_{t+1} \in \tilde{\mathcal{A}}_0$, then the edge (A_t, A_{t+1}) is part of \tilde{G}_1 .

Suppose this holds for all edges in the path. Then the path to A_n is in \tilde{G}_1 , and A_n contains $d_n \in Z_2$. Since d_n has an empty neighborhood, by Lemma 1 we are done.

Suppose on the other hand that there is some edge whose endpoint is not in any $A \in \tilde{\mathcal{A}}_0$, and let d_t be the earliest such endpoint.

The demographic d_t is in \mathcal{A}_0 . That it is not in any $A \in \tilde{\mathcal{A}}_0$ means that there is some $A' \in \tilde{\mathcal{A}}_0$ such that $A' \tilde{\succeq}_0 d_t$ and $d_t \notin A'$. By Proposition 3, $Z_1^* \tilde{\succeq}_1 A'$ implying $Z_1^* \tilde{\succeq}_1 d_{t-1}$. ■

References

- BLUME, L. (1993): “The Statistical Mechanics of Strategic Interaction,” *Games and Economic Behavior*, 4, 387–424.
- ELLISON, G. (1993): “Learning, Local Interaction, and Coordination,” *Econometrica*, 61, 1047–1071.
- FOSTER, D., AND P. YOUNG (1990): “Stochastic Evolutionary Game Dynamics,” *Journal of Theoretical Population Biology*, 38, 219–230.
- FREIDLIN, M., AND A. WENTZELL (1984): *Random Perturbations of Dynamical Systems*. Springer-Verlag, New York.
- FUDENBERG, D., AND C. HARRIS (1992): “Evolutionary Dynamics with Aggregate Shocks,” *Journal of Economic Theory*, 57, 420–442.
- HARSANYI, J., AND R. SELTEN (1988): *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge.

- KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 65, 383–414.
- KIM, Y.-G., AND J. SOBEL (1995): "An Evolutionary Approach to Pre-Play Communication," *Econometrica*, 63, 1181–1193.
- MAILATH, G., L. SAMUELSON, AND A. SHAKED (2001): "Endogenous Interactions," in *The Evolution of Economic Diversity*, ed. by A. Nicita, and U. Pagano. Routledge, New York.
- MATSUI, A. (1991): "Cheap-Talk and Cooperation in a Society," *Journal of Economic Theory*, 54, 245–258.
- NÖLDEKE, G., AND L. SAMUELSON (1993): "An Evolutionary Analysis of Backward and Forward Induction," *Games and Economic Behavior*, 5, 425–454.
- OECHSSLER, J. (1999): "Competition Among Conventions," *Mathematical and Computational Organization Theory*, 5, 31–44.
- ROBSON, A. (1993): "The Adam and Eve Effect and Fast Evolution of Efficient Equilibria in Coordination Games," mimeo.
- ROMANOVSKY, A. (1970): *Discrete Markov Chains*. Wolters-Noordhoof, Groningen.
- SAMUELSON, L. (1994): "Stochastic Stability in Games with Alternative Best Replies," *Journal of Economic Theory*, 64, 35–65.
- YOUNG, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.