The Evolution of Conventions with Mobile Players *

Tone Dieckmann
Department of Economics
National University of Ireland,
Maynooth

Revised: August 1997

Abstract

This paper analyzes the evolution of conventions in a society with local interaction and mobile players. Three innovative aspects are introduced: Imperfect observability of play outside a player's home location, friction in the strategy adjustment process, and restricted mobility. It is shown that, if mobility is unrestricted, only efficient conventions are stochastichally stable. If there are barriers on mobility, the coexistence of different conventions can be observed. While imperfect observability and friction alone cannot prevent society from reaching an overll efficient outcome, restricted mobility can.

Keywords: Evolution of Conventions; Mobility; Imperfect Observability
JEL classification: C72, D82

1 Introduction

In recent years, there has been considerable interest in coordination problems in games (e. g. Kandori et al. (1993), Vega–Redondo (1995, 1996), Young (1993), among many others). When a game possesses several strict equilibria, the traditional refinements of Nash equilibrium are not applicable. On the other hand, strong selection results have been derived from evolutionary learning models. While most models are based on the hypothesis of *random matching* among the

*The main part of this work was done while I visited CentER, Tilburg. I am grateful to Andreas Blume, Eric van Damme, Martin Dufwenberg, Frederic Palomino, and Alex Possajennikov for helpful discussions and comments. Financial support was provided by a TMR grant from the European Commission. CentER's hospitality is gratefully acknowledged.

members of the population (e. g. Binmore and Samuelson (1994), Kandori et al. (1993), Young (1993)), an increasing number of papers deals with *local interaction structures* (e. g. Anderlini and Ianni (1996), Berninghaus and Schwalbe (1994a,b), Boyer and Orléan (1992), Blume (1993), Ellison (1993), Eshel et al. (1996)). In these models, each player is matched only with a subset of the population, her *neighbourhood.* With respect to equilibrium selection results, however, the matching mechanism seems to be irrelevant: For a large class of adaptation rules and generic modelling of mutations, the risk dominant equilibrium is selected (a concept introduced by Harsanyi and Selten (1988)) as the unique stable outcome, independent of the underlying interaction structure.[1] On the other hand, recent work by Bhaskar and Vega-Redondo (1996), Ely (1995), and Oechssler (1997) has shown that the selection of the risk dominant equilibrium is due to the presumption of an exogenously fixed interaction structure. In their models, the players can *choose* their neighbourhoods by moving between locations. In this setting, the payoff dominant equilibrium can be the unique stable outcome. The explanation for this result is very intuitive: The inefficient equilibrium can be upset by a single player's moving to an unoccupied location and playing the strategy corresponding to the efficient equilibrium. All other players will then follow as soon as they get the chance to move. Conversely, in a situation where all players initially play the efficient equilibrium, no player would follow if any number of players (short of the entire population) moved to another location to play the risk dominant but inefficient equilibrium. As a consequence, the risk dominant equilibrium is much more likely to be destabilized by 'trembles' than the efficient equilibrium.

The purpose of this paper is to examine under what conditions this efficiency result can be generalized to less restrictive assumptions concerning the players' behaviour. I focus on three aspects: Imperfect observability of play, friction in the adjustment process, and limited mobility.

Ely (1995) and Oechssler (1997) presume the ability of each player to observe the individual strategies and payoffs of all other players in the entire population and, what is more, to take all this information into account when choosing his[2] strategy. However, in the context of local interaction, and especially in the case of large populations usually considered in evolutionary models, it is natural to assume that players have *local information* in the sense that they are better informed about their neighbours' doings than about what is going on at other locations. Moreover, I believe that the limited capacity to gather, memorize and make use of information is an essential feature of bounded rationality. Therefore, the present model departs from the framework of Ely and Oechssler by assuming that the players can only imperfectly observe the play of the game outside their own

[1] Exceptions are Vega–Redondo (1995), where players' strategy choices are based on expectations, and Vega–Redondo (1996), where a different matching mechanism is considered.

[2] As this is a model of bounded rationality, I assume the players to be male.

neighbourhood. Of the various conceivable ways to model imperfect observability, I consider two szenarios which seem the most straightforward to me. In the first setting, the players can observe play at distant locations only with a certain (positive) probability, which may vary across locations. The choice of locations is then restricted to the ones where play is actually observed.

In the second setting, play at distant locations cannot be observed at all prior to moving. All players observe only a statistic about play at any other than their present location, namely the average payoff across players at that location. Players cannot observe the strategy distributions prevailing at other locations. Location choice has to be made under imperfect information. Only after moving, i. e. after the location choice has been irrevocably fixed for that period, the player learns the strategy distribution at the new location. However, it is assumed that the player might not be able to adjust his strategy immediately after moving to a new location. He can do this only with a certain (positive) probability.

The second szenario bears some similarity to the assumptions made by Bhaskar and Vega–Redondo (1996). However, there is a fundamental difference with respect to the way friction is introduced. In their model, play at all locations is perfectly observable by all players at all times. However, knowing that they can adjust their strategies only with a certain probability p after moving, the players compute their *expected* payoffs for all locations, taking p into account. Location choice is then carried out on the basis of expected payoffs. As a consequence, if p is small, players can be deterred from moving to a location where the efficient convention is played because they expect to be unable to adjust their strategies at the new location. In the present model, in contrast, location choice is independent of p. The probability affects only the players' actual strategy choice, not their decision to move. It will be seen that the assumption of the players' taking friction into account prior to moving has the same effect as restricted mobility (as introduced in section 7 of this paper) in that it can prevent players from moving to a preferred location.

It will be shown that, in both settings, the selection of the payoff dominant equilibrium can be sustained as long as mobility is unrestricted. This result reinforces the folk theorem that mobility promotes efficiency. However, if mobility is limited, e. g. due to limited capacity of locations to accomodate players, the coexistence of different conventions may be observed. Further, in the case of two locations, it will be shown that all stochastically stable states involve at least one location where the risk dominant convention is played. That is, restricting the players' mobility actually prevents them from achieving overall efficient play.

The paper is organized as follows. The next section provides an informal description of the model. Section 3 presents the basic model, the dynamics of which are discussed in section 4. In section 5, I briefly describe a method to determine stochastically stable sets, which was developed by Glenn Ellison (1995). In section 6, equilibrium selection results are derived. Section 7 analyzes the impact of mobility restrictions. The final section concludes.

2 Informal Description of the Model

Imagine a network of adjoining cities or regions (locations), inhabited by the members of a large but finite population of players. In each period, each inhabitant of a location is pairwise matched with all his neighbours, i. e. all other inhabitants of that location, and (in the case of overlapping neighbourhoods) also with inhabitants of neighbouring locations, to play a coordination game. Familiar examples which fit to this setting include the circular interaction structure (e. g. Ellison (1993)), where interaction takes place between neighbouring locations on a circular line; or the grid structure (e. g. Berninghaus and Schwalbe (1994b)), where each location is represented by a node of a lattice, or in fact any allocation of locations in Euklidean space.

Players can choose their neighbourhoods by moving freely between locations. After locations have been chosen at the beginning of each period, the strategies for the coordination game are determined by an *imitation rule*: Each player adopts the strategy which gained the highest average payoff at his present location in the previous period.

A crucial assumption of the model is the imperfect observability of play outside a player's own neighbourhood. Two different settings are considered. First, I assume that a player presently situated at some loction L be able to observe the average payoff gained by each strategy at another location L' with probability $p(L, L') > 0$. For instance, this probability could be a decreasing function of the geographical distance between locations L and L': A player is more likely to observe play at locations nearer to his home, and less likely to observe play at more distant locations. Of all observed locations, the player then picks the one with the highest average payoff gained by any strategy. If this is not unique, the player randomizes, placing positive probability on each location where the observed payoff is maximal. After moving, the player simply adopts the strategy that gained the highest average payoff at the new location in the previous period.

In the second setting, the players are unable to observe the average payoff gained by each strategy at other locations. They observe only the average payoff across all strategies at any location. That is, for each location (other than his home location), the player can only observe the average payoff to the inhabitants of that location, but not the payoffs gained by the individual strategies. Each player will then choose the location with the highest average payoff, and randozmize in case of ties. Once a player has moved to a new location, he learns the average payoffs gained by all strategies at that location. However, he is able to adjust his strategy according to the imitation rule only with probability $p > 0$. For simplicity, I assume that this probability is constant across players, locations, and periods.

As a first result, it will be shown that, in an absorbing state of the dynamical system generated by the individual adaptation processes, all players live at the same location and play the same strategy. The main result states that, given that all players live at the same location, an efficient convention is less likely to be destabilized by simultaneous mistakes on the part of the players than an inefficient convention, irrespective of the underlying assumptions concerning imperfect observability and friction. In the case of restricted mobility, however, this result no longer holds. Instead, overall efficiency as a stable outcome is precluded.

3 The Model

There is a large but finite population $I = \{1, \ldots, N\}$ of players and a finite set $\mathcal{L}$, $|\mathcal{L}| \geq 2$, of locations which are represented by the nodes of a graph. Two locations $L, M \in \mathcal{L}$ are *connected* if the respective nodes of the graph are connected by an edge. Let $C(L) \subset \mathcal{L}$ denote the union of the set of all locations connected to L with L itself. We place no further restrictions on the interaction structure except for excluding the trivial case of full connectivity, i. e. I assume that there exist at least two locations L, M with $L \notin C(M)$.

Time is discrete. At the beginning of each period t, each player $i \in I$ chooses a location $\lambda_i^t \in \mathcal{L}$. The *neighbourhood* of player i in t is defined by the set of players $C(\lambda_i^t)/\{i\}$, the elements of which are referred to as her *neighbours*. After locations are chosen, each player chooses a strategy $s_i^t \in \{x, y\}$. Each player is then characterized by a strategy–location pair $\sigma_i^t := (s_i^t, \lambda_i^t)$. The *state of the system* at time t is a vector $\sigma^t = (s^t, \lambda^t)$, where the vector $s^t = (s_i^t)_{i=1}^N$ indicates a strategy for each player, and $\lambda^t = (\lambda_i^t)_{i=1}^N$ indicates the locations of the players. Let Σ denote the set of all possible states of the system. Further, let n_L^t denote the number of inhabitants of location L in state σ^t.

In each period t, after locations are determined, each player is matched with each of his neighbours to play a coordination game

	x	y
x	a, a	b, c
y	c, b	d, d

. (1)

Coordination games are characterized by the existence of two strict Nash–equilibria, (x, x) and (y, y), i. e. $a > c$ and $d > b$. The most interesting case is the one where one of the equilibria, say (x, x), is risk dominant, whereas the other one is payoff dominant, i. e. $(a - c) > (d - b)$ and $a < d$. The focus of the present paper is on this case. As an example, consider the game

	x	y
x	$3,3$	$2,0$
y	$0,2$	$4,4$

. (2)

3.1 The Imitation Rule

Whenever a player gets the chance to adjust his strategy for the coordination game, he will do so according to an imitation rule. This rule prescribes to choose the strategy that gained the highest average payoff at the player's location in the previous period. Let $c_L := \Sigma_{M \in C(L)} n_L$ denote the number of players living in $C(L)$ in state σ, and suppose that v_L of these players chose x, $v_L \leq c_L$.[3] Then, the average payoff gained by an x–player at location L in state σ is

$$\pi_x(L, \sigma) := \frac{(v_L - 1)a + (c_L - v_L)b}{c_L - 1}, \tag{3}$$

and the average payoff gained by a y–player is

$$\pi_y(L, \sigma) := \frac{v_L c + (c_L - v_L - 1)d}{c_L - 1}. \tag{4}$$

Loners, i. e. players who have no neighbours to interact with, receive a *reservation payoff* $\pi_r < \min\{a, b, c, d\}$.

Now the imitation rule prescribes to play strategy x if $\pi_x(L, \sigma) > \pi_y(L, \sigma)$, which is equivalent to

$$v_L > \frac{c_L(d - b) + a - d}{a - c + d - b} =: v_L^*, \tag{5}$$

and they choose y if the reverse inequality holds, which is equivalent to

$$c_L - v_L > \frac{c_L(a - c) - a + d}{a - c + d - b}. \tag{6}$$

If both strategies earned the same payoff, the player randomizes, placing positive probability on both strategies. Loners are supposed to pick a strategy at random.

Note that the assumptions on risk dominance and efficiency implay that $v_L^* > c_L - v_L^*$, i. e. $v_L^* > (c_L/2)$, and $c_L - v_L^* < (c_L/2)$, for all locations L. This means that it takes less than half the population at any given location to make x the better strategy, and more than half the population has to play y in order to make y the better strategy.

3.2 Location Choice

I consider two settings, the first of which focuses on imperfect observability alone, while the second one also deals with friction.

[3] For ease of notation, we write v_L for $v_L(\sigma)$, etc.

3.2.1 Setting S1

At the end of each period, each player i at location L observes the average payoffs gained by both strategies at L, $\pi_x(L,\sigma)$ and $\pi_y(L,\sigma)$. Further, he observes the average payoffs at all other locations M, $\pi_x(M,\sigma)$ and $\pi_y(M,\sigma)$, with respective probabilities $p(L,M) > 0$. These probabilities are independent across players and periods. The locations for which player i's realization of $p(L,\cdot)$ is positive are referred to as *the locations observed* by player i, and the set of these locations is denoted by $\mathcal{L}(i,L,\sigma)$. Player i then chooses the location with the highest observed payoff for any strategy, i. e. he chooses a location

$$\max_{M\in\mathcal{L}(i,L,\sigma)} \max\{\pi_x(M,\sigma),\pi_y(M,\sigma)\}. \tag{7}$$

If the maximizer of (7) is not unique, the player randomizes, placing positive probability on each.

Once locations are determined, strategies are chosen according to the imitation rule.

3.2.2 Setting S2

At the end of each period, each player i at location L observes the average payoffs gained by both strategies at L, $\pi_x(L,\sigma)$ and $\pi_y(L,\sigma)$. Further, he observes the average payoff across all players at each location $M \neq L$:

$$\pi(M,\sigma) := \frac{z_M \pi_x(M,\sigma) + (n_M - z_M)\pi_y(M,\sigma)}{n_M} \tag{8}$$

(where z_M is the number of x–players at M) for $n_M \geq 2$, $\pi(M,\sigma) = \pi_r$ for $n_M = 1$, and it is undefined otherwise. This average payoff at M is referred to as the *location payoff* at M. The player then chooses the location with the highest location payoff, i. e.

$$\arg\max_{M\in\mathcal{L}} \pi(M,\sigma).$$

If there is more than one location where the location payoff is maximal, the player randomizes between all these locations, choosing each with positive probability.

After locations have been determined, each player who has not moved ($\lambda_i^t = \lambda_i^{t-1}$) adjusts his strategy according to the imitation rule. Each player who has moved ($\lambda_i^t \neq \lambda_i^{t-1}$) gets the opportunity to adjust his strategy with probability p. If the opportunity arises, he uses the imitation rule. Otherwise, he sticks to his previous strategy s_i^{t-1}.

4 The Dynamics

The behavioural rules described in S1 and S2 in the previous section each define a stochastic process on the state space Σ. As the state space is finite, and each player's decisions concerning strategy and location choice depend only on the state in the previous period, each of these processes constitutes a Markov chain.

We are interested in *absorbing states* of the system, i. e. states that, once entered, cannot be left again. Such states are stationary over time in the sense that no player changes either his strategy or location any more.

Definition 1 *A state* $\sigma^* = (s^*, \lambda^*) \in \Sigma$ *is* absorbing *if* $\text{prob}(\sigma^*|\sigma^*) = 1$.

The following proposition states that, in an absorbing state, all players must live at the same location and play the same strategy.

Lemma 1 *If* $\sigma^* = (s^*, \lambda^*) \in \Sigma$ *is absorbing, then* $\lambda_i^* = \lambda_j^*$ *and* $s_i^* = s_j^*$ *for all* $i, j \in I$.

Proof. Suppose that two locations are inhabited in an absorbing state. Either the maximal average (location) payoffs differ, in which case players will observe this with positive probability and move, or the relevant payoffs at both locations are equal, in which case players randomize between the locations. In either case, there is a positive probability of movement between the locations, a contradiction to the state being absorbing. Hence, absorbing states are characterized by a single inhabited location. It is obvious that, in such a state, all players must earn the same payoff (or else some players would want to change their strategies). This in turn requires that all play the same strategy. □

There are two different types of absorbing states: Those where only x is played, and those where only y is played. We call such states *conventions*.

Definition 2 *A state* $\sigma^* = (s^*, \lambda^*)$ *is a* convention *if* $s_i^* = s_j^*$ *and* $\lambda_i^* = \lambda_j^*$ *for all* $i, j \in I$.

Proposition 1 states a well known result from the theory of Markov chains[4], namely that the process converges to one of the absorbing states with probability one as time goes to infinity.

Proposition 1 *As time tends towards infinity, the process converges to a convention with probability one, irrespective of the initial state.*

However, *which* of the possible conventions will be reached in the long run depends on the initial state of the system. In order to derive equilibrium selection results independently of initial conditions, I employ the concept of *stochastic stability*.

[4] e. g. Kemeny/Snell (1976), Theorem 3.1.1, p.43.

5 Stochastic Stability

Suppose that in each period there is a small probability (independent across players and periods) $\epsilon > 0$ that each player 'trembles' or 'makes a mistake' i. e. chooses another than the intended strategy and/or location, such that all strategy–location pairs have positive probability[5]. It follows that each state of the system can be reached from every other state. That is, we have an ergodic Markov chain which has a unique stationary distribution. We then obtain the *limit distribution* of the process from the unique stationary distribution by letting the probability of mistakes vanish. By the now well known method of perturbed Markov chains introduced by Freidlin and Wentzel (1984), and further developed by Kandori et al. (1993), and Young (1993), it can be shown that the states which have positive probability in the limit distribution, referred to as *stochastically stable states*, form a subset of the set of absorbing states of the model with no noise. Further, the stochastically stable states are those absorbing states with the largest basins of attraction. A *basin of attraction* of a state σ is the set of all states from which the unperturbed process converges to σ with probability one, and its size is inversely related to the minimal total number of mistakes required to reach this basin from all other absorbing states. That is, the concept of the limit distribution allows us to select a subset of the set of absorbing states of the model with no noise. The focus of this paper is on determining stochastically stable conventions. To this end, I use a method developed by Ellison (1995), which will be briefly described for the reader's convenience.

Define the set of all x–conventions by

$$\Sigma_x := \{\sigma = (s, \lambda) | s_i = x, \lambda_i = \lambda_j \forall i, j \in I\},$$

and analogously the set of all y–conventions by

$$\Sigma_y := \{\sigma = (s, \lambda) | s_i = y, \lambda_i = \lambda_j \forall i, j \in I\}.$$

The set of absorbing states of the process with no trembles ($\epsilon = 0$) is $\Sigma_x \cup \Sigma_y$. In order to sort out those conventions that have positive probability in the limit distribution, I employ the concepts of *radius* and *coradius* of a basin of attraction of a set of absorbing states, introduced by Ellison (1995). First, some additional notation has to be introduced.

Let $\Omega \subset \Sigma_x \cup \Sigma_y$ be a set of absorbing states of the model with no noise ($\epsilon = 0$). The *basin of attraction* of Ω, denoted by $D(\Omega)$, is the set of all states from which the Markov process converges to a state in Ω with probability one:

$$D(\Omega) := \{\sigma \in \Sigma | \text{prob}(\exists \tau \text{ such that } \sigma^t \in \Omega \ \forall t > \tau | \sigma^0 = \sigma) = 1\}.$$

Intuitively speaking, the *radius* of the set $D(\Omega)$ is the number of trembles (in the model with noise) necessary to leave this set, starting from a state in Ω. Write

[5]In evolutionary models, such trembles are referred to as *mutations*.

$c(\sigma, \sigma')$ for the number of independent simultaneous trembles necessary for the system to transit from state σ to state σ'. That is, $c(\cdot)$ measures the *transition cost* between these states in terms of these trembles. Further, define a *path* by a finite sequence $(\sigma^1, \sigma^2, \ldots, \sigma^\tau)$ of distinct states. The *cost* of such a path is defined by

$$c(\sigma^1, \sigma^2, \ldots, \sigma^\tau) = \Sigma_{t=1}^{\tau-1} c(\sigma^t, \sigma^{t+1}).$$

The radius of Ω is the 'cheapest' path that leads from any state in Ω to any other state outside the basin of attraction of Ω.

Definition 3 *The* radius *of the basin of attraction of a set* $\Omega \subset \Sigma_x \cup \Sigma_y$ *is defined by*

$$R(\Omega) := \min_{(\sigma^1, \ldots, \sigma^\tau)} c(\sigma^1, \ldots, \sigma^\tau) \ s.\ t.\ \sigma^1 \in \Omega, \sigma^\tau \notin D(\Omega).$$

The path $(\sigma^1, \ldots, \sigma^\tau)$ defining the radius of $D(\Omega)$ thus describes the 'cheapest way out' of that set. In most cases, it will be seen that this path involves but a single transition, i. e. it consists of two states only. In this case, the cheapest way out is realized by a direct transition from a state in Ω to a state outside $D(\Omega)$. Intuitively, the larger the radius, the 'costlier' (in terms of trembles) it is to leave the set. Put differently: The larger the radius, the more improbable is the event that simultaneous mistakes by individual players shift the system away from this set and thus into the basin of attraction of another absorbing set.

Conversely, the *coradius* of the basin of attraction of a set of absorbing states is defined by the number of trembles necessary to reach this set from the most 'unfavourable' absorbing state outside the set, i. e. from the state where the minimum number of trembles required to reach $D(\Omega)$ is maximized.

Definition 4 *The* coradius *of the basin of attraction of a set* $\Omega \subset \Sigma_x \cup \Sigma_y$ *is defined by*

$$CR(\Omega) := \max_{\sigma^1 \notin \Omega} \min_{(\sigma^1, \ldots, \sigma^\tau)} c(\sigma^1, \ldots, \sigma^\tau) \ s.\ t.\ \sigma^\tau \in D(\Omega).$$

The coradius thus measures the minimum number of trembles required to reach $D(\Omega)$ from the most 'unfavourable' state outside that set. The smaller the coradius, the likelier is the event that simultaneous trembles shift the system from any absorbing state to some state in $D(\Omega)$. Ellison's main result[6] states that a set of states is stochastically stable if the radius of its basin of attraction exceeds the coradius. Formally, let $\mu(\epsilon)$ denote the unique stationary distribution of the ergodic Markov chain defined by the transition probabilities $\text{prob}(\sigma^{t+1}|\sigma^t)$ together with the probabilities of trembles given by ϵ, and write $\mu^* := \lim_{\epsilon \to 0} \mu(\epsilon)$ for the

[6] Ellison's result actually refers to a more stringent concept, the *modified* coradius, and the theorem below is just a corollary of Ellison's main theorem. However, the modified coradius is not needed in the present context.

limit distribution of this process. The *limit set*, i. e. the set of *stochastically stable states*, is defined by

$$\Sigma^* := \{\sigma \in \Sigma_x \cup \Sigma_y | \mu^*(\sigma) > 0\}.$$

If we define $\mu^*(\Omega) := \Sigma_{\sigma \in \Omega} \mu^*(\sigma)$, we can write $\mu^*(\Sigma^*) = 1$.

Theorem 1 *For any $\Omega \subset \Sigma_x \cup \Sigma_y$, if $R(\Omega) > CR(\Omega)$, then $\Sigma^* \subset \Omega$, i. e. $\mu^*(\Omega) = 1$.*

Proof. Ellison (1995), Theorem 1 on page 16. □

The theorem states that, if the radius of the basin of attraction of a set Ω exceeds the coradius, all stochastically stable states are contained in Ω. The intuitive explanation is that it is easier to reach Ω from any other absorbing set, than to reach any other absorbing set from Ω. Thus, Ω is less likely to be destabilized by mutations than all other absorbing sets. At the same time, Ω is more easily reached than other absorbing sets, given any initial state. We will now use the theorem to determine the stochastically stable conventions in our model.

6 Equilibrium Selection

Theorem 1 stated above enables us to select those subsets from the set of absorbing states of the model with no noise which are stochastically stable. Further, Lemma 1 ensures that the states within one set of conventions Σ_s, $s \in \{x, y\}$, differ only with respect to the inhabited location. That is, all elements of Σ_s are identical up to a relabeling of locations. We can infer that $R(\Sigma_x) = CR(\Sigma_y)$, and vice versa. The reason is that, since all states within one set of conventions Σ_s are identical up to a relabeling of locations, the 'most unfavourable' state from which the complement basin of attraction can be reached, i. e. the state from which the minimum number of trembles required to reach the complement set is maximal (which is needed for the computation of $CR(\cdot)$), is equal to the 'most favourable' state from which the basin of attraction of that set can be left: If all states are identical, all states are equally 'favourable' or 'unfavourable'. As the respective basins of attraction of the two sets of conventions form a partition of the state space, leaving one basin of attraction, say $D(\Sigma_x)$, automatically means entering the other basin of attraction, $D(\Sigma_y)$. It follows that the number of trembles required to leave the basin of attraction of one set of conventions is equal to the number of trembles required to reach the basin of attraction of the other set of conventions, which is the complement set. Stated formally:

Lemma 2 *In the present model, $R(\Sigma_x) = CR(\Sigma_y)$ and $R(\Sigma_y) = CR(\Sigma_x)$, independent of the assumptions concerning observability.*

Proof. Since all elements of Σ_x are identical up to a relabeling of locations, the following holds. Given any $\sigma_y \in \Sigma_y$,

$$c(\sigma', \sigma_y) = c(\sigma'', \sigma_y) \quad \forall \sigma', \sigma'' \in \Sigma_x.$$

This implies

$$\min_{\sigma_y} c(\sigma', \sigma_y) = \min_{\sigma_y} c(\sigma'', \sigma_y) = R(\Sigma_x) \quad \forall \sigma', \sigma'' \in \Sigma_x. \tag{9}$$

Since the minimum cost $c(\sigma, \sigma_y)$ to get out of $D(\Sigma_x)$ is identical for all $\sigma \in \Sigma_x$, it is also identical to the maximum over $\sigma \in \Sigma_x$ of this cost, i. e.

$$\min_{\sigma_y} c(\sigma, \sigma_y) = \max_{\sigma \in \Sigma_x} \min_{\sigma_y} c(\sigma, \sigma_y) = CR(\Sigma_y). \tag{10}$$

Combining (9) and (10) yields $R(\Sigma_x) = CR(\Sigma_y)$. A similar argument shows that $R(\Sigma_y) = CR(\Sigma_x)$. □

We will now state our main result that only efficient conventions are stochastically stable, provided that the population is not too small. The model S2 contains the extreme case of perfect observability ($p = 1$) considered by Ely (1995) and Oechssler (1997) as a special case. We show that the limit set corresponds to the set of (efficient) y–conventions by proving that the radius of its basin of attraction exceeds its coradius.

Proposition 2 *In the model defined by either S1 or S2, the following holds: if $N > [(a + d - 2c)/(d - b)] + 2$, only efficient conventions are stochastically stable, i. e. $\mu^*(\Sigma_y) = 1$. Otherwise, only risk dominant conventions are stochastically stable, i. e. $\mu(\Sigma_x) = 1$.*

Proof. We show that the radius of $D(\Sigma_y)$ exceeds its coradius if N exceeds the threshold indicated in the proposition. First, we compute $CR(\Sigma_y)$, the minimum number of trembles required to reach $D(\Sigma_y)$ from the most unfavourable state outside $D(\Sigma_y)$. Because of Lemma 2, this number is equal to $R(\Sigma_x)$, the minimum number of trembles required to leave $D(\Sigma_x)$, starting from a state in Σ_x. To compute this number, consider a state σ in Σ_x. Lemma 1 ensures that such states are characterized by a single inhabited location. Call this location L. Suppose two players tremble: they simultaneously move to another location $L' \notin C(L)$ and there switch to strategy y. Call the resulting state σ'.

First, consider S1. Players at L observe the average payoff $\pi_y(L', \sigma') = d > \pi_x(L, \sigma') = a$ with probability $p(L, L') > 0$. There is a positive probability ($p(L, L')^{N-2}$) that *all* players at L observe this, and thus move to L' and play y. Hence, σ' is outside the basin of attraction of Σ_x, because there is a positive probability that a y–convention will be reached.

Now consider S2. The location payoff at L' in state σ' is

$$\pi(L', \sigma') = d > \pi_x(L, \sigma') = a,$$

which implies that all x–players will move from L to L' in the next period. Further, as $p > 0$, there is a positive probability (p^{N-2}) that *all* $N-2$ players will be able to adjust their strategies at L', and thus switch to y. Again, σ' is outside the basin of attraction of Σ_x, because there is a positive probability that a y–convention will be reached. As two trembles suffice to reach σ', $R(\Sigma_x) = CR(\Sigma_y) = 2$.

Now compute $R(\Sigma_y)$, the minimum number of trembles necessary to leave $D(\Sigma_y)$, starting from any y–convention. According to Lemma 2, this number is equal to $CR(\Sigma_x)$. Consider a state $\sigma \in \Sigma_y$. Lemma 1 ensures that such states are characterized by a single inhabited location. Call this location L. Leaving the basin of attraction of σ cannot be achieved by some players' moving to a location outside $C(L)$ and there switching to x since these players would gain a lower payoff than those remaining at L and, as a consequence, no player would follow to the location where x is played. Therefore, the trembles must be such that the trembling players switch to x but stay at the location L, or within $C(L)$. Consequently, the number of trembles must be high enough to ensure that strategy x gains a higher payoff than strategy y in $C(L)$. That is, a necessary condition for a state σ' to be in $D(\Sigma_x)$ is that $v_L \geq v_L^*$. To transit from σ to $D(\Sigma_x)$ thus requires enough players' switching to x such that

$$v_L > v_L^* = \frac{N(d-b)+a-d}{a-c+d-b}.$$

This can be achieved by

$$R(\Sigma_y) = CR(\Sigma_x) = \left\lfloor \frac{N(d-b)+a-d}{a-c+d-b} \right\rfloor + 1$$

trembles. It is easily seen that $R(\Sigma_y)$ exceeds $CR(\Sigma_y) = 2$ for $N > [(a+d-2c)/(d-b)] + 2$, and $R(\Sigma_x)$ exceeds $CR(\Sigma_x)$ otherwise. The result follows from Theorem 1. □

In the case of large populations usually considered in evolutionary models, N will typically exceed the lower bound $[(a+d-2c)/(d-b)]+2$ necessary to establish an efficient equilibrium configuration as a convention. However, the restrictiveness of the lower bound on N depends on the parameters of the model. In the stage game (2), this bound would be 5.5, such that efficient conventions are stable in populations of six or more players. Note that, while the lower bound on N is sufficient for the set of efficient conventions to be stable, a *necessary* condition for this result is $N > 4$. This is because $v_L^* < N/2$, since the equilibrium (x, x) is risk dominant ($a-c < d-b$). Consequently, a necessary condition for $R(\Sigma_y) > 2$ is $N/2 > 2$. This implies that, in populations of less than five players, only risk dominant conventions prevail in the long run.

6.1 No Information

This section deals with the extreme case of $p = 0$ under S2, i. e. players who move to a new location cannot adjust their strategies and stick to their previous strategy after moving. In this case, only risk dominant equilibrium configurations can be stochastically stable. The reason is that $CR(\Sigma_y)$ is now much larger than in the case of $p > 0$. If the players have no chance at all to adjust their strategies after moving to a new location, two trembles do not suffice to reach the basin of attraction of a y–convention from an x–convention. To see this, consider a state $\sigma \in \Sigma_x$ where all players live at some location L, and suppose two players move to another location $L' \notin C(L)$ and switch to y. Call the resulting state σ'. As before, the 2 players at L' gain the maximal payoff d. Consequently, all other $N-2$ players will move to L' in the next period. But, since they are unable to adjust their strategies at the new location, they will continue to play x. Thus, the number of x–players at L' will be $N-2$, which is larger than $v^*_{L'}$ for $N \geq 4$. As a consequence, σ' is in $D(\Sigma_x)$. Thus, two trembles do not suffice to leave the basin of attraction of any x–convention. Instead, the number of trembles must be such that y gains a higher payoff than x in the new state σ'. This will be the case if more than $n_L - v^*_L$ players switch to y. It follows that

$$R(\Sigma_x) = CR(\Sigma_y) = \left\lfloor \frac{N(a-c)+d-a}{a-c+d-b} \right\rfloor + 1,$$

where the first equality follows from Lemma 2. An analogous argument yields

$$R(\Sigma_y) = CR(\Sigma_x) = \left\lfloor \frac{N(d-b)+a-d}{a-c+d-b} \right\rfloor + 1.$$

Note that the assumption concerning p does not make any difference with respect to the computation of $R(\Sigma_y)$ and $CR(\Sigma_x)$. The reason is that p affects the players' ability to adjust their strategies only if they move to a location outside their current neighbourhood. As y–conventions yield a higher payoff than x–conventions, there is always a strong incentive to move to a neighbourhood where only y is played. Conversely, no player would contemplate moving from a location where only y is played to one where x is played. Therefore, an x–convention cannot be reached from a y–convention by some players moving to another location and switching to x, because no player would follow. Hence, the value of p is relevant only with regard to the computation of $CR(\Sigma_y)$ and $R(\Sigma_x)$.

Since $(a-c) > (d-b)$ (the equilibrium (x,x) is risk dominant), simple computations show that $R(\Sigma_x) > N/2$, whereas $CR(\Sigma_x) < N/2$. According to Theorem 1, this proves the following proposition.

Proposition 3 *In the model S2, if $p = 0$, only risk dominant conventions are stochastically stable, i. e. $\mu^*(\Sigma_x) = 1$.*

The impossibility to adjust one's strategy at a new location thus restores the familiar result that only risk dominant equilibria prevail in the long run. It is easy to see that the same result holds in S1 for the trivial case of $p(L, M) = 0$ for all $M \neq L$: If players cannot observe anything outside their own neighbourhood, they will never move, and the model is equivalent to one with no mobility.

7 Restricted Mobility

We have seen that, under the assumption of unrestricted mobility, absorbing states are characterized by all players living at the same location. We will now consider the case that locations have a limited capacity to accomodate players. Suppose the number of inhabitants of location L in any period is bounded above by $\bar{L}$, $N < \sum_{L \in \mathcal{L}} \bar{L}$. In this case, the players' freedom of choice with respect to locations might be restricted: A player who wants to move might be prevented from doing so because the location of his choice is fully occupied. If this happens, we simply assume that the fully occupied locations are dropped from the players' choice set. In this variant of the model, Lemma 1 no longer holds. Instead, absorbing states are now characterized by several occupied locations. As a consequence, the *coexistence* of both conventions can be observed in absorbing states. To see this, consider the following simple example. Suppose there are only two locations, and the capacity of one of them is restricted to half the population. Then, the state where each location is inhabited by half of the population, and the y–convention prevails at the one location while the x–convention prevails at the other, is absorbing. This is because, given the impossibility of moving for the x–players, no player has an incentive to change his strategy. Hence, restricting the players mobility might lead to the coexistence of conventions in the society.

In the case of restricted mobility, we have to deal with absorbing *sets* rather than absorbing states. This is because, as $N \leq \sum_{L \in \mathcal{L}} \bar{L}$, there is always at least one location where x is played that is not fully occupied (locations where the y–convention prevails must be fully occupied because otherwise x–players would move to that location). Now if there are several locations where x is played, the x–players are indifferent between these locations, and there is a positive probability of movement between these locations.

Any absorbing set with coexistence of conventions is characterized by all players at the same location playing the same strategy, and all locations where y is played being fully occupied. The set of all such states is denoted by Σ_{co}.

Unfortunately, results concerning stochastic stability of absorbing sets can be derived only under very restrictive conditions concerning the population size relative to the capacity constraints. However, for the case of two locations, it can easily be shown that states involving *only* efficient play cannot be stochastically stable. We prove the result by showing that all stochastically stable states are contained in $\Sigma_{co} \cup \Sigma_x$. That is, every stochastically stable state contains at least one location where the inefficient convention prevails.

Let $\Sigma_{cox} := \Sigma_{co} \cup \Sigma_x$.

Proposition 4 *In the model with restricted mobility, if* $|\mathcal{L}| = 2$*, then* $\mu^*(\Sigma_{cox}) = 1$.

Proof. Suppose there are 2 locations, L and M, and $\bar{M} \leq \bar{L}$. We first compute the radius of Σ_{cox}. The easiest way to leave $D(\Sigma_{cox})$, i. e. to transit to $D(\Sigma_y)$, is to start from the state in $\sigma_{co} \in \Sigma_{co}$ where the y–convention is played at L, the location with the less restrictive capacity constraint. Then there are $N - \bar{L}$ players at M playing x. In order to reach $D(\Sigma_y)$, at least $n_M - v_M^*$ of these x–players must switch to y. The other x–players will then imitate y in the next period. Thus, the radius is

$$R(\Sigma_{cox}) = \left\lfloor \frac{(N - \bar{L})(a - c) + a - d}{a - c + d - b} \right\rfloor + 1.$$

The coradius is the minimum number of trembles required to reach $D(\Sigma_{cox})$ from a state in Σ_y. We start from the state $\sigma_y \in \Sigma_y$ where location L is fully occupied, such that $n_M = N - \bar{L}$[7]. If at least v_M^* of the players at M switch to x, the other inhabitants of M will follow in the next period. Thus, the coradius is

$$CR(\Sigma_{cox}) = \left\lfloor \frac{(N - \bar{L})(d - b) + a - d}{a - c + d - b} \right\rfloor + 1.$$

The assumptions on risk dominance and efficiency ensure that the radius exceeds the coradius. □

While states exhibiting overall efficiency of play are excluded from the set of stochastically stable states, states where only the inefficient equilibrium is played are not. The reason is that a state where all play x can easily be reached from a state in Σ_{co}. Thus, restricted mobility actually prevents the players from reaching an overall efficient outcome; it might even prevent efficient play altogether.

A thorough analysis of a similar model with restricted capacity can be found in Anwar (1996). His model is based on the best reply rule instead of imitation, and the locations have equal capacity constraints. In that case, equilibrium selection results depend on the parameters of the model, i. e. stage game payoffs and capacity constraints. However, overall efficient states cannot be stochastically stable. This shows that the qualitative result that overall efficiency is excluded is independent of the exact specification of the learning process and the parameters of the model.

I have chosen to model restricted mobility by imposing capacity constraints on locations. Alternatively, one could imagine a variety of other possible restrictions to a player's mobility. For instance, one could assume that moving incurs a cost which increases in the number of players already residing at the destination. Then there is a limit to the number of players beyond which moving is no longer

[7] This state can be reached from any other state in Σ_y with zero trembles.

worthwhile because it is too costly. This limit would be equivalent to a capacity constraint. Another possibility would be to allow the inhabitants of a location to refuse potential immigrants. Also, friction in the adjustment process can be modeled in a way that it deters players from moving, as explained in the conclusion to this paper. As a matter of fact, the exact specification of mobility restrictions plays no role. What exactly prevents the players from moving is just a question of interpretation, and thus irrelevant with respect to the formal model.

8 Conclusion

The models of Ely (1995) and Oechssler (1997) show that mobility promotes efficient play in coordination games. The present paper reinforces that result. Even if players can only imperfectly observe play at distant locations, or if the strategy adjustment process is impeded by friction, the set of efficient conventions remains stable. However, the exact specification of the way friction is introduced into the adaptation rule can be crucial, as a comparison with Bhaskar and Vega–Redondo (1996) shows. As far as the underlying assumptions are concerned, theirs is the model probably most closely related to my own. Bhaskar and Vega–Redondo assume that, in each period, each player gets the opportunity to adjust his strategy for the coordination game only with a certain probability, p. A player's decision whether to move to another location or not is then based on *expected* payoffs, and a player will move to another location only if his expected payoff at that location exceeds his current (secure) payoff. For instance, consider a player currently residing at a location L where the risk dominant convention prevails, his payoff being a. Suppose there is another location, say M, where the efficient convention is played. If the player moved to M, his expected payoff would be pd, the probability of his getting the opportunity to adjust his strategy times the payoff presently gained at M. Now, if the expected payoff falls short of his present payoff a, the player will not move, even though he *might* get the higher payoff d if he did. Thus, the fact that friction is taken into account prior to moving can actually deter players from moving to a location where the efficient convention prevails. In my model, in contrast, the friction is not considered in the location choice, and players will always move to a location where the better convention prevails, unless there are capacity constraints. Hence, the way friction is modeled by Bhaskar and Vega–Redondo serves the same purpose as the capacity constraints in my model: both may prevent a player from moving to a preferred location.

References

Anderlini, L., and A. Ianni (1996): 'Path Dependence and Learning from Neighbors', Games and Economic Behavior 13, 141-177.

Anwar, A. W. (1996): 'When does Immigration Facilitate Efficiency?', University College London working paper.

Berninghaus, S. K., and U. Schwalbe (1994a): 'Evolution, Interaction and Nash Equilibria', Journal of Economic Behavior and Organization 29, 57-86.

Berninghaus, S. K., and U. Schwalbe (1994b): 'Conventions, Local Interaction, and Automata Networks', Journal of Evolutionary Economics 6, 297-312.

Bhaskar, V., and F. Vega-Redondo (1996): 'Migration and the Evolution of Conventions', University of Alicante discussion paper.

Binmore, K., and L. Samuelson (1994): 'Muddling Through: Noisy Equilibrium Selection', SFB 303 Discussion Paper No. B-275, Universität Bonn.

Blume, L. (1993): 'The Statistical Mechanics of Strategic Interaction', Games and Economic Behavior 5, 387-424.

Boyer, R., and A. Orléan (1992): 'How do conventions evolve?', Journal of Evolutionary Economics 2, 165-177.

Ellison, G. (1993): 'Learning, Local Interaction, and Coordination', Econometrica 61, 1047-1071.

Ellison, G. (1995): 'Basins of Attraction, Long Run Equilibria, and the Speed of Step–by–Step Evolution', MIT working paper.

Ely, J. (1995): 'Local Conventions', University of California at Berkeley working paper.

Eshel, I., L. Samuelson, and A. Shaked (1996): 'Altruists, Egoists and Hooligans in a Local Interaction Model', University of Bonn SFB 303 Discussion Paper.

Eshel, I., E. Sansone, and A. Shaked (1996): 'Evolutionary Dynamics of Populations with a Local Interaction Structure', working paper, University of Bonn.

Freidlin, M. I., and A. D. Wentzell (1984): Random Perturbations of Dynamical Systems. New York: Springer Verlag.

Harsanyi, J. C., and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, Cambridge, Massachusetts: MIT Press.

Kandori, M., G. Mailath, and R. Rob (1993): 'Learning, Mutation, and Long Run Equilibria in Games', Econometrica 61, 29-56.

Kemeny, J. G., and J. L. Snell (1976): Finite Markov Chains. New York: Springer Verlag.

Oechssler, J. (1997): 'Decentralization and the Coordination Problem', Journal of Economic Behavior and Organization 32, 119-139.

Vega–Redondo, F. (1995): 'Expectations, Drift, and Volatility in Evolutionary Games', Games and Economic Behavior 11, 391-412.

Vega–Redondo, F. (1996): 'Efficient Equilibrium Selection in Evolutionary Games with Random Matching', Journal of Economic Theory 70, 65-92.

Young, P. H. (1993): 'The Evolution of Conventions', Econometrica 61, 57-84.