Comments welcome

COMPETITION AMONG CONVENTIONS*

Jörg Oechssler**

Department of Economics

Columbia University

December 5, 1993

First draft: January 1993

Abstract:

Journal of Economic Literature Classification number: C 72

A convention can be seen as the standard way of playing a game. If different conventions exist in various geographical, social or other entities (called "towns") and if there is some mobility between these towns, which conventions, if any, will emerge as the successful ones? A simple evolutionary process is suggested and it is shown that the process converges to a Nash equilibrium for all games satisfying weak acyclity or a condition called *evolutionary stable with respect to pure strategies* (ESPS). Further, if the process converges, it converges to an efficient convention for all games in which the Pareto optimal symmetric equilibria are strict. Hence, the paper presents an explanation for the endogenous evolution of efficiency. In contrast to most recent studies in evolutionary game theory, the conclusions do not rely on random "mutations". Instead, the driving force is the tendency of players to have increased interaction with member of their own group (viscosity).

* I am indebted to Omar Azfar, David Canning, Susan Elmes, Tom Krebs, Aki Matsui and seminar audiences in Bonn and Hamburg for very helpful comments and many discussions. Financial support from the German Academic Exchange Service is gratefully acknowledged.

** 362Riverside Dr. #9B8, New York, NY 10025, (212) 749-1578

email: jo33@columbia.edu.

I. INTRODUCTION

In different societies, regions, clubs, classes or other groups, there can be quite different conventions about how to behave and how not to behave. There are conventions concerning which language to speak, what kind of dress to wear to the opera, what side of the road to drive on, whether to collude in a market or engage in cut-throat competition, etc. The acceptance of these conventions may change over time, some may be adopted by the entire population while others are only honored in some small fraction of the population, or a convention may die out completely.

This essay, using an approach from evolutionary game theory, tries to explain why some conventions are more successful than others. In contrast to Young (1993) who addresses the origin of conventions, this paper takes the existence of competing conventions in different social, geographical or economic entities called "towns" as a given, and tries to explain which, if any, emerge as the successful conventions.

According to David Lewis' (1969) definition, a convention is a regularity in behavior in a repeated coordination game to which (1) almost everyone conforms, and (2) almost everyone expects (and prefers) everyone else to conform. Similarly, I will define a *convention* to be a pure strategy that is played and is expected to be played by every member of some group in a 2-person random matching game.

Taking this definition as a starting point, the basic premise of the approach taken in this paper is the assumption that people who share a common convention, i.e. belong to the same "town", have most of their interaction with members of this group rather than with outsiders. This tendency to have increased interaction with neighbors, termed *viscosity* by biologists, is responsible for the central conclusion of the paper; namely, that the evolutionary process allows societies to coordinate endogenously on an efficient convention.[1]

Since this paper is primarily interested in explaining the competition among existing conventions, I will assume that at the beginning of the game each possible convention is being played somewhere in society. Once the game starts players are free to adjust their strategies and move between towns. However, I will assume that there is some *inertia* in the sense that players are locked into their strategy for some time due to adjustment and information costs or the existence of legal and/or contractual barriers. I will assume that each player has the chance to adjust his strategy with some constant probability every period.

---

[1] A good example is the evolution of free-trade areas. In Europe the efficient convention "free trade" emerged as the dominant convention as more and more countries left the "protectionism" camp and joined the EC. One important factor in this development was the fact that the countries in question have most of their trade interactions with each other.

Once a player has the opportunity to adjust, he chooses a best response against the current strategy profile, i.e. he picks a town and a strategy that maximizes his expected payoff in the current environment. I will further assume that a player switches towns or changes strategies only if he can strictly improve on his current position. In contrast to much of the literature in evolutionary game theory I will *not* make the assumption that players make mistakes.

The main conclusions of the paper can be summarized as follows. First, the evolutionary process, despite the absence of 'mutations', converges for a very broad class of games. Specifically, the process converges for all games that satisfy either weak acyclity (Young, 1993) or a condition, called *evolutionary stability with respect to pure strategies* (ESPS). This result is independent of the town structure of the model and applies to all models with the same best response dynamics.

ESPS is the an adaptation of the concept of an evolutionary stable strategy (Maynard Smith, 1982) to games in which players are restricted to playing pure strategies. Games that satisfy either weak acyclity or ESPS include, among others, coordination games, common interest games, 2x2 games, and games with only pure Nash equilibria.

The second main result is, that, if the process converges, it converges to an efficient convention (rather than to the risk dominant one - if they differ) whenever the Pareto optimal among the symmetric Nash equilibria are strict. This result, which is due to the assumed town structure, shows how important neighborhood effects are for the emergence of efficient coordination.

Some of the ideas in this paper can be traced back to Axelrod's *Evolution of Cooperation* (1984). More closely related, however, are several recent contributions to learning and equilibrium selection. In the first, which can be termed the "ask around" or "adaptive play" approach (Young, 1993) players ask around to find out how the game was played in the past, i.e. they take a random sample of past play.[2] Players then choose an optimal strategy against this sample.

The second approach is based on the "inertia" assumption mentioned above, first introduced by Kandori, Mailath and Rob (1993).[3] In their model, the opportunity to change a strategy comes up with some constant probability every period. If players get the opportunity, they myopically choose a best strategy against the current strategy profile which they are assumed to know.

It has been shown that with both, the "ask around" and the "inertia" approach, play converges for some classes of games to certain long-run states which are in general Nash equilibria

---

[2] See Canning (1992) for an application of this approach to signaling games.

[3] See also Kandori and Rob (1992). Nöldeke and Samuelson (1992) apply this approach to extensive form games.

of the game. Through the introduction of mutations or mistakes (players play a suboptimal strategy with some probability) a unique equilibrium can be selected for some games in the sense that this equilibrium will be observed "almost all the time" in the long-run. This follows from the fact that it takes a certain minimum number of mutations to move from one long run state to another. If mutations are rare, and if it takes fewer mutations to move from A to B than from B to A, then A will be observed more often. As the mutation rate goes to zero, A will be observed almost all the time. With respect to 2x2 games with two strict Nash equilibria both approaches select the risk dominant equilibrium (Harsanyi and Selten, 1988).

It may be questioned whether risk dominance is a good equilibrium selection criterion in the context of repeated interaction. In games where the Pareto efficient and the risk dominant equilibrium do not coincide, it is reasonable to select the risk dominant equilibrium in a *one shot situation* because no one can be sure about the other player's strategy. In the context of repeated interactions in some population, it seems less risky to play the Pareto efficient equilibrium once a conventions is established since, by definition, if a convention is in place, everyone expects everyone else to behave according to it.[4]

The third approach in the literature, which, contrary to the above, yields Pareto efficient outcomes, is based on cheap-talk (Matsui 1991, Kim and Sobel 1992, and Sobel 1993). Again, players are randomly matched to play a normal form game. Before they play, however, each player can send a costless and non-binding message regarding the strategy he proposes to play. Cheap-talk works as a "secret handshake" (Robson, 1990) that allows players to escape from inefficient equilibria in common interest games and symmetric coordination games. The problem with this approach is that, because cheating is costless, there is no pressure for any pathological liars that may exist in the initial population to change their strategies. For them to disappear it may take a long time and the coincidence of many mutations.

My paper diverges from these three approaches in the way I model how conventions can be overturned. In the models discussed above, conventions (or long-run states) are upset if a sufficient number of mutations in a certain direction occur, which may take a very long time to happen, especially in large populations.[5] During the time leading up to a change, any suboptimal payoff resulting from the use of an inefficient convention has no effect on the success of the convention in question. I take the opposite approach in this paper. I disregard mutations and concentrate instead on the payoffs associated with certain conventions. Conventions change

---

[4] The efficient equilibrium would also be selected by Harsanyi and Selten (1988, p. 356) as their theory gives clear precedence to *payoff* dominance over *risk* dominance.

[5] See Kandori et al. (1993) for an example of how the required number of mutations increases with population size. Ellison (1993), however, shows that in a model with local interaction (with players spatially distributed around a circle) convergence may occur at a much faster rate.

because players will abandon a convention if it is not as successful as some alternative in another town.[6] This produces the result that, for a broad class of games, the evolutionary process converges globally to a Pareto efficient convention.

The remainder of this paper is organized as follows. In the next section I present the basic framework for the analysis. Section 3 presents the efficiency result and in section 4 the conditions for global convergence are given. In section 5 the model is generalized to the case with interaction between towns. Section 6 contains a conclusion.

II. THE EVOLUTIONARY PROCESS

Having defined conventions as regularities in strategic behavior to which almost every member of a "town" adheres, I will now suggest an evolutionary process that governs the development of different conventions in the total population. To be precise, I will consider a large population from which pairs of two players are randomly matched to play a normal form game. As defined above, a town is an entity whose members have most of their interactions amongst themselves. In the simple model considered below, players are only matched against players from their own town. In section 5, below, I will relax this assumption and allow for interaction between towns.

In symmetric games the definition of a convention in a random matching process poses no problems: it is simply a pure strategy which (almost) everybody in that town plays; and everybody expects that almost everybody else plays this strategy. However, conventions are by no means restricted to symmetric situations. For example, a convention may be that a driver on the main road has the right of way and a driver on the side road has to wait. Thus, each driver has conditional strategies for each of his possible roles. This suggests how asymmetric situations can be handled in the current framework. Following a common practice in evolutionary game theory (see e.g. Selten, 1980), an asymmetric game is symmetrized by assuming that all players have strategies conditional on their role. Assuming that it is equally likely to be in either role, the payoffs in the symmetrized game (c_{ij}) are simply the average of what a role-one player and a role-two player would get in the asymmetric game (denoted by a_{ij} and b_{ij}, respectively).

Since the primary focus of this paper is to explain the competition among existing conventions, I will assume that at the beginning of the game each possible convention is already established somewhere in society, i.e. in every town there is an historical way of playing the game.

---

[6] In this sense the model is not unlike the "voting with one's feet" in Tiebout's (1956) model of local public goods and its most recent generalization to the choice between institutions by Caplin and Nalebuff (1992).

Assumption 1: The set of pure strategies in the symmetric or symmetrized game is S = {S_1, S_2, ... S_n}. At the beginning of the game, each strategy is being played as a convention somewhere in the population.

Towns are labeled by the convention that existed there at the beginning of the game, that is, a town in which the original population played strategy A is called "town A" throughout the entire game regardless of the current strategy profile of that town (hence, there are n non-empty towns T = {T_1, T_2, ... T_n} at the beginning of the game). Players do not have to play the current convention of their town if it is not in their interest to do so. Once the game has started they may switch to another strategy or to another town if they can expect a strictly higher payoff. However, following Kandori et al. (1993), I will assume that only a fraction of players may switch their strategies or town in every given period. This inertia may be justified by a variety of reasons: there could exist legal and/or contractual barriers against changing one's strategy too often; social pressure might discourage too frequent changes; and finally, there may exist significant adjustment and information costs of finding a better strategy and changing to it. These costs might be lower at some random points in time. For example, if the new convention in some group becomes to drive a pink car, I am much more likely to join in when my old car breaks down.

Assumption 2: Every period the opportunity to change one's strategy/town comes up with constant probability $p_f \in (0,1)$ for each player f.

Whenever a player has the chance to adjust his strategy, he can choose a new town and a new strategy to play in that town. He does so by maximizing his payoff in the *current* environment without taking into account possible future developments of the game. This myopia is justified to some extent by the inertia assumption. Since only a small fraction of the population changes strategies in each period, a strategy that proves successful in the current environment will remain successful for some time in the future.[7] Note that I assume that players observe only the current strategy profile. Players may be ignorant about their opponents' utility functions and no common knowledge about any of the parameters of the game like payoffs, priors, etc. is assumed.

To be precise, let N denote the number of players in the total population, N_j the number of players in town j, and N_{ij} the number of players choosing strategy i in town j. Let the proportion of players in town j playing strategy i be denoted by $\alpha_{ij} \equiv N_{ij}/N_j$. Let $\theta_j \equiv N_j/N$ be the proportion of players located in town j. The strategy profile in each town j is given by $\boldsymbol{\alpha_j} = (\alpha_{1j}, \alpha_{2j}, \ldots, \alpha_{nj})$.

[7] See also Kandori et al. (1993) for this argument.

For any given period t the *state* of the system can then be summarized by a vector $\boldsymbol{\alpha}^{(t)} \equiv (\{\alpha_{ij}^{(t)}\}_{i=1...n,\, j=1...n}, \{\theta_j\}_{j=1...n})$. The *state space* A is given by

$$A = \left\{ \boldsymbol{\alpha}^{(t)} \in [0,1]^{n^2+n} \,\middle|\, \sum_i \alpha_{ij} = 0 \text{ if } \theta_j = 0, \sum_i \alpha_{ij} = 1 \text{ otherwise} \right\}.$$

To avoid arbitrary payoff discontinuities driven by the possibility of a town becoming extinct and the impossibility of resettlement, I assume that in each town there exists a dummy player, who cannot change his strategy. Thus a town cannot die out completely.

A rational player would exclude himself from the profile $\boldsymbol{\alpha}$ and would then play a best reply to the remaining profile. Mainly for notational convenience I will assume that players are not sophisticated enough to do this calculation.[8] Thus, the expected payoff from playing strategy i in town j is

$$\Pi_{ij}(\boldsymbol{\alpha}) = \sum_{k=1}^{n} \alpha_{kj} c_{ik}$$

where c_{ik} is the payoff from playing strategy i against strategy k in the symmetrized game.[9] A *best response* $\beta(\boldsymbol{\alpha}^{(t)}) \in S \times T$ against the current strategy profile is a pure strategy S_i and a town T_j in which to play S_i such that $\beta(\boldsymbol{\alpha}) \in B(\boldsymbol{\alpha}) \equiv \mathit{argmax}_{ij}\ \Pi_{ij}(\boldsymbol{\alpha})$. A best response against the strategy profile in town j is $\beta(\boldsymbol{\alpha}_j) \in B(\boldsymbol{\alpha}_j) \equiv \mathit{argmax}_i\ \Pi_{ij}(\boldsymbol{\alpha}_j)$.

Assumption 3: Players who have the opportunity to adjust their strategy know the current strategy profile $\boldsymbol{\alpha}^{(t)}$ and switch to a best response $\beta(\boldsymbol{\alpha}^{(t)})$ in the next period if and only if their current strategy is not a best response. If there are several best responses, any of them can be chosen (according to some stationary probability distribution) as long as none is excluded a priori.

The above defines a stationary Markov chain on the state space A. I will say that the process has converged to an *absorbing state* in period τ if $\boldsymbol{\alpha}^{(t)} = \boldsymbol{\alpha}^{(t+1)}$ for all periods $t \geq \tau$. The process will be said to *converge globally* if the same absorbing state is reached from all initial partitions of the population into towns satisfying Assumption 1.

III. THE EVOLUTION OF EFFICIENT CONVENTIONS

---

[8] The main result in this paper (concerning the efficieny of conventions) is independent of this assumption. See also Kandori and Rob (1992) for same simplification. There are, however, important differences for the case of convergence to mixed equilibria.

[9] Payoffs of dummy players need not be defined as they cannot change their strategies.

Before proceeding it might be helpful to look at an example of how the process works. I will use an asymmetric coordination game, the "battle of sexes". With payoffs of (4,2) and (2,4) on the main diagonal and zeros elsewhere, the symmetrized version of the game is the following.

	SS	SB	BS	BB
SS	3,3	1,2	2,1	0,0
SB	2,1	0,0	4,2	2,1
BS	1,2	2,4	0,0	1,2
BB	0,0	1,2	2,1	3,3

Figure 1

where BS, for example, means to go to the ballet if you are a type 1 player and to a sporting event if you are a type 2 player. At the beginning of the game, everybody plays the strategy of his town. That is, the initial distribution is $\alpha_{ii} = 1$ and $\alpha_{ij} = 0$ if $i \neq j$. The θ_j are arbitrary as long as $\theta_j > 0$, $\forall i$ and $\Sigma\theta_j = 1$. The first players who have the opportunity to change strategies move to BS-town and play SB, which yields a payoff of 4. However, as soon as 25% of the population in BS-town is playing SB, the expected payoff drops below what a player can get in SS or BB-town (because the SB players now meet more often someone playing also SB in which case they receive a payoff of zero) Hence, every player who has the opportunity to switch will move to SS- or BB- town. The process converges to an absorbing state α with $\alpha_{11} = \alpha_{44} = 1$ and $\theta_1 + \theta_2 = 1$. Since (SS,SS) and (BB,BB) yield the same payoffs, both conventions will coexist in their respective towns.

While the example might suggest that the process always converges to a pure Nash equilibrium, this is not so. Convergence to a mixed strategy equilibrium is possible even if a strict, pure strategy Nash equilibrium exists. Mixed strategy equilibria, however, pose some interpretational problems since individuals are constrained to choosing pure strategies. A mixed strategy equilibrium, therefore, implies that different strategies are played by a homogeneous population in the same town. This seems to contradict the term "convention" according to which everybody is supposed to take the same action (remember that the game is already symmetric). Thus, although the process reaches an absorbing state when it converges to a mixed strategy equilibrium, I will not call this a convention.

Of course, the process can only converge to a mixed strategy equilibrium if N is divisible by the appropriate factor to produce the equilibrium profile α^*; otherwise the process is going to oscillate around the equilibrium forever. For example, in the matching pennies game an absorbing state can never be reached if the number of players is odd. In the following, I will assume that the population size is such that every desired mixed strategy can be achieved by the appropriate population proportions.

Proposition 1: A state α^* *is absorbing only if it induces a symmetric Nash equilibrium in each town, i.e.* α^* *is absorbing only if* (α_j^*, α_j^*) *is a Nash equilibrium for all j. Moreover, if* α^* *induces different Nash equilibria in different towns then they all yield the same expected payoff.*

Proof: (a) Assume to the contrary that the process reaches an absorbing state α which does not induce a Nash equilibrium in at least one town, say town j. Thus, there is at least one player playing a strategy $S_k \notin B(\alpha_j)$. As soon as this player has the opportunity to change strategies he will switch to a best response $\beta(\alpha)$. Since S_k was not a best response, no player switches to S_k and α_{kj} must fall. Hence α was not an absorbing state.

(b) Symmetry of the equilibrium follows from the fact that in any absorbing state α^*,

$$\alpha_{ij}^* > 0 \Rightarrow S_i \in B(\alpha_j^*), \quad \forall j.$$

This implies that any convex combination of the S_i must be a best response against α_j^*. In particular α_j^* must be a best response which shows that (α_j^*, α_j^*) is a symmetric Nash equilibrium.

(c) Assume to the contrary that α^* induces a Nash equilibrium with different expected payoffs in different towns. Obviously, players in towns with lower payoff do not play a best response against α^*. They will switch to a town in which an equilibrium with higher payoff is played as soon as they have the chance. Hence, α^* cannot be an absorbing state.∎

The result of Proposition 1 is not surprising since I have assumed that players possess a degree of rationality that allows them to determine a best response to the current strategy profiles in the different towns, which rules out all non-Nash equilibria absorbing states.[10] The interesting question is to *which* Nash equilibrium the process converges if there are several. In the "battle of sexes" example the process converged to an efficient convention in each town. The next proposition addresses the generality of this result.

Definition 1: A convention is called *efficient* if it is Pareto optimal among all symmetric Nash equilibria.

---

[10] This is an important difference to the biological model in which players (animals) inherit strategies and are genetically bound to use them. In models with local or kinship interaction non-Nash equilibria outcomes become possible because players cannot simply switch to better strategies. Altruistic behavior up to ultimate self sacrifice can thus become evolutionary successful (see e.g. Hamilton, 1964).

Proposition 2: Consider the class of games for which the Pareto optimal symmetric Nash equilibria are strict. If the evolutionary process converges, then an efficient convention is established in all non-empty towns.

Proof: Assume to the contrary that the process converges to an inefficient Nash equilibrium (μ,μ) (which must be symmetric by Proposition 1) in some town, even though an efficient, strict equilibrium, say (S_k,S_k), exists. Since the latter equilibrium is strict, the equilibrium strategy is the unique best response against the original population in k-town. Hence, no player will enter k-town and play a strategy other than S_k (this holds even if $\alpha_{kk} = 0$ at some point during the game). But this contradicts the supposition that (μ,μ) is an absorbing state: at the end players would switch to playing S_k in k-town which yields a strictly higher payoff∎

IV. CONDITIONS FOR CONVERGENCE

We can now turn to the conditions under which the evolutionary process converges to an absorbing state. The results in this section do not depend on the town structure assumed in the previous section. Nor do they depend on the initial conditions stipulated in Assumption 1. The results apply, therefore, to all best reply dynamics that satisfy Assumptions 2 and 3. Sufficient conditions needed to assure convergence are either weak acyclity (Young, 1993) or a condition similar in spirit to the concept of an evolutionary stable strategy (ESS) (Maynard Smith, 1982).

Adapting Young's (1993) definition of acyclity to symmetric games, let $\Gamma(G)$ be the best reply graph of a game G whose vertices consist of pure strategies $S_i \in S$. For any two vertices S_i, S_j, there exists a directed edge $S_i \rightarrow S_j$ if and only if $S_j \in B(S_i)$.

Definition 2: A symmetric game G is *weakly acyclic* if, from any initial vertex S_i, there exists a directed path to some vertex S_k that is part of a strict, symmetric Nash equilibrium (S_k, S_k).

Given a weakly acyclic game convergence of the best reply process can easily be shown. However, even if a game has a very cyclical structure like the game in figure 2, convergence is possible if the game possesses an equilibrium that satisfies a stability condition similar to ESS.

An ESS tests the stability of a (possibly mixed) strategy that is initially played by the entire population against the invasion of a small number of mutants playing a different strategy. A strategy satisfies ESS if it does better on average, when matched randomly against a member of the new population, than does the entrant strategy. The entrants are selected against and the population returns to its original state.

The ESS concept is based on a monomorphic population where players are able to play mixed strategies. The following definition is an adaptation of the ESS concept to polymorphic populations where players are restricted to pure strategies.[11]

If a game is not weakly acyclic, the best reply process may enter an ergodic set of states from which there is no escape. That is, for an ergodic set E, the probability of eventually reaching state α' from α is

$$f_{\alpha\alpha'} = \begin{cases} 0 & \text{if } \alpha \in E, \alpha' \notin E \\ 1 & \text{if } \alpha, \alpha' \in E \end{cases}$$

Let K(E) denote the "carrier" of E, that is

$$K(E) = \left\{ S_i \in S \middle| \alpha_{ij} > 0 \text{ for some } \boldsymbol{\alpha} \in E \right\}.$$

Consider now the reduced game consisting only of strategies contained in K(E). Since it is a symmetric game, it possesses at least one symmetric equilibrium μ^*. Note that μ^* is also an equilibrium of the entire game because E is closed under best replies, that is,

$$S_i \notin K(E) \Rightarrow S_i \notin B(\boldsymbol{\alpha}), \forall \boldsymbol{\alpha} \in E.$$

Let

$$O_{\mu^*}(\mu) \equiv \left\{ S_k \in K(E) \mid \mu_k - \mu_k^* \geq \mu_i - \mu_i^*, \forall i \right\}$$

be the set of strategies in K(E) that are the most over-represented in the strategy profile μ relative to the equilibrium profile μ^*.

Definition 3: A game is *evolutionary stable with respect to pure strategies* (ESPS) if for all ergodic sets E and the corresponding reduced games consisting only of strategies contained in K(E), there exists a strategy profile μ^* s.t.

1. (μ^*, μ^*) is a symmetric Nash equilibrium and
2. $\forall\ \mu \in \Delta K(E), \mu \neq \mu^*, \ \exists\ S_k \in O_{\mu^*}(\mu)$ and $S_j \notin O_{\mu^*}(\mu)$ s.t. $S_k \notin B(\mu)$ and $S_j \in B(\mu)$.

A simpler, but overly restrictive, stability condition would be to require that none of the over-represented strategies can be a best reply. This way it would be easy to guarantee that the process converges to an equilibrium. ESPS requires only that *at least one* of the strategies that are over-represented *the most* in μ relative to the equilibrium profile is not a best reply to μ; and that at least one of the best replies does not belong to the strategies that are over-represented the most. In particular, ESPS implies that if just one player deviates from the equilibrium profile, his new strategy would not be a best response against the new strategy profile. Note also that ESPS implies that μ^* is the unique equilibrium on K(E). As an example consider the following game.

[11] This is not a refinement of ESS as there are equilibria that satisfy ESPS but not ESS, and vice versa.

0,0	x,1	1,x
1,x	0,0	x,1
x,1	1,x	0,0

Figure 2

For all $x \geq 0$, the mixed equilibrium $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ satisfies ESPS. However, if $x < 0$, it is easy to find a μ for which there does not exist a $S_j \notin O_{\mu^*}(\mu)$ s.t. $S_j \in B(\mu)$, e.g. $\mu = (0, \frac{1}{2}, \frac{1}{2})$.

Theorem: The evolutionary process converges to an absorbing state for all finite, symmetric (or symmetrized) games that are either weakly acyclic or satisfy ESPS.

The process as defined by Assumptions 2 and 3 is a stationary Markov chain on a finite state space A. A can therefore be partitioned into transient and (positive) recurrent states. With time, either the process reaches an absorbing state (a symmetric Nash equilibrium by Proposition 1) or it enters a recurrent class (ergodic set) of states. Thus, the process may fail to converge only if there exists a non-singleton ergodic set $E \subseteq A$.

The following lemma will be useful for the proof of the Theorem.

Lemma: If E is a non-singleton ergodic set, then there exists an $\alpha \in E$ *such that* $\alpha_{kl} = 1$ *and* $\alpha_{ij} = 0, \forall\ ij \neq kl$ (excluding the dummy players).

Proof: In the appendix.

Proof of Theorem: The strategy for proving the Theorem is to show that all ergodic sets are singletons, that is, they are absorbing states themselves. The proof is divided into two parts. In the first step I will consider weakly acyclic games and in a second step games that are not weakly acyclic but satisfy ESPS. For both cases I will show that the supposition of an ergodic set being a non-singleton yields a contradiction.

(i) Assume to the contrary that in a weakly acyclic game there exists an ergodic set E that is not a singleton. From the Lemma we know that there exists an $\alpha^{(0)}$ such that $\alpha^{(0)}_{ij} = 1$ for some ij. Since the game is weakly acyclic, there is a directed path in $\Gamma(G)$ that leads from the starting vertex S_{ij} in m steps to a strict, symmetric Nash equilibrium (S_{kj}, S_{kj}), which is an absorbing state. There exists a positive probability that all players have the opportunity to change their strategies simultaneously and switch to the next vertex on the directed path. Hence, there exists a probability $q > 0$ that all players have the opportunity to change their strategies m times in a row such that the process reaches the strict Nash equilibrium.

Since E is a non-singleton ergodic set, $\alpha^{(0)}$ is visited infinitely often. Therefore, the probability of *not* reaching the absorbing state after t visits to $\alpha^{(0)}$, $(1-q)^t$, goes to zero as $t \to \infty$. Hence E contains an absorbing state contradicting the assumption that E is an ergodic set.

(ii) In the second step I will now consider the class of games that are not weakly acyclic but satisfy ESPS. Again, assume to the contrary that there exists an ergodic set E that is not a singleton. Consider the reduced game consisting of the strategies in K(E). By assumption there exists a symmetric equilibrium μ^* satisfying ESPS. I will show by construction that the process converges to an equilibrium μ^*, which is an absorbing state, contradicting the assumption of E being an ergodic set.

As a starting point take some $\alpha \in E$ such that $\alpha_{kl} = 1$, which exists by the Lemma. I will denote the strategies in K(E) by T_i, starting with $T_1 = S_k$. The starting profile in town *l* is thus $\mu^0 = (1,0,0,...,0)$.

If (T_1, T_1) is not a Nash equilibrium, there either exists some strategy in town *l*, call it T_2, that is a best response to μ^0, or there exists a best response T'_2 in some other town. Note that in the latter case $T'_2 \in$ *argmax*$_i$ c_{ii}. Also, if T'_2 is a best response in some town $j \neq l$, then it is a best response in all towns $j \neq l$. Hence, there is a positive probability that N/n players move to each town and play T'_2. Hence, the starting profile in each town looks like μ^0 and the following construction can be continued parallel for each town.

With positive probability q_1 exactly $(1-\mu_1^*)$ percent of players have the chance to switch strategies and adopt T_2 while the remaining μ_1^* percent have to retain T_1. The resulting profile is $\mu^1 = (\mu_1^*, 1-\mu_1^*, 0, 0,...,0)$.

If $\mu^1 = \mu^*$, an absorbing state is reached. If not, $T_2 \notin B(\mu^1)$ because μ^* is assumed to satisfy ESPS: since $\mu^1_2 = 1-\mu_1^* > \mu_2^*$, $\{T_2\} = O(\mu^1)$, which implies that $T_2 \notin B(\mu^1)$.

The remaining possibilities are either,

[1] $\{T_1\} = B(\mu^1)$, i.e. T_1 is the unique best response. In this case players will switch from T_2 to T_1 and increase μ_1 until a profile $\overline{\mu}^1$ is reached, such that either

[1.1] $T_1, T_2 \in B(\overline{\mu}^1) \to$ absorbing state, or

[1.2] $T_1, T_2 \notin B(\overline{\mu}^1)$. Both, T_1 and T_2 are over-represented in $\overline{\mu}^1$ (If $\overline{\mu}^1_2$ became smaller than μ_2^* at some point before reaching $\overline{\mu}^1$, T_1 would be the only over-represented strategy and, hence, by ESPS could not be a best response) $\to$ proceed with the construction; or

[2] $\{T_1\} \neq B(\mu^1) \to$ proceed with the construction.

In the case of [1.2] or [2] there must be some other strategy, call it T_3, that is a best response against $\overline{\mu}^1$ or μ^1, respectively. There exists a strictly positive probability q_2 that $1-\mu_1^*-\mu_2^*$ percent of the players switch to T_3 while the others stay at T_1 and T_2 respectively, yielding the new profile $\mu^2 = (\mu_1^*, \mu_2^*, 1-\mu_1^*-\mu_2^*, 0, 0,...,0)$.

If $\mu^2 = \mu^*$, an absorbing state is reached. If not, $T_3 \notin B(\mu^2)$ due to ESPS. By the same logic as above, either

[1] $\{T_1, T_2\} \not\supseteq B(\mu^2) \rightarrow$ proceed.

[2] $\{T_1, T_2\} \supseteq B(\mu^2)$, i.e., either T_1 or T_2 or both are the only best responses.

[2.1] Consider first the case $\{T_1, T_2\} = B(\mu^2)$. With positive probability the process moves to a state $\mu' = (\mu_1', \mu_2', \mu_3', 0, 0, ..., 0)$, where

$$\mu_1' - \mu_1^* = \mu_2' - \mu_2^* = \mu_3' - \mu_3^*. \qquad (1)$$

Since $O(\mu') = \{T_1, T_2, T_3\}$, by ESPS there exists a $T_4 \in B(\mu')$ and at least one $T_i \in O(\mu')$, s.t. $T_i \notin B(\mu')$. Three cases can occur:

[2.11] $T_i \notin B(\mu')$, $\forall i = 1,2,3 \rightarrow$ proceed.

[2.12] Only one strategy remains a best reply. Assume, without loss of generality, that $T_1, T_2 \notin B(\mu')$. Due to the upper hemi-continuity of the best reply correspondence,

$$\exists\, \mu'' \equiv (\mu_1' - \varepsilon, \mu_2' - \varepsilon, \mu_3', 2\varepsilon, 0,.., 0), \text{ s.t. } T_1, T_2 \notin B(\mu'').$$

Since $O(\mu'') = \{T_3\}$, ESPS implies that $T_3 \notin B(\mu'')$. Furthermore, there exists a $T_4' \in B(\mu'')$. If $T_4 \in B(\mu'')$, let $T_4' = T_4$. Otherwise, everyone currently playing T_4 switches to $T_4' \rightarrow$ proceed.

[2.13] Two strategies remain a best reply. This case can be handled by applying [2.12] twice.

[2.2] The remaining cases are $\{T_1\} = B(\mu^2)$ or $\{T_2\} = B(\mu^2)$. Assume, WLOG, $\{T_1\} = B(\mu^2)$. With positive probability the process moves to $\mu''' = (\mu_1''', \mu_2^*, \mu_3''', 0, 0, ..., 0)$, where the μ_i''' are defined as in equation (1). By ESPS either T_1 or T_3 is not a best reply to μ'''. Assume, WLOG, $T_1 \notin B(\mu''')$.

[2.21] $T_2 \in B(\mu''')$. If $T_2 \in B(\mu''')$, $\hat{\mu} = (\mu_1''' - \varepsilon', \mu_2^*, \mu_3''', 0, ..,0)$ can be reached where $\hat{\mu}$ is such that $T_1 \notin B(\hat{\mu})$ and $T_2 \in B(\hat{\mu})$. Since $O(\hat{\mu}) = \{T_3\}$, by ESPS, $T_3 \notin B(\hat{\mu})$. Hence, μ' can be reached and one can proceed as in [2.1].

[2.22] $T_2 \notin B(\mu''')$. In this case ESPS implies that there exists a $T_4 \in B(\mu''')$. [2.11] and [2.12] can then be applied analogously.

For all the above cases I have shown that a state can be reached in which none of the strategies in $\{T_1, T_2, T_3\}$ is a best reply. The construction can be continued from this state, since there is a positive probability q_3 that $1-\mu_1^*-\mu_2^*-\mu_3^*$ percent of players switch to T_4 while all others stay put, yielding $\mu^3 = (\mu_1^*, \mu_2^*, \mu_3^*, 1- \Sigma\, \mu_i^*, 0, 0,...,0)$.

Continuing in this fashion it is clear that an absorbing state must be reached starting from μ^0 in a finite number of steps with some probability $q > 0$. Since $\mu^0 \in E$, it will be visited infinitely often if E is an ergodic set as assumed. Hence, the probability of not reaching an absorbing state from μ^0 after u visits to μ^0, $(1-q)^u$, approaches zero as $u \rightarrow \infty$. Therefore, E must be a singleton.∎

In much of the literature the discussion of conventions centers on two kinds of games, coordination games and games of common interest. For example, Lewis (1969) based his original definition of conventions on coordination games. A *coordination game* is an n×n game in which each player prefers his j^{th} strategy if and only if the other player also prefers to play his j^{th} strategy.[12] A game of *common interest* can be defined as a game in which there exists a payoff vector that strongly Pareto dominates all other feasible payoffs (see e.g. Aumann and Sorin, 1989).

Proposition 3: For all coordination games and common interest games, the evolutionary process converges *globally* (for all initial conditions satisfying Assumption 1) to an efficient convention in each non-empty town.

Proof: Convergence for symmetric common interest games is obvious. For asymmetric games rename strategies such that (a_{11}, b_{11}) strictly Pareto dominates all other payoff vectors. In the symmetrized game $c_{11} = (a_{11} + b_{11})/2$ Pareto dominates all other payoffs c_{ij}. Hence, from the beginning players switch only to that strategy. The process has converged as soon as every player had the chance to move once.

Convergence to an efficient convention follows from Proposition 2 since in both cases the Pareto optimal symmetric Nash equilibria are strict.

By definition of a coordination game, all strategy pairs on the main diagonal yield strictly higher payoffs for both players than any other strategy pair in their respective row or column, i.e., $\forall$ k

$$\begin{aligned} a_{kk} &> a_{kj} \quad \text{and} \quad a_{kk} > a_{ik} \\ b_{kk} &> b_{kj} \quad \text{and} \quad b_{kk} > b_{ik} \quad \forall\, i, j \neq k. \end{aligned} \tag{2}$$

For symmetric coordination games (2) implies weak acyclity. Hence, convergence is guaranteed by the Theorem. Furthermore, (2) implies that all pure Nash equilibria are strict. Thus, Proposition 2 can be applied to show convergence to an efficient convention.

To prove convergence for the symmetrized version of asymmetric coordination games rename, without loss of generality, the strategies in the original game such that

$$a_{11} + b_{11} \geq a_{ii} + b_{ii} \quad \forall\, i.$$

Together with (2) this implies that

$$a_{11} + b_{11} \geq a_{ij} + b_{ij} \quad \forall\, i,j \tag{3}$$

It is easy to see that Nash equilibria in the asymmetric game become the symmetric Nash equilibria of the symmetrized game. Furthermore, the process of symmetrizing the game implies that the payoffs in the symmetrized game are averages from two cells of the asymmetric game matrix, i.e., if $c_{ij} = (a_{kl} + b_{mn})/2$ then $c_{ji} = (a_{mn} + b_{kl})/2$. Hence,

---

[12] Coordination games are sometimes also defined as games with positive payoffs on the main diagonal and zeros off the diagonal. The results below apply to this definition a fortiori.

$$c_{ij} + c_{ji} = (a_{kl} + b_{kl})/2 + (a_{mn} + b_{mn})/2 \leq (a_{11} + b_{11}) = 2c_{11} \qquad (4)$$

where the inequality follows from (3). In the symmetrized game (c_{11}, c_{11}) is therefore a Pareto optimal equilibrium payoff.

If $a_{ii} \neq b_{ii}$ for some i, there may be payoffs c_{ij} off the diagonal higher than c_{11}. Assume to the contrary that the process converges to one of these payoff vectors. Due to (4) the payoff in any of these vectors can only be higher than c_{11} for one player and must be lower for the other player (the "exploited").[13] Hence, the exploited players will switch to other strategies or towns with the consequence that eventually no one is left to be exploited. Thus, the expected payoff from exploiting falls below c_{11}. From then on all agents with an opportunity to change strategies switch to strategy S_1 (or to a strategy that yields the same payoff). An absorbing state (an efficient convention) is reached once all players had an opportunity to switch to one of these strategies ■

V. INTERACTION BETWEEN TOWNS

The assumption maintained so far that players only interact with players from their own group or town was very strong, but it served well to demonstrate the basic point of the model. Now, I will show that the main results remain valid even if there is interaction between towns. Although most of social or economical interactions take place within a town, players may have to deal with outsiders who might follow different conventions. I will assume that players cannot discriminate between players from their own town and outsiders (i.e. there are no labels on the forehead). Following Hamilton (1964) who defined *viscosity* to be the tendency of animals to have a higher degree of interaction with members of their own (genetical) group, I will define a *viscosity parameter* δ as the probability that a matching partner is drawn exclusively from one's own town (see also Myerson et al., 1991). Accordingly, $1-\delta$ is the probability that a matching partner is drawn from the population at large.

A player who plays strategy i in town j receives the following expected payoff:

$$\Pi_{ij}(\alpha) = \delta\left(\sum_{k=1}^{n} \alpha_{kj} c_{ik}\right) + (1-\delta)\sum_{l=1}^{n} \theta_l \left(\sum_{k=1}^{n} \alpha_{kl} c_{ik}\right)$$

where again α_{kl} is the proportion of players in town l playing strategy k and $\theta_l \equiv N_l/N$ is the share of players located in town l. For $\delta = 1$ one gets the payoffs associated with the model in the first part of this paper. Lower δ's correspond to higher degrees of interaction with players from other towns. A δ equal to zero corresponds to a model without a town structure; matching partners are randomly drawn from the population at large.

---

13 This is the (4,2) payoff in the "battle of sexes" example

Since the town structure played no role in the proofs of the Theorem or Proposition 1, I can state immediately the following proposition.

Proposition 4: For all $\delta \in [0,1]$ *, and for all games that are either acyclic or satisfy ESPS, the evolutionary process converges to a symmetric Nash equilibrium.*

The town structure was the driving force behind the result of Proposition 2. Therefore, one cannot expect the efficiency result to hold if the neighborhood effects are too weak. Clearly, the required δ will depend on the payoff structure and the original partition of players into towns. However, there will always be a $\delta < 1$ such that Proposition 2 continues to hold.

Proposition 5: Consider all games in which the Pareto optimal among the symmetric Nash equilibria are strict. Then $\exists\ \delta^* < 1$, *s.t.* $\forall\ \delta > \delta^*$, *an efficient convention is established in all non-empty towns whenever the process converges.*

Proof: Again, assume to the contrary that the process converges to an inefficient Nash equilibrium (μ,μ) (which must be symmetric by Proposition 1) in some town, even though an efficient, strict equilibrium, say (S_k,S_k), exists. At the beginning of the game all players play the convention of their respective town, i.e. $\alpha_{kl} = 0$, $\forall$ $k \neq l$. Hence, the payoffs from playing S_k in town k are given by

$$\Pi_{kk} = \delta\, c_{kk} + (1-\delta)\sum_{l=1}^{n} \theta_l c_{kl} \qquad (5)$$

Since (S_k,S_k) is strict, i.e. $c_{kk} > c_{ik}$ $\forall$ i,

$$\exists\ \delta^* < 1, \text{ s.t. } \forall\ \delta > \delta^*,\ \Pi_{kk} > \Pi_{ik}\ \forall\ i.$$

Hence no player is going to enter k-town and play a strategy other than the S_k during the entire game (even if $\alpha_{kk} = 0$ at some point during the game). Therefore, (μ,μ) cannot be an absorbing state since at the end players would switch to playing S_k in k-town which yields a strictly higher payoff■

As an example consider the following game and assume that at the beginning of the game $N_A = N_B$.

	A	B
A	8,8	0,4
B	4,0	6,6

Figure 3

For all $\delta > \frac{2}{3}$ strategy A in A-town is the best response against the starting profile $\boldsymbol{\alpha}_A^{(0)} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Hence, some players will switch to A in A-town, with the consequence that A becomes even more attractive. Therefore, the efficient convention (A,A) is the absorbing state.

This should be compared to the results of Young (1993) and Kandori et al. (1993) in which the process driven by random mutations settles on the risk dominant equilibrium (B,B) even if originally everybody plays A. This is due to the fact that there is a small chance of enough players making simultaneously a mistake such that society switches to the other equilibrium. Since it requires fewer mistakes to switch from (B,B) to (A,A) than vice versa, the former will be observed almost all the time as the probability of mistakes goes to zero.

It is interesting to note that experimental evidence seems to reject the view that players would accept inefficient equilibria in coordination games. In the experiments by Van Huyck, Gillette, and Battalio (1992) subjects played a coordination game in which an arbiter gave non-binding recommendations regarding which equilibrium to play. Van Huyck et al. found that most subjects rejected the recommendation of an inefficient equilibrium in favor of the efficient one even though they were given a clear focal point. Unfortunately, the games in Van Huyck et al. (1992) were pure coordination games in which the efficient and the risk dominant equilibrium coincide. It would be interesting to test whether players would still choose the efficient equilibrium even if does not coincide with the risk dominant equilibrium.

VI. CONCLUSION

The purpose of this paper was to demonstrate how an efficient equilibrium in a two player, random matching game can emerge endogenously through an evolutionary process. I assumed that a player in a repeated game adapts to his environment; however, he adapts with a certain inertia in the sense that he cannot always react instantaneously to changes in the environment. The key assumption was that a player has most of his interaction with players from his own neighborhood ("town").

The two main results were the following. First, I showed that the evolutionary process - despite the absence of 'mutations' - converges to a Nash equilibrium for a broad class of games satisfying weak acyclity or *evolutionary stability with respect to pure strategies* (ESPS). This result was independent of the town structure.

Secondly, I proved that an efficient convention emerges for all games in which the Pareto optimal among the symmetric Nash equilibria are strict.

While this paper explains efficiency among the equilibria of the game, it does not yield efficiency among all *payoffs* because only Nash equilibria can be absorbing states of the stochastic process. It will be interesting to study the evolution of cooperation in a prisoner's dilemma situation in which the efficient payoff combination is not a Nash equilibrium of the game. I reserve this question for future work.

APPENDIX

Lemma: If E is a non-singleton ergodic set, then there exists an $\alpha \in E$ *such that* $\alpha_{kl} = 1$ *and* $\alpha_{ij} = 0 \; \forall \; ij \neq kl$.

Proof: Take any $\alpha^* \in E$. Define $E' \equiv \{ \alpha \in E \mid \alpha_{ij} \geq \alpha_{ij}^*$ if $S_{ij} \in B(\alpha^*)$ and $\alpha_{ij} = 0$ if $S_{ij} \notin B(\alpha^*)\}$. Note that any $\alpha' \in E'$, and hence any convex combination of α^* and α', can be reached from α^* in one step with positive probability and, therefore, will be reached with probability one eventually.

(i) If $\exists \; \alpha' \in E'$ and $S_{kl} \in B(\alpha^*)$ s.t.

$$\{S_{kl}\} = \overline{B}(\alpha') \equiv \underset{S_{ij} \in B(\alpha^*)}{argmax} \; \Pi(\alpha')$$

then $\exists \; \varepsilon > 0$, s.t.

$$\{S_{kl}\} = B(\alpha'') \tag{6}$$

with $\alpha'' = \varepsilon\alpha' + (1-\varepsilon)\alpha^*$. Since α'' can be reached from α^*, there exists a positive probability that starting from α'' all players switch to S_{kl}, which produces the desired α.

(ii) If there does not exist an α'' such that (6) holds, then all $S_{ij} \in \overline{B}(\alpha')$ must have the same payoff against any $S_{mn} \in B(\alpha^*)$,

$$\Pi(S_{ij}, S_{mn}) = \Pi(S'_{ij}, S_{mn}), \; \forall \; S_{ij}, S'_{ij} \in \overline{B}(\alpha') \text{ and } \forall \; S_{mn} \in B(\alpha^*). \tag{7}$$

Let $\alpha''' \in E''' \equiv \{\alpha \in E \mid \alpha_{ij} \geq \alpha_{ij}'$ if $S_{ij} \in \overline{B}(\alpha')$, and $\alpha_{ij} = 0$ if $S_{ij} \notin \overline{B}(\alpha')\}$. With positive probability the process moves $\alpha^* \rightarrow \alpha' \rightarrow \alpha'''$. Due to (7) either,

$\overline{B}(\alpha') \subset B(\alpha''')$, in which case α''' is an absorbing state contradicting the assumption that E is a non-singleton ergodic set, or

$\overline{B}(\alpha') \cap B(\alpha''') = \varnothing$. In the latter case, with positive probability all players adopt some $S_{kl} \in B(\alpha''')$, which gives the desired α. ■

REFERENCES

AUMANN, R. and SORIN, S. (1989), "Cooperation and Bounded Recall", *Games and Economic Behavior*, **1**, 5-39.

AXELROD, R. (1984) *The Evolution of Cooperation* (New York: Basic Books).

CANNING, D. (1992), "Evolution of Language Conventions in Common Interest Signaling Games" (mimeo: Columbia University).

CAPLIN, A and NALEBUFF, B. (1992), "Competition among Institutions", Columbia University Discussion Paper No. 621.

ELLISON, G. (1993), "Learning, Local Interaction, and Coordination", *Econometrica*, **61**, 1047-1071.

HAMILTON, W. (1964), "The Genetical Evolution of Social Behavior", *Journal of Theoretical Biology*, **7**, 1-52.

HARSANYI, J. and SELTEN, R. (1988) *A General Theory of Equilibrium Selection in Games* (Cambridge, Mass.: MIT Press).

LEWIS, D. (1969) *Convention: A Philosophical Study* (Cambridge: Harvard University Press).

KANDORI, M. and ROB, R. (1992), "Evolution of Equilibria in the long Run: A General Theory and Applications" (mimeo, Princeton University).

KANDORI, M., MAILATH, G. and ROB, R. (1993), "Learning, Mutation, and Long Run Equilibria in Games", *Econometrica*, **61**, 29-56.

KIM, Y. and SOBEL, J. (1992),"An Evolutionary Approach to Pre-Play Communication" (mimeo, University of California, San Diego).

MATSUI, A. (1991), "Cheap-Talk and Cooperation in a Society", *Journal of Economic Theory*, **54**, 245-258.

MATSUI, A. and ROB, R. (1992), "Evolution, Rationality and Equilibrium Selection in Societal Games" (mimeo: University of Pennsylvania).

MAYNARD SMITH, J. (1982) *Evolution and the Theory of Games* (Cambridge, UK: Cambridge University Press).

MYERSON, R., POLLOCK, G. and SWINKELS, J. (1991), "Viscous Population Equilibria", *Games and Economic Behavior,* **3**, 101-109.

NÖLDEKE, G. and SAMUELSON, L. (1993), "An Evolutionary Analysis of Backward and Forward Induction", *Games and Economic Behavior*, **5**, 425-454.

ROBSON, A. (1990), "Efficiency in Evolutionary Games: Darwin, Nash, and the Secret Handshake", *Journal of Theoretical Biology*, **144**, 379-396.

SAMUELSON, L. (1991), "Limit Evolutionary Stable Strategies in Two-Player, Normal Form Games", *Games and Economic Behavior,* **3**, 110-128.

SELTEN, R. (1980), "A Note on Evolutionary Stable Strategies in Asymmetric Animal Conflicts", *Journal of Theoretical Biology,* **84**, 93-101.

SOBEL, J. (1993), "Evolutionary Stability and Efficiency" (mimeo, University of California, San Diego).

TIEBOUT, C. (1956), "A Pure Theory of Local Expenditures", *Journal of Political Economy,* **64**, 416-424.

VAN HUYCK, J., GILLETTE, A. and BATTALIO, R. (1992), "Credible Assignment in Coordination Games", *Games and Economic Behavior*, **4**, 606-626.

YOUNG, P. (1993), "The Evolution of Conventions", *Econometrica,* **61**, 57-84.