MIGRATION AND THE EVOLUTION OF CONVENTIONS[*]

Venkataraman Bhaskar & Fernando Vega-Redondo[**]

---

[*] This paper was initiated during a visit of V. Bhaskar to the Univeristy of Alicante, financed by the Sabattical Program of the Spanish Ministry of Education. It was completed during a simultaneous visit of both of us to the Instituto de Análisis Económico (CSIC) at Barcelona. We are grateful to this instituttion for their hospitality. Finally, Fernando Vega-Redondo also acknowledges financial support by the Spanish Ministry of Education, CICYT Projec 94-1504.

[**] Venkataraman Bhaskar: University of St. Andrews
Fernando Vega-Redondo: University of Alicante and Instituto Valenciano de Investigaciones Económicas.

MIGRATION AND THE EVOLUTION OF CONVENTIONS

Venkataraman Bhaskar & Fernando Vega – Redondo

ABSTRACT

This paper analyzes an evolutionary model where agents are locally matched to play a coordination game and can adjust both their strategy and location. Their decisions are subject to friction, so that an agent who migrates to a different location may be unable to adjust her strategy optimally to the new environment. A condition on off-equilibrium payoffs introduced by Aumann (1993) plays a major role in our characterization (for general coordination games) of the long-run outcomes. For the particular 2×2 case, this condition (which is unrelated to risk dominance) implies that the possibility of medium term simultaneous *co-existence* of conventions at different locations depends on whether the game is of "pure" coordination (where co-existence is always possible) or of the stag-hunt type (where it is not). When we introduce noise (i.e. mutations) into the model, this distinction continuous to play a crucial role in the selection of the long-run equilibria: for large friction, *both* equilibria are stochastically stable in the former case, whereas only the efficient one is so in the latter.

KEYWORDS: Migration; Evolution; Conventions.

1 Introduction

The ubiquity of multiple equilibria is perhaps one of most important insights offered by game theory to economics. Simple 2×2 games can have multiple Pareto-ranked equilibria, and this suggests that differences in economic performance between societies may be related to different conventions, i.e. the playing of different equilibria. Several questions arise: How are conventions established? How do they change? Can different conventions co-exist? Under what conditions can one expect efficient conventions to prevail?

The recent evolutionary literature on equilibrium selection in games has shed important light on these issues. Most of this literature (e.g. see Foster and Young [6], Kandori, Mailath & Rob [9], or Young [13]) has focused on models where agents in a single homogeneous society interact over time. This literature shows that a single convention must prevail. Furthermore, in 2×2 games, the stochastically stable convention is the risk-dominant one, so that efficiency is not ensured.

These conclusions continue to hold if interaction is local rather than global, as Ellison [3] shows. Ellison studies a model in which players are *fixed* at different locations, and the pattern of interaction is pre-specified by an overlapping neighborhood structure. By indirectly linking the actions of different players, this neighborhood structure ensures that, as in the case of global interaction, only a single convention can prevail. The key difference arising in this case concerns the rate of convergence towards the risk-dominant convention: in contrast with the context with global interaction, this convergence rate turns out to be much faster (and independent of population size). While allowing for local interaction is clearly of substantial interest, it is also important to take into account that agents typically have discretion with regard to location choice, and hence some freedom in choosing their neighbors. In other words, rather than fixing exogenously the pattern of interaction, we need to understand how interaction patterns and strategic behavior *co*-evolve.

Three recent papers by Oechssler [11], Ely [5], and Mailath, Samuelson & Shaked [10] explore models with endogenous location interactions. Oechssler studies an evolutionary model where agents are initially distributed over a given set of independent "cities" and, over time, may *freely* adjust both their strategic and locational decisions. Assuming that all conventions are represented at the start of the process (i.e. are adopted by some city), he shows that the efficient one will eventually prevail throughout. The intuitive reason why this will occur is simple: any agent, when given the opportunity to adjust, will immediately shift to a city where the efficient convention is played (if not already there). So doing (and changing her strategy accordingly), she is sure to meet only agents who play the efficient strategy, thus achieving the maximum available payoff.

The paper by Ely considers a model somewhat similar to Oechssler's in the context of 2×2 coordination games. He does away with the assumption that both conventions are initially present, focusing upon stochastically stable states when players decisions are subject to small mutations. As in Oechssler's scenario, conventions cannot co-exist, since players playing the inefficient equilibrium will migrate to a location playing the efficient convention when the opportunity arises. For this reason, destabilizing the inefficient convention is easy – it is sufficient that a single player move to an empty location and start playing the efficient convention, in order for migrants formerly playing the inefficient convention to switch away from the latter to the former. On the other hand, destabilizing the efficient convention requires a large number of mutations, since this requires that a substantial number of players switch to the inefficient strategy at a single location, thus making it optimal for the others to switch their strategy as well. Hence only the efficient convention is stochastically stable.

The paper by Mailath *et al.* contemplates a quite different context with two continuum populations. They assume that the players of each population are fixed at some respective finite set of locations (or activities). This affects whom they meet as the result of two types of considerations. On the one hand, each player of population $i = 1, 2$, makes a decision as to what location of population j to visit. Reciprocally, she may also receive some visits by individuals of population j who decide to visit her location. They establish that if agents have some control over their interaction pattern (in particular, if they are able to avoid any undesired matching), then every configuration which is locally stable in terms of a suitably monotonic evolutionary system must be efficient.[1]

The present paper shares some similarities with those summarized above, but also displays some key differences. The interaction framework is the simple one of Oechssler [11] or Ely [5], with players distributed among a certain number of locations ("cities"). By choosing to locate in one of them, an

---

[1] Another interesting result of Mailath *et. al.* is their demonstration that an equilibrium with *fixed* interactions is formally equivalent to a correlated equilibrium of the underlying game.

agent selects the corresponding matching pattern induced by its population profile. A key difference resides in the fact that, in our model, players' adjustment is subject to friction. In a given period, a player may or may not be able to migrate, but if she does migrate to another location, she is never fully sure that she will be able to adjust their strategy in the optimal direction in the new environment. Although there are several possible reasons for allowing uncertainty about future strategy adjustment, one of the most natural interpretations is that migrants are never sure that they will be able to adapt to a foreign "culture". This uncertainty has the consequence that the model no longer produces the unqualified efficiency conclusion of Oechssler [11] or Ely [5]. In particular, and depending on the particular payoff structure, one may have convention *co*-existence in the medium term (i.e. in the absence of mutations), and even in the long run efficiency may not be ensured.

The condition on off-equilibrium payoffs which plays a critical role in our model is orthogonal to considerations of risk dominance, but is related to a distinction enunciated by Aumann [1] in a completely different context. Aumann discussed the effectiveness of pre-play communication between rational players who are playing a 2×2 coordination game, and argued that communication leads to efficiency if the game is one of pure coordination (i.e. one where the off-equilibrium payoffs are uniformly lower than the payoffs in either equilibrium), but not if the game is of the stag-hunt kind. Our results for 2×2 coordination games hinge upon the same distinction, but are, in a sense, the opposite of Aumann's. We find that convention co-existence is possible in the *medium term* if the game is one of pure coordination and friction is significant, but is not possible in a stag-hunt game.

Once we allow for mutations (i.e. in the ultra-long run),[2] *simultaneous* co-existence of convention is not possible. However, in a pure coordination game, *both* equilibria are stochastically stable (thus, in a sense, co-existence is displayed over time, with both conventions being played at comparable time frequencies). In contrast, the efficient equilibrium turns out to be the unique stochastically stable convention in the stag-hunt game. In our formal analysis, these conclusions will be generalized in a natural way to the *full class* of coordination games with arbitrarily many strategies.

Finally, we explore the robustness of our conclusions by considering a "slight" modification of the model where a small fraction of players is unable to migrate in some (possibly very few) locations. Our previous conclusion

---

[2] See Binmore *et al.* (1995) for a discussion of the different time horizons (medium, long, and ultra-long runs) that are typically considered in evolutionary models.

for stag-hunt games is maintained, i.e. efficiency prevails in the long run, irrespectively of risk considerations. For pure coordination 2 × 2-games, on the other hand, the formerly exposed weakness of the efficient convention becomes drastically reinforced. For, in this latter case, it turns out that *only* the risk dominant convention (possibly inefficient) is stochastically stable when friction is significant. This again stresses the limitations of migration as a mechanism for achieving efficient outcomes when the underlying game is of pure coordination.

The rest of the paper is organized as follows. The following section presents the basic model, where all agents are potentially mobile across locations. Section 3 presents and discusses our results for this context. Section 4 turns to the case where some locations have a few immobile players. The final section concludes with a summary. For the sake of smooth discussion, all formal proofs are relegated to an Appendix.

2 The Basic Model

Consider a population of n (≥ 3) individuals who are matched in pairs to play a finite bilateral symmetric game with strategy set $S = \{s_1, s_2, ..., s_Q\}$ and payoff function $\pi : S \times S \to \Re$, where $\pi(s_q, s_r)$ is interpreted as the payoff earned by a player adopting s_q when matched with an opponent who plays s_r.

Matching between players takes place at specific locations. Let $\mathcal{L} = \{\ell_1, \ell_2, ..., \ell_I\}$ stand for the set of possible locations ($I \geq 2$). A sufficient description of the current state of the system specifies how many agents are playing each of the Q actions at everyone of the I locations. Thus, we can think of a typical state as an $I \times Q$ matrix ω whose entries $\omega_{iq} \in \mathbb{N} \cup \{0\}$ satisfy

$$\sum_{\substack{i=1,...,I \\ q=1,...,Q}} \omega_{iq} = n.$$

The set of all possible states will be denoted by Ω.

For each of the encounters taking place at any given location, matching will be assumed randomly uniform across all possible realizations. In other words, we shall assume that the *ex-ante* probability that a given player meets any other one in her own location is identical across all other players, and independent of any other encounters. Further assuming that the number of encounters in a certain location is proportional to the number of individuals in this location,[3] the expected payoff of playing strategy s_q at location ℓ_i is captured by (i.e. is proportional to) the following expression:

$$\hat{\pi}(s_q, \omega_i) = \frac{\pi(s_q, s_q) \cdot (\omega_{iq} - 1) + \sum_{r \neq q} \pi(s_q, s_r) \cdot \omega_{ir}}{n_i - 1}$$

where $\omega_i \equiv (\omega_{iq})_{q=1,...,Q}$ stands for the *ith* row of the matrix ω, and $n_i \equiv \sum_{q=1,...,Q} \omega_{iq}$ denotes the total number of players located at ℓ_i.

It remains to specify the payoff of a player who happens to be alone at a certain location and, therefore, cannot be matched against anyone else. We shall normalize the payoff of such "sad loner" to zero, further assuming that playing the game always entails positive payoffs (i.e. $\pi(s_q, s_r) > 0$ for all $q, r = 1, 2, ..., Q$). This is a convenient assumption that will simplify matters thereafter by ensuring that, in any limit state, all players will be part of some "city." It could be generalized at the cost of some formal complexity without affecting the essential nature of our analysis.

We now turn to the dynamics of the model. Time is measured discretely and indexed by $t = 0, 1, 2, ...$ At each t, players may adjust both their strategy and location in the following two-stage fashion. First, every player is subject to an independent draw of a dichotomous random variable which determines whether or not she can change her location (i.e. "migrate"). Once this first stage is over and the agent has irreversibly decided (when given the option) whether and where she migrates to, she is assumed subject to an additional independent draw of another dichotomous random variable. This second draw determines whether she will be able to adjust her strategy within the location determined in the first stage. For simplicity, we assume that all individuals face and identical and stationary probability $\lambda \in (0, 1)$ of being able to change location every period. Analogously, at every t, all players are assumed to face an identical and stationary probability $\theta \in (0, 1)$ of being able to change their strategy.

The novel and key feature of our adjustment scenario is that action and location revision possibilities are not deterministically coupled. This implies that, in general, agents might not be able to evaluate "safely" the merits of a locational decision, since the realization of its potential benefits could require some subsequent adjustment of their action. We think of this as some stylized form of *friction*, which admits a number of different interpretations. For example, a natural one is based on the idea that attempting to adopt a new action generally requires undergoing a learning process, the outcome

---

[3] Note that this formulation implies that the expected number of encounters enjoyed by any player is constant across all locations where there are *any* partners.

of which cannot be guaranteed *ex ante*. If it were unsuccessful, the agent would have to stay with her previous action, then possibly "regretting" having migrated.

Given the adjustment scenario described, we now consider how players will react to it. We shall assume that when agents are able to revise their former decisions, they select their new location/action pairs so as to increase their expected payoff. In carrying out the required payoff comparisons, players are taken to hold the expectation that the configuration of the system remains fixed (except, of course, the part of it which concerns herself). Thus, in other words, we postulate that agents hold "static expectations" on the evolution of the process, in line with what is the customary approach in much of received evolutionary theory.

Proceeding now formally, we start by considering the second stage of the adjustment process, with a given individual at location ℓ_i receiving a strategy revision opportunity at some t. Let s_q be this player's current strategy and denote by $\tilde{\omega}_i$ her expectations on the strategy configuration prevailing at ℓ_i. Since these expectations are assumed formed in a static fashion, we write

$$\tilde{\omega}_i = \omega_i(t-1) + \delta\, e_q, \tag{1}$$

where $\delta = 1$ if the player concerned has migrated in period t, $\delta = 0$ otherwise, and e_q stands for the Q-dimensional vector $(0, ..., 0, 1, 0, ..., 0)$ whose q-th component is 1. Consider now the set $\{s_r : \tilde{\pi}(s_r, \tilde{\omega}_i + e_r - e_q) \geq \tilde{\pi}(s_q, \tilde{\omega}_i)\}$, where recall that $\tilde{\pi}(s_q, \omega_i)$ stands for the expected payoff of strategy s_q facing configuration ω_i. Every strategy in this set (consisting of those strategies whose payoff is expected to be no smaller than that of s_q) is assumed chosen by the player in question with positive probability.

Next, we address the first-stage migration decision for an agent who is currently playing action s_q at location ℓ_i and is contemplating migrating to ℓ_j. It is assumed that if this agent is given the opportunity to migrate, she will choose location ℓ_j with positive probability if, given static expectations, her expected payoff from migration is greater than the expected payoff resulting at her current location. In order to compute these expected payoffs, we suppose that the agent has already decided at this point upon the particular action that she will choose in each of these alternative locations *if* she is able to adjust her strategy. This clear-cut formulation is adopted here for the sake of simplicity, but could be substantially generalized at some formal complexity. For example, we could admit that the agent has non-degenerate expectations about her future choice, provided that there is sufficient stochastic variability in how these expectations are formed.

Formally, consider the individual who is currently adopting strategy s_q at location ℓ_i , and who anticipates changing her action to $s_{q'}$ in the event that she does not migrate. Then, her expected payoff from not migrating is

$$(1-\theta)\hat{\pi}(s_q, \omega_i(t-1)) + \theta\hat{\pi}(s_{q'}, \omega_i(t-1) + e_{q'} - e_q). \tag{2}$$

Since $s_{q'}$ must be a weakly better response than her current strategy s_q, this implies that the minimum expected payoff from *not* migrating is $\hat{\pi}(s_q, \omega_i(t-1))$, i.e. the payoff obtained when the player simply stays with s_q as the (weakly) "better response". On the other hand, if the player migrates to location ℓ_j and anticipates adopting strategy s_r in response to a revision opportunity, her expected payoff is

$$(1-\theta)\hat{\pi}(s_q, \omega_j(t-1) + e_q) + \theta\hat{\pi}(s_r, \omega_j(t-1) + e_r). \tag{3}$$

The maximum of this latter expression is achieved when s_r is a *best response* to the strategy configuration $\omega_j(t-1)$ predicted at location ℓ_j. Therefore, the player will migrate to ℓ_j with positive probability (in general, less than one since there may be other locations which are also better than ℓ_i) if, and only if, the maximal expected payoff from migration to ℓ_j (achieved when s_r in (3) coincides with a *best* response) is greater than the minimum payoff from not moving (achieved by making $s_{q'}$ equal to s_q in (2)).

Given s_q, define

$$\pi^*(s_q, \omega_j(t-1)) \equiv (1-\theta)\,\hat{\pi}(s_q, \omega_j(t-1)+e_q) + \theta \max_{s_r \in S} \hat{\pi}(s_r, \omega_j(t-1)+e_r).$$

for every $\omega_j(t-1)$, $j \neq i$. Combining all of the above considerations, the set of locations ℓ_j to which the agent in question (originally in location ℓ_i) can migrate may be identified as follows:

$$\{\ell_j \in \mathcal{L} : \pi^*(s_q, \omega_j(t-1)) \geq \hat{\pi}(s_q, \omega_i(t-1))\}\ .$$

The process described defines a finite Markov chain on the state space Ω. In the following section we analyze this process and characterize its limit states. Subsequently, we also analyze a perturbation of this Markov chain in which individuals occasionally experiment (or mutate).

3 Analysis

Our analysis will be concerned with strict coordination games, i.e. games where the set of (strict) pure strategy equilibria consists of all diagonal elements of $S \times S$. The results may be substantially extended to more general contexts.[4] However, for the sake of focus, it has seemed best to restrict our discussion to the standard context of coordination games, where our main points stand out in a simpler and clearer fashion.

The analysis is decomposed in two parts. Firstly, we characterize the limit (or minimally absorbing) sets of the unperturbed dynamics. Secondly, we address the issue of stochastic stability, i.e. we identify those limit sets which are selected in the long run when the original dynamics is perturbed by some small mutational noise.

3.1 Limit Sets of the unperturbed dynamics

Let $T \in \Re^{|\Omega|} \times \Re^{|\Omega|}$ denote the transition matrix of the unperturbed dynamics, as described in the previous section. We start by defining the concept of an absorbing set. Given $\omega \in \Omega$, let $\Gamma(\omega)$ denote the set of states that are accessible from ω (i.e. if $\omega' \in \Gamma(\omega), T^u(\omega, \omega') > 0$ for some $u \in \mathbb{N}$).[5] Extend now these notions to sets, so that $\Gamma(A)$ is the set of states accessible from some $\omega \in A$. Then, a set A is called *absorbing* if $\Gamma(A) \subset A$. Furthermore, a set A is said to be a *limit set* of the Markov chain if it is a minimal absorbing set.

First, we abstract from considerations of migration and show that play at any location must settle down to a convention. To do this, assume for the moment that there is only a single location, so that the vector ω consists only of the actions played at location 1, i.e. $\omega = \omega_1$.

Proposition 1 *Let there be only a single location. If A is a limit set of the unperturbed dynamics, then $A = \{\omega\}$ and $\omega_{1q} = n$ for some $q \in \{1, 2, ...Q\}$, i.e. all players play the same strategy.*

Proof. See the Appendix.

Proposition 1 shows that in the absence of migration, the process of strategy adjustment converges to a convention, with all players playing the same strategy. This proposition may be of some independent interest, since

---

[4] For example, we have studied games with general symmetric curbsets (i.e., which are not necessarily singletons, as for strict coordination games). If the above formulation is generalized by allowing players to rely on sufficiently long past history in forming their expectations, the essential gist of our results extends to this context.

[5] The notation used here is standard. $T(\omega, \omega')$ stands for the transition probability from ω to ω' (or the (ω, ω')-entry of the matrix T). On the other hand, $T^{(u)}$ stands for the transition matrix resulting from u iterations of the process.

the result holds for a wider class of games than the result proved by Kandori and Rob [8], who prove a similar result for the subclass of pure coordination games where the payoffs to off-diagonal elements of $S \times S$ are all zero.[6]

As we shall see below, not only is it the case that play at any location converges to a convention, i.e. at most one strategy is played in each location. Furthermore, any strategy is played in at most one location. To *characterize* the set of strategies which may be played, we define the following binary relation on S. Given θ (the strategy-revision probability), define a binary relationship $R(\theta)$ on S as follows. For any pair of strategies s, $s' \in S$,

$$s\, R(\theta)\, s' \Leftrightarrow \pi(s,s) > \theta\, \pi(s',s') + (1-\theta)\, \pi(s,s') \ .$$

Observe that $R(\theta)$ is necessarily irreflexive, and is not transitive. However, if s and s' are distinct, then we have either $sR(\theta)s'$ or $s'R(\theta)s$. To see this, note that $\pi(s,s) \geq \pi(s',s')$ implies $s\, R(\theta)\, s'$, since all elements of S are strict Nash equilibria.

In terms of the relation $R(\theta)$, we define the following collection of subsets of S :

$$\Phi(\theta) \equiv \{D \subseteq S : \forall s, s' \in D : s \neq s',\ s\, R(\theta)\, s' \ \wedge\ s'\, R(\theta)\, s\} \ . \tag{4}$$

Of course, $\Phi(\theta)$ includes all the singleton subsets of S. The interesting fact is that, in general, it will also contain subsets with more than one strategy. A subset of S with more than one element belongs to $\Phi(\theta)$ if, and only if, all of the strategies included in it are maximal with respect to $R(\theta)$ restricted to this set.

Later on, we shall also be interested in another construction derived from the relation $R(\theta)$:

$$\varphi(\theta) \equiv \{s \in S : s\ R(\theta)\ s',\ \forall s' \in S, s' \neq s\},$$

which defines the set of maximal elements of $R(\theta)$ in S.

To illustrate these concepts, we turn to a simple 2×2-game, whose general payoff structure is represented in the following table:

2 1	H	L
H	a, a	b, c
L	c, b	d, d

Table 1

---

[6] These authors study a somewhat different dynamics, i.e. a best-response dynamics, rather than the better-response dynamics that we consider here.

Suppose that we are in the presence of a coordination game, so that $a > c$ and $d > b$, further assuming w.l.o.g. that $a > d$, i.e. H is the efficient strategy. Consider now the following two alternative scenarios. First, if $c \geq d$, we call it a *Stag-Hunt Game.* On the other hand, if $c < d$ we call it a *Pure Coordination Game.* The essential difference between both cases does not dwell at all in the standard considerations of risk-dominance which are so prevalent in the literature of equilibrium selection. (In fact, one could say that they are "orthogonal" to them, since risk dominance depends upon the relative sizes of the differences $a - c$ and $d - b$, while our categorization depends only on the *sign* of $d - c$.)

The distinction between stag-hunt and pure-coordination games was first elucidated by Aumann [1] in his discussion of the stability of the inefficient equilibrium (L, L) when players are able to communicate prior to the play of the game. Aumann argued that in a pure coordination game, the inefficient equilibrium is easily de-stabilized, since an announcement by one player, say player 1, that he intends to play H should be believed by player 2. This is because if player 1 were in fact intending to play L, then he would like player 2 to also play L, and hence has no interest in persuading player 2 to play H. On the other hand, in a stag-hunt game, if player 1 plans to play strategy H , he cannot credibly communicate this intention to the other player. For, even if he were planning to play L, he would be interested in convincing player 2 that he will play H. Thus, since the partner should understand this, whatever signal or communication he receives from the player in question carries no informative content whatsoever and she might as well fully ignore it. Hence pre-play communication does not ensure efficiency in a stag-hunt game, where both (H, H) and (L, L) are stable. To summarize, Aumann's argument implies that one should observe only H being played in a society where players can communicate if the game is a pure coordination game, but may observe both (H, H) and (L, L) if the game is a stag-hunt game.

Let us now relate this discussion to the above defined constructs, $\Phi(\theta)$ and $\varphi(\theta)$. In a stag-hunt game $\Phi(\theta) = \{\{H\}, \{L\}\}$ for every value of θ. In contrast, in a pure coordination game, there is a small enough value of θ (specifically, it is enough that $\theta < \frac{d-c}{a-c}$) for which $\Phi(\theta) = \{\{H\}, \{L\}, \{H, L\}\}$, i.e. not only the two singletons but also the set that includes both strategies is included in $\Phi(\theta)$. As we shall see shortly, this implies that two conventions may co-exist if the game is of pure coordination, but cannot co-exist if it is a stag-hunt game. In other words, our evolutionary analysis will turn out to depend upon exactly the same condition on payoffs as Aumann's, but with diametrically opposite results – co-existence is possible in our model precisely when it is impossible in Aumann's case, and vice-versa.

This disjunction also applies in the long run, when we allow for mutations, as shown in section 3. In the long run, the simultaneous co-existence of two conventions is not possible; however inter-temporal co-existence may well be possible, depending upon payoffs. The relevant construct here is $\varphi(\theta)$. For a stag-hunt game, $\varphi(\theta) = \{H\}$ for all values of θ, whereas $\varphi(\theta) = \{H, L\}$ in the case of pure coordination for values of $\theta < \frac{d-c}{a-c}$. Again, as discussed below, this will manifest itself in the different performance of the model concerning its stochastically stable outcomes. In the case of a stag-hunt game, the unique stochastically stable outcome is one where every player plays H. In contrast, both monomorphic states (everyone playing H or everyone playing L) will turn out to be stochastically stable in a pure coordination game, provided θ is small enough.

We our now in a position to characterize the limit sets of the unperturbed dynamics. Such characterization is contained in the following result.

Theorem 1 *Any given $A \subset \Omega$ is a limit set of the unperturbed dynamics if, and only if, it is a singleton of the form $A = \{\omega\}$, and the state ω satisfies:*
a) For each location ℓ_i, $i = 1, 2, .., I$, there exists at most one strategy s_q such that $\omega_{iq} > 0$.
b) For each strategy s_q, $q = 1, 2, ..., Q$, there is at most one location ℓ_i such that $\omega_{iq} > 0$.
c) For all $i = 1, 2, ..., I$, $n_i = \sum_{q=1} \omega_{iq} \geq 2$.
d) The set of strategies played at ω, $\{s_q \in S : \omega_{iq} > 0, i = 1, ..., I, \}$, is an element of $\Phi(\theta)$.

Proof. See the Appendix.

This theorem fully characterizes the set of states (and corresponding strategies) played at any limit set of the unperturbed dynamics. We interpret this result as indicating what type of outcomes we could expect from the model in the medium run, i.e. within a time frame in which the mutations contemplated in the following subsection may be thought to have little impact on the process. Under these conditions, it is clear that the initial conditions will generally have an important effect – *a priori*, we should expect that different outcomes could materialize depending on where the process starts.

An important conclusion in this respect is that the structure of payoffs off-equilibrium will play a critical role in determining whether co-existence of conventions can be observed across different locations. In particular, if we focus again on the illustrative context provided by 2×2-coordination games, co-existence may arise in a game of pure-coordination if θ is not too large, but not in a stag-hunt game. What is, for our purposes, the essential distinction between these two kinds of games? Heuristically, the main feature at work in stag-hunt games is that any player adopting strategy L prefers to interact with other players choosing H (despite being "uncoordinated" with them) rather than with those of her own kind. This obviously represents a strong force towards the disruption of co-existence, even for small θ, through the migration of agents originally playing the L convention to locations where the H convention prevails. In contrast, such strong force for migration away from an inefficient convention does not materialize any longer if the game is of pure coordination. In this case, the fear to suffer a penalty from uncoordination will deter migration (and thus the disruption of co-existence) if friction is sufficiently significant. These considerations would seem to suggest a possible explanation for different economic performance in neighboring regions (e.g. North vs. South Italy) based on the persistence of different conventions being played in each of them despite the possibility of migration.

Finally, the question may be raised of whether our result on possible co-existence is robust to the specification of the neighborhood structure. More specifically, suppose that instead of isolated cities, we had an overlapping neighborhood structure where agents at each location also interact with those of neighboring locations, in a way reminiscent of Ellison [3]. Might this alternative formulation not destroy the possibility of co-existence? While not providing a full analysis of such a model, we believe this not to be the case.

Suppose, for example, that we have a finite set of locations (say $2n$ of them) sequentially numbered and arranged along a circumference. Further assume that agents at location ℓ_i are not only matched among themselves but also with agents located at locations ℓ_{i+1} and ℓ_{i-1}. Consider now any initial state with a concentration of agents near location ℓ_1 playing s_1, and a concentration around ℓ_n playing s_2, intermediate locations having possibly a mixed strategy profile. If agents play a pure coordination game and θ is sufficiently low, then any agent at an intermediate location will want to migrate (to location ℓ_1 if she is currently playing s_1, or location ℓ_n if she is currently playing s_2). Hence with positive probability, the final outcome will consist of separated cities with distinct conventions, much like our model. Contrast this with Ellison's analysis where, since locations are fixed, strategies must adjust until they are uniform throughout.

3.2 Stochastic Stability

We now discuss a perturbation of the contemplated Markov chain where individuals mutate. More precisely, it is postulated that, with some stationary and independent (small) probability ϵ, every individual adopts an arbitrary locational and strategic decision, unrelated to any of the considerations described above. For concreteness, it will be assumed that, in this event, the individual in question adopts any pair of strategic and locational decisions with uniform probability.

This perturbation may be rationalized on different (non-exclusive) grounds. One possibility is to think of the mutation process as modelling a very gradual process of population turnover. When a newcomer replaces an incumbent, she operates from *tabula rasa,* thus choosing any strategy-location profile with the same probability. Another alternative way of motivating the exercise is as a test of robustness. Specifically, we explore what kind of long-run behavior is robust to the introduction of small noise into the process.

In any case, one important practical implication of the contemplated perturbation is that it obviously ensures that the process is ergodic. Thus, in particular, it has a unique well defined invariant distribution which fully summarizes the long run behavior of the system. Denote by $\mu_\epsilon \in \Delta(\Omega)$ such invariant distribution, in order to explicitly reflect its dependence on the mutation probability. Since, as explained, we want to think of ϵ as small, we shall be interested in the limit invariant distribution which results when $\epsilon \longrightarrow 0$. Such *limit invariant distribution*

$$\mu^* \equiv \lim_{\epsilon \to 0} \mu_\epsilon \tag{5}$$

will be shown to be a well defined element of $\Delta(\Omega)$.

When enlarging out the time scale so that mutations can significantly affect the dynamics of the process, the appropriate limit notion to study is that of stochastic stability. Heuristically, a state is *stochastically stable* if it is visited a significant fraction of time in the long run, when a small rate of mutation does affect the dynamics of the process. This is precisely the idea captured by the limit invariant distribution of the process, as defined in (5). Since, as explained, the stochastic process is ergodic for any given $\epsilon > 0$, this probability distribution identifies the long-run frequencies observed along any sample path of it, almost surely (a.s.).

Note, of course, that in order for a state to be stochastically stable, it must necessarily belong to some limit set of the unperturbed dynamics. Thus, in view of Theorem 1, we may equivalently speak of a stochastically stable state or set. For the sake of completeness, we provide the following formal definition.

Definition 1 *The state* $\omega \in \Omega$ *is* stochastically stable *if* $\mu^*(\omega) > 0$, *where* $\mu^* \in \Delta(\Omega)$ *is as defined in (5).*

We are now in a position to state our main result.

Theorem 2 *A state* $\omega \in \Omega$ *is stochastically stable if, and only if, it satisfies:*
(i) $\exists i \in \{1, ..., I\}$, $\exists q \in \{1, ..., Q\}$ *s.t.* $\omega_{iq} = n$.
(ii) The (unique) strategy s_q *played at* ω *belongs to the set* $\varphi(\theta)$.

Proof. See the Appendix.

In contrast with Theorem 1, the above result establishes that no simultaneous co-existence of conventions may prevail in any of the long-run states. However, depending on the nature of the payoff structure (which affects the set $\varphi(\theta)$), the process is consistent with what we could interpret as inter-temporal coexistence. For example, in the context of the 2×2-game described in Table 1, the two states where the whole population alternatively play H or L will be observed a significant fraction of the time in the long run if we have a pure coordination game and $\theta < \frac{d-c}{a-c}$. (In this case, $\varphi(\theta) = \{H, L\}$.) This contrasts with those results in the literature (e.g., those by Oechssler [11], Ely [5] or Mailath, Samuelson & Shaked [10] summarized in the Introduction) where, by allowing agents to adjust their locational choices as the process evolves, efficient configurations arise as the long-run outcomes.

To gain some basic intuition for our result, recall our former heuristic discussion (following Theorem 1) on the different implications of pure-coordination and stag-hunt games with respect to the possible co-existence of conventions. As it was explained, stag-hunt games induce strong migration forces towards locations where the efficient convention is being played, whereas these forces do not exist for the alternative kind of games when friction is significant. It is precisely these same considerations which underlie the differential fragility of each type of convention in the presence of mutations. For games of pure coordination, both conventions (H and L) are comparably robust to mutations since both exhibit a similar difficulty in being disrupted by migration. In contrast, if the game is of the stag-hunt type, the inefficient convention is extremely fragile to the appearance of someone playing the efficient strategy in some other location. If this happens by way of mutation, a wave of migration can be immediately triggered (with some positive probability), leading to the eventual collapse of the inefficient convention.

New light will be shed on these ideas with the analysis carried out in the next section of a context with a few immobile players. To conclude the present discussion, it is important to emphasize that it is the introduction of friction into the adjustment framework which crucially underlies our conclusions. For, if θ is made sufficiently large, thus largely removing the practical implications of friction, our framework approximates that of former literature, leading as well to similar conclusions. (Specifically, it is clear that, if θ is large enough, $\varphi(\theta)$ includes only the efficient strategy.)

Remark: Even though issues related to rate of convergence are peripheral to our concerns in this paper, it may be worthwhile to make the following point. In our model, the maximum expected waiting time to observe any long-run (i.e. stochastically stable) outcome is of the order ϵ^{-1}, where ϵ is the mutation probability.[7] Thus, in contrast with most of existing evolutionary models, the rate of convergence towards long-run outcomes is very fast and *independent* of either the population size or the payoff structure of the game. To the extent that this rate of convergence is a natural basis for assessing the relevance of a long-run prediction, those of our model appear to be on relatively firm grounds.

As advanced, in order to attain a clearer understanding of the nature and robustness of our results (specifically, on the role played by player mobility), it is useful to consider a context where some players (possibly very few) may have their mobility restricted. This is our purpose in the following section.

4 A Few Immobile Players

Consider a variation of the original model where *some* individuals are assumed immobile, i.e. may adjust their action but *not* their original location. For a sharper contrast with our previous analysis, we shall focus our attention on simple coordination games with just two strategies: s_1 and s_2, where, say, s_1 is taken to be the efficient strategy. That is:

$$\pi(s_1, s_1) > \pi(s_2, s_2).$$

This is a game as described in Table 1 above, where s_1 is identified with H and s_2 with L.

---

[7] An easy way to confirm this statement is to observe that, in the language of Ellison [4], the co-radius of any of our stochastically stable equilibria is 1.

Suppose now that some (arbitrary and non-empty) subset of locations $\mathcal{L}_0 \subseteq \mathcal{L}$ have a few immobile players. Since the particular number of immobile players per location have no implications for our results we shall simplify formal details by assuming that all locations in $\mathcal{L}_0$ have m immobile individuals (the same number, assumed odd[8] and larger than 2, in all of them). Moreover, we shall think of the immobile players as representing only a small minority relative to the whole population (i.e. $\frac{m|\mathcal{L}_0|}{n}$ being very small), although this is not essential.

Since the model is, in every other respect, fully parallel to that described in Section 2, we dispense here with a formal description of it. It only needs to be noted that, due to the additional restriction experienced by immobile players, the state space has to specify separately the matrices of action-location entries for each of the two types of individual, mobile and immobile. That is, a state ω requires the specification of two $I \times Q$-matrices (ω^1, ω^2) with the obvious adding-up conditions to be satisfied by each of them. (For example, if ω^1 stands for the matrix corresponding to immobile players and location $\ell_i \in \mathcal{L}_0$, then we must have $\sum_{q=1}^{Q} \omega_{iq}^1 = m$.)[9]

In this context, the essential gist of our previous analysis will be reinforced. In particular, it will continue to be the case that in games of pure coordination the efficient convention is not necessarily selected in the long run (provided, of course, friction is significant). However, an interesting new aspect introduced by the presence of immobile players is that, in those games of pure coordination, a *single* convention is now selected in the long run, in contrast with what occurred in the original context where all players were assumed mobile (recall our discussion following Theorem 2). Moreover, this convention turns out to coincide with the risk dominant strategy s_r satisfying

$$\pi(s_r, s_r) + \pi(s_r, s_{r'}) > \pi(s_{r'}, s_r) + \pi(s_{r'}, s_{r'}) \ ,$$

for $r \neq r'$. Note, of course, that it may well happen that $s_r \neq s_1$, i.e. the risk dominant strategy is the inefficient one.

On the other hand, if the game is of the stag hunt type, the analysis remains essentially unaffected by the introduction of a few immobile players. In this case, that is, the (single) long-run strategy continues to be (for almost

---

[8] The assumption that m is odd is just a simplifying assumption which guarantees that $\frac{m-1}{2}$ is an integer. It could be dispensed with in Theorem 3 below by assuming that (given the payoff structure of the game) m is large enough.

[9] Since no confusion should arise, we keep the same notation as before for both the state space and the particular states.

all players, including every mobile one) the efficient one, as in the original context.

The above summarized conclusions are formally presented in the following result.

Theorem 3
(a) Assume $\theta < \frac{\pi(s_1,s_1)-\pi(s_1,s_2)}{\pi(s_2,s_2)-\pi(s_1,s_2)}$,[10] *and let* s_r *be the risk-dominant strategy. Then a state* $\omega \in \Omega$ *is stochastically stable if, and only if, it satisfies:*[11]

$$\sum_{\ell_i \in \mathcal{L}} (\omega_{ir}^1 + \omega_{ir}^2) = n. \tag{6}$$

(b) Conversely, if $\frac{\pi(s_1,s_1)-\pi(s_1,s_2)}{\pi(s_2,s_2)-\pi(s_1,s_2)} \leq \theta$*, then a state* $\omega \in \Omega$ *is stochastically stable if, and only if, it satisfies:*

$$\sum_{\ell_i \in \mathcal{L}} (\omega_{i1}^1 + \omega_{i1}^2) = n - m(|\mathcal{L}_0| - 1). \tag{7}$$

Proof. See the Appendix.

Part (b) of Theorem 3 simply establishes that, for stag hunt games, the conclusion obtained in the previous section is robust to the introduction of a few immobile players. Part (a), on the other hand, is formally different to the conclusion obtained before, although it displays a similar spirit. It indicates that when the game is of pure coordination, the risk-dominant convention (possibly inefficient) is achieved in the long run. Since clarifying the mechanism leading to this latter conclusion should prove useful in understanding the general approach proposed in this paper, we discuss it next in some detail.

In every one of our different contexts (both here and in the previous section), the easiest (i.e. most likely) transitions across mutation-free stationary points always involve the opening of "beach-heads" in unoccupied locations. Once some such location has become colonized by a self-sustaining population (i.e. one whose members do not prefer to return), the rest of the population may migrate to it in one of two different fashions. Either they

---

[10] For the present 2 × 2-context, this inequality is equivalent to saying that $\varphi(\theta) = \{s_1, s_2\}$. Analogously, the reciprocal inequality considered below implies that $\varphi(\theta) = \{s_1\}$.

[11] As can be seen from the argument used in the proof of the theorem, both of the summatories included in (6) and (7) can have their locations ℓ_i restricted to lie in $\mathcal{L}_0$. That is, every long-run state has all players located in one of the locations where there are inmobile players.

move by fast adjustment if it is profitable to migrate when the opportunity arrives or, alternatively, they do it by the slower process of accumulating one-step mutations. In any case, the event which triggers the process is always the establishment of such a beach-head.

When a given location has immobile players, it *always* remains a potential beach-head to be used by mobile players in implementing a transition towards it and its corresponding convention. If only by drift, such transition is bound to occur sooner or later. But, once it has occurred and other locations have consequently become depopulated, setting up new beach-heads in previously empty locations will turn out to be a much more difficult (i.e. unlikely) event: they will require some *simultaneity* of mutations, an infrequent but unavoidable event if a stable enclave is to arise. Thus, in the end, the considerations which underlie transitions within *permanently* established beach-heads (i.e. within those locations with immobile players) should prove crucial in shaping the long-run dynamics of the model. In pure coordination games, with sufficiently large friction in strategy adjustment, it is not difficult to see that the driving force operating within these locations must be the same as in standard evolutionary models (whose players are all "fixed"). Consequently, as in these models, risk dominance becomes the long-run selection criterion. On the other hand, in stag-hunt games, the establishment of a small enclave where the efficient strategy is played is sufficient to attract all mobile players, and hence our earlier analysis is essentially unaffected.

5 Summary

This paper has been motivated by the premise that endogenizing the structure of interactions is an important step forward in extending the evolutionary analysis of equilibrium selection in games. In comparison with the literature that contemplates a fixed interaction structure (e.g. Kandori, Mailath & Rob [9], Young [13], or Ellison [3]), our results emphasize the positive role of migration as a force for efficiency. In particular, efficiency is always ensured in stag-hunt games, irrespective of risk considerations and the presence or otherwise of some immobile players.

We are not alone in seeking to endogenize location choice, as the papers by Oechssler [11], Ely [5] and Mailath, Samuelson & Shaked [10] summarized in the Introduction indicate. However, our substantive conclusions are quite different from these papers, since our efficiency result is significantly qualified. Specifically, it does not apply to pure coordination games. In this case, not only has it been shown that co-existence of conventions is possible in the medium term but that, in the long term, efficiency is not ensured since either both equilibria are stochastically stable (if all players mobile) or the risk dominant equilibrium is selected (when a few players are immobile).

Clearly, the critical reason for this different conclusions derives from our assumption that the opportunities to change locations and adjust strategies are not fully correlated. Consequently, an agent always experiences some uncertainty on her ability to adapt optimally to some new environment she might consider migrating to. We believe that this feature of the model is realistic and captures an important feature of the phenomenon of migration in the real world. In our view, its inclusion leads to some interesting qualifications of the clear-cut conclusions of long-run efficiency that can be found in other related models of the existing literature.

Appendix

Proof of Proposition 1. If at a state ω a single strategy s_q is played by all players at the single location ℓ_1, then clearly for any player, the unique "better-response" is s_q, since (s_q, s_q) is a strict Nash equilibrium. Hence $\Gamma(\omega) = \omega$, i.e. no other state is accessible from ω,thus proving that $\{\omega\}$ is a limit set.

To prove the converse, we show that if m (≥ 2) distinct strategies are played at ℓ_1, then with positive probability we have a transition to a state where $m-1$ distinct strategies are played. By induction, this ensures that some state where a single strategy is played is accessible from any given state ω.

Let s_q and s_r be distinct strategies played at any state ω, and denote by $v(s_q, s_r, \omega_1)$ the expected payoff from strategy s_q to a player who is currently playing s_r at location ℓ_1. (Note that since we have a single finite population, the payoff to a player from changing his strategy when the system is at any given state also depends upon the player's former strategy.). We may write:

$$v(s_q, s_r, \omega_1) = \sum_{u=1}^{Q} \pi(s_q, s_u)\, \omega_{1u} - \pi(s_q, s_r)$$

and therefore:

$$\begin{aligned} v(s_q, s_r, \omega_1) - v(s_q, s_q, \omega_1) &= \pi(s_q, s_q) - \pi(s_q, s_r) \\ v(s_r, s_q, \omega_1) - v(s_r, s_r, \omega_1) &= \pi(s_r, s_r) - \pi(s_r, s_q) \end{aligned}$$

Adding the latter two expressions, one obtains:

$$\begin{aligned} & [v(s_q, s_r, \omega_1) - v(s_q, s_q, \omega_1)] + [v(s_r, s_q, \omega_1) - v(s_r, s_r, \omega_1] \\ = \; & [\pi(s_q, s_q) - \pi(s_r, s_q)] + [\pi(s_r, s_r) - \pi(s_q, s_r)] > 0 \end{aligned}$$

Hence either the s_q-player would like to adjust to s_r or the s_r-player would like to adjust to s_q. Assume w.l.o.g. that the former applies. With positive probability all s_q-players get the chance to adjust strategy, and choose s_r, while all players who are not playing s_q do not get the chance to adjust. Hence from a state where m distinct strategies are played, we have a positive-probability transition to a state where $m-1$ distinct strategies are played. By induction, we therefore have a transition with positive probability to a state where a single strategy is played. Since this state is absorbing, no further transitions are possible. ■

Proof of Theorem 1. Let $\omega \in \Omega$ satisfy conditions (a)-(d) of the theorem. Suppose that a player who is currently playing strategy s_q at location ℓ_i is given the opportunity to move. By (a) and (c), the player's current payoff is $\pi(s_q, s_q)$. If the player moves to a non-empty location ℓ_j, then (by (b)) a distinct strategy s_r must be played at this location. Since (d) implies that

$$\pi(s_q, s_q) > \theta\pi(s_r, s_r) + (1-\theta)\pi(s_q, s_r)$$

the player's payoff from migration is strictly less than the expected payoff from remaining at ℓ_i. Given that a player does not migrate, he will clearly not change his strategy either. Hence $\Gamma(\omega) = \{\omega\}$.

To prove the converse, let ω be any given state. With positive probability, no player gets the opportunity to migrate for Q periods but gets instead the opportunity to adjust his strategy. Hence, by Proposition 1, we have a positive-probability transition to a state where the same strategy is played at any location with two or more players. On the other hand, with positive probability, all loners (if any exist) get to migrate to some location and also adjust to the prevailing strategy. Combining these considerations, we may assert that there exist positive probability that if state ω did not satisfy (a) or (c), a (multi-step) transition takes place to a state ω' which does satisfy both of them.

Suppose now that the resulting state ω' displays two distinct locations, ℓ_i and ℓ_j, where strategies s_q and s_r are being played, with s_q being possibly identical to s_r. If either (b) or (d) is not satisfied at ω', then for the players at one of these locations (say, at ℓ_j), it is weakly better to migrate to ℓ_i and adjust strategy to s_q (if this is different from s_r). This occurs with positive probability until we are at a state satisfying all of conditions (a)-(d), which is absorbing. This shows that no state violating one of the conditions contemplated in the theorem can be part of a limit set of the unperturbed dynamics, thus completing the proof of the theorem. ■

Proof of Theorem 2. Since a stochastically stable state must obviously belong to a limit set of the unperturbed dynamics, Theorem 1 already narrows down substantially the states that need to be considered. Specifically, we just need to consider states ω which satisfy (a)-(d) in Theorem

1. For future reference, denote the set of those states by $U(\theta)$, letting $V(\theta)$ ($\subseteq U(\theta)$) stand for the set of states satisfy (i) and (ii) in the statement of the theorem.

As customary in recent evolutionary literature, our proof shall rely on the graph-theoretic framework developed by Freidlin & Wentzel [7]. Particularized to our context, it may be briefly summarized as follows.

For each $\omega \in \Omega$, define a ω-*tree* H as a collection of ordered pairs ("arrows") (ω', ω'') such that:

- Every $\omega' \in \Omega \setminus \{\omega\}$ is the first element of exactly one pair;
- From every $\omega' \in \Omega \setminus \{\omega\}$ there exists a *path* $\{(\omega^0, \omega^1), (\omega^1, \omega^2), ..., (\omega^{r-1}, \omega^r)\}$ such that $\omega^0 = \omega'$ and $\omega^r = \omega$. The set of all such ω-trees is denoted by $\mathcal{H}_\omega$.

Denote by T_ϵ the transition matrix of the perturbed evolutionary dynamics when the mutation probability is ϵ. Define, for each $\omega \in \Omega$,

$$q(\omega) \equiv \sum_{H \in \mathcal{H}_\omega} \prod_{(\omega', \omega'') \in H} T_\epsilon(\omega', \omega''). \tag{8}$$

Then, as established by Freidlin & Wentzel [7, Lemma 3.1, p. 177], we have:

$$\mu_\epsilon(\omega) = \frac{q(\omega)}{\sum_{\omega' \in \Omega} q(\omega')} \ . \tag{9}$$

Each $q(\omega)$ is a polynomial in ϵ. Thus, the *limit invariant distribution* defined by (5) is well-defined and, therefore, unique. To compute each $q(\omega)$, it is useful to introduce a cost function on transitions:

$$c : \Omega \times \Omega \to \mathcal{N} \cup \{0\} \ ,$$

which for each pair (ω, ω') specifies the minimum number of mutations $c(\omega, \omega')$ needed for the transition to occur with positive probability *via* mutation-free dynamics. That is, if $d(\omega, \omega')$ denotes the number of individuals whose location or action differ between ω and ω', then

$$c(\omega, \omega') \equiv \min_{\omega'' \in \Omega} \{d(\omega, \omega'') : T(\omega'', \omega') > 0\}.$$

The function $c(\cdot)$ is extended to every path h and every tree H, by simply adding the cost of all their constituent links. It is easy to see that the order of each $q(\epsilon)$, as a polynomial in ϵ, is simply given by $\min_{H \in \mathcal{H}_\omega} c(H)$. Thus, from (9), it follows that the set of stochastically stable states are precisely those whose minimum cost trees are themselves minimum across all possible states in Ω.

From the previous considerations, the proof of the theorem simply requires verifying the following lemma.

Lemma 1 *For all* $\hat{\omega} \in U(\theta)\backslash V(\theta)$ *and every* $\hat{\omega}$*-tree* $\hat{H} \in \mathcal{H}_{\hat{\omega}}$*, there exists some* $\tilde{\omega} \in V(\theta)$ *and an* $\tilde{\omega}$*-tree* $\tilde{H}$ *such that* $c(\tilde{H}) < c(\hat{H})$.

Proof of the Lemma. Given any state $\hat{\omega} \in U(\theta)\backslash V(\theta)$, consider the following two possibilities:
(a) $S(\hat{\omega}) = \{s_q \in S : \omega_{iq} > 0\} = \{s_{\hat{q}}\}$, i.e. is a singleton;
(b) $|S(\hat{\omega})| \geq 2$.

For any of these two cases, we first construct an auxiliary path $\tilde{h}$ from $\hat{\omega}$ to some $\tilde{\omega} \in V(\theta)$. This path connects $\hat{\omega}$ and $\tilde{\omega}$ with arrows (ω', ω'') which satisfy:

$$\begin{array}{rcll} c(\omega', \omega'') & = & 0 & \text{if } \omega' \notin U(\theta) \text{ (i.e. } \omega' \text{ is not a limit state)} \\ c(\omega', \omega'') & = & 1 & \text{otherwise.} \end{array} \quad (10)$$

In case (a), consider any state ω^o which is derived from $\hat{\omega}$ by one individual mutating to some new location ℓ^o and adopting any strategy $s_{q^o} \in \varphi(\theta)$. Since $s_{\hat{q}} \notin \varphi(\theta)$, it follows that a transition from ω_0 to a state $\tilde{\omega}$ where every individual plays strategy s_{q^o} in location ℓ^o is costless. Therefore, the path $\{(\hat{\omega}, \omega^o), (\omega^o, \tilde{\omega})\}$ satisfies condition (10), as desired.

In case (b), let $\check{\ell}$ denote any of the locations occupied in state $\hat{\omega}$ with its (at least two) individuals adopting some given strategy $s_{\check{q}}$. Consider a chain of transitions in which, in every one of them, *one* (and only one) individual in locations other than $\check{\ell}$ mutates to this location, adopting action $s_{\check{q}}$. All of these transition involve a unit cost. Eventually, all individuals are at location $\check{\ell}$, adopting action $s_{\check{q}}$. Let $\check{\omega}$ be the resulting state. If $s_{\check{q}} \in \varphi(\theta)$, the desired path has been constructed since then $\check{\omega} \in V(\theta)$. Otherwise, simply proceed as described for case (a), connecting $\check{\omega}$ to some other state $\tilde{\omega} \in V(\theta)$ through a path which verifies the contemplated requirements.

To complete the proof of the lemma, consider now any $\hat{\omega}$-tree $\hat{H} \in \mathcal{H}_{\hat{\omega}}$. By "tree-pruning" operations we want to transform $\hat{H}$ into an $\tilde{\omega}$-tree $\tilde{H}$ of lower cost, where $\tilde{\omega}$ is a state in $V(\theta)$ to which $\hat{\omega}$ may be connected through a path $\tilde{h}$ satisfying (10). This can be done through the following steps:

(1) Eliminate the arrow $(\tilde{\omega}, \omega')$ in $\hat{H}$ which starts at $\tilde{\omega}$.

(2) For all states ω which can be costlessly connected to $\tilde{\omega}$ (i.e. belong to its "basin of attraction") replace the arrows which started from these states in $\hat{H}$ by direct arrows to $\tilde{\omega}$, i.e. with corresponding $(\omega, \tilde{\omega})$.

(3) For all states in the path $\tilde{h}$ (including $\hat{\omega}$) simply replace the original arrows starting at them (or add, in the case of $\hat{\omega}$) with their respective arrow included in this path.

It is immediate to verify that, once steps (1)-(3) have been conducted, the original $\hat{\omega}$-tree $\hat{H}$ has been transformed into a well-defined $\tilde{\omega}$-tree. Moreover, by (1) and (2), at least a cost of 2 has been saved since, in order to escape the basin of attraction of $\tilde{\omega}$, at least two mutations are needed. On the other hand, by adding the arrow from $\hat{\omega}$ included in path $\tilde{h}$, the cost has increased by only one unit. Combining these considerations, the proof of the lemma, and thus of the theorem, is complete. ■

Proof of Theorem 3. As in the proof of Theorem 2, we shall rely on the graph-theoretic techniques developed by Freidlin & Wentzel [7]. For expositional simplicity, we shall assume that $\mathcal{L}_0 = \{\ell_1, \ell_2\}$. The argument for any non-empty set of locations with fixed players proceeds in a fully analogous fashion.

We start by introducing some concepts and facts which shall be used throughout the proof. The first concept is analogous to Samuelson's [12] notion of *mutation-connected component of Nash equilibria.* For each $i, j = 1, 2$, define:

$$\mathcal{C}_{ij} \equiv \{\omega = (\omega^1, \omega^2) : (\omega^1_{1i} + \omega^2_{1i}) + (\omega^1_{2j} + \omega^2_{2j}) = n\},$$

i.e. the set of states where everyone is either in location 1 or 2, those in the first location playing strategy s_i and those in the second location playing strategy s_j. We shall call each $\mathcal{C}_{ij}$ the *ij- component.* As for Samuelson's original concept, it is easy to see that all (mutation-free stationary) states in a given component are connected *via* one-step mutations. (Note that this crucially depends on θ satisfying the inequality included in the statement of the theorem.) Moreover, it follows by an argument fully parallel to that used in the proof of Theorem 2 above that only states in some $\mathcal{C}_{ij}$ may be stochastically stable. Indeed, one can adapt a result in [12, Theorem 2, p. 48] to conclude that if any state in a given component is stochastically stable so is every other one in this component.

We are now interested in extending the previous concept as follows. Let $\mathcal{A}$ denote the set of limit (or minimally absorbing) states of the mutation-free process. Given any $i, j = 1, 2$, we define

$$\mathcal{D}_{ij} \equiv \{\omega = (\omega^1, \omega^2) \in \mathcal{A} : \omega^1_{1i} > 0;\ \omega^1_{2j} > 0\}\ ,$$

as the *extended ij-component*, i.e. the set of states which have all individuals which are in location 1 play strategy i and all those which are in location 2 play strategy j. (Therefore, the configurations prevailing in other locations are unrestricted, provided, of course, that all individuals in each location play the same strategy.)

Suppose w.l.o.g. that strategy s_2 is the risk dominant strategy. First, we address Part (a) of the theorem. To prove the desired conclusion, it is enough to show that for every $\omega \notin \mathcal{C}_{22}$, and any ω-tree $H \in \mathcal{H}_\omega$ which is of minimum cost, there is some $\omega' \in \Omega$ for which an ω'-tree H' exists whose $c(H') < c(H)$. In fact, by the mentioned fact that all states in a given component are *jointly* included in (or excluded from) the set of stochastically stable states, it is sufficient to prove the previous statement for *some* (i.e. not every) state ω in each component $\mathcal{C}_{ij}$, with i or/and j equal to 1. For concreteness, we shall make $i = 1$ and $j = 2$ (the argument is fully analogous for the other two cases) and choose the state $\omega \in \mathcal{C}_{12}$ with $\omega^2_{11} = 0$, (i.e. the state in $\mathcal{C}_{12}$ where *all* mobile players are in location 2 playing strategy s_2).

We shall find it useful to restrict our attention to ω-trees H defined only in states belonging to $\mathcal{A}$. Note that, as argued by Kandori & Rob [8], if one focuses only on those periods where the process visits states in $\mathcal{A}$, the corresponding "imbedded process" can be described as a Markov chain on $\mathcal{A}$. And clearly, for the sake of identifying the stochastically stable states, it is sufficient to apply the graph-theoretic techniques of Freidlin & Wentzel [7] to such imbedded process.

Let $\hat{\omega} \in \mathcal{C}_{22}$ be the state uniquely characterized by the equality $\hat{\omega}^2_{12} = 0$ (thus, in analogy with ω above, $\hat{\omega}$ is the state in $\mathcal{C}_{11}$ where all mobile players are in location 2). In the ω-tree H under consideration (restricted to $\mathcal{A}$), there must exist a path, say h, linking $\hat{\omega}$ to ω. Consider now the following transformations on this ω-tree:

(1) Choose along path h the first state whose successor is *not* an element of the extended component $\mathcal{D}_{22}$ and identify this state with ω' above. Eliminate the arrow starting at ω' in H.

(2) Add the arrow $(\omega, \hat{\omega})$.

The previous set of operations transforms the original ω-tree H into a well-defined ω'-tree. On the cost side, step (1) decreases cost by no less than $\frac{m+1}{2}$ (i.e. $\frac{m-1}{2} + 1$), whereas step (2) increases cost by exactly $\frac{m-1}{2}$. Thus, overall, the cost of the resulting tree must be strictly less than that of the original ω-tree H. This completes the proof of Part (a) of the theorem.

Finally, to prove Part (b), it simply needs to be realized that if the game is of the stag hunt type (or θ is large enough), every limit state where the efficient strategy is played in *some* location must have *every* mobile player choosing this strategy. Thus, a transition to a state satisfying (7) can be materialized after just $\frac{m+1}{2}$ mutations from a state $\omega \in \mathcal{C}_{11}$ where, say, every mobile player is concentrated in *one* of the locations. On the other hand, a transition away from the collection of states satisfying (7) must necessarily involve a number of mutations which grows unboundedly with n, the size of the total population. Therefore, if n is large enough relative to m, it should be clear how we can rely on the techniques repeatedly used in this paper to prove rigorously the desired conclusion. To avoid unnecessary redundancy, the formal details are left to the reader. ■

References

[1] Aumann, R. (1993): "Nash equilibria are not self-enforcing", in J. Gabszewicz, J.-F. Thisse and L. Wolsey (eds.), *Economic Decision Making: Games, Econometrics and Optimization,* Elsevier B.V.

[2] Binmore, K., L. Samuelson & R. Vaughan (1995): "Musical chairs: modelling noisy evolution", *Games and Economic Behavior* **11**, 1-35.

[3] Ellison, G. (1993): "Learning, local interaction, and coordination", *Econometrica* **4**, 1047–73.

[4] Ellison, G. (1995): "Basins of attraction and long run equilibria", mimeo, Massachusetts Institute of Technology.

[5] Ely, J. (1995): "Local conventions", mimeo, University of California, Berkeley.

[6] Foster, D. & P. Young (1990): "Stochastic evolutionary game dynamics," *Theoretical Population Biology* **38**, 219-32.

[7] Freidlin, M.I. & A.D. Wentzel (1984). *Random Perturbations of Dynamical Systems*, New York: Springer Verlag.

[8] Kandori, M. & R. Rob (1995): "Evolution of equilibria in the long run: a general theory and applications", *Journal of Economic Theory* **65**, 383-414.

[9] Kandori, M., G. Mailath & R. Rob (1993): "Learning, mutations, and long-run equilibria in games", *Econometrica* **61**, 29-56.

[10] Mailath, G., L. Samuelson & A. Shaked (1995): "Evolution and endogenous interactions", Social Systems Research Institute, working

paper no. 9426, University of Wisconsin-Madison.

[11] Oechssler, J. (1993): "Competition among conventions", mimeo, Columbia University.

[12] Samuelson, L. (1994): "Stochastic Stability in Games with Alternative Best Replies", *Journal of Economic Theory* **64**, 35-65.

[13] Young, P. (1993): "The evolution of conventions", *Econometrica* **61**, 57-84.