ARTICLE IN PRESS

Journal of Economic Behavior & Organization
xxx (2004) xxx–xxx

JOURNAL OF Economic Behavior & Organization

www.elsevier.com/locate/econbase

Migration and the evolution of conventions

V. Bhaskar [a,*], Fernando Vega-Redondo [a,b,c]

[a] *Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK*
[b] *Departamento de Fundamentos del Análisis Económico and Instituto Valenciano de Investigaciones Económicas, Universidad de Alicante, 03071 Alicante, Spain*
[c] *Departament de Economía i Empresa, Universitat Pompeu Fabra, 08005 Barcelona, Spain*

Received 30 May 2001; accepted 12 March 2002

Abstract

We study an evolutionary model where agents are locally matched to play a symmetric coordination game. Opportunities to adjust strategy and location arrive asynchronously and infrequently, and cannot be coordinated. Our results on the short-run co-existence of different conventions and long-run efficiency depend upon a condition on off-equilibrium payoffs introduced by Aumann. In a *pure coordination game*, efficient and inefficient conventions may co-exist at different locations in the short-run, and inefficient conventions are stochastically stable. In a *stag-hunt game*, there is neither short-run coexistence nor long-run inefficiency.

JEL classification: C72; C73

Keywords: Evolution; Equilibrium selection; Migration

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 suggesting that differences in economic performance between societies may be related to different conventions (i.e. the playing of different equilibria). This raises an important question: under what conditions can one expect the efficient convention to prevail? If efficiency is not ensured, is it possible for different conventions to co-exist within the same society, thereby providing an explanation for differential economic performance?

* Corresponding author. Tel.: +44-1206-872-744; fax: +44-1206-872-724.
E-mail addresses: vbhas@essex.ac.uk (V. Bhaskar), vega@merlin.fae.ua.es (F. Vega-Redondo).

0167-2681/$ – see front matter
doi:10.1016/j.jebo.2002.03.001

The recent evolutionary literature on equilibrium selection in games has shed important light on these issues. Much of this literature (e.g. Foster and Young, 1990; Kandori et al., 1993, or Young, 1993) has focused on models containing agents from a single homogeneous society who interact over time. This literature shows that, in the medium-term, the best response dynamic ensures that a single convention prevails, making co-existence not possible. In the long-run, if agents occasionally make mistakes, the *stochastically stable* convention often coincides with the risk-dominant equilibrium in 2×2 games, so that efficiency is not ensured. These conclusions continue to hold if interaction is local rather than global, as Ellison (1993) shows, in a context where 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.

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, 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. Four recent papers by Oechssler (1997), Ely (1995), Mailath et al. (1995), Dieckmann (1998) 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). By 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 mutations with small probability. As in Oechssler's scenario, conventions cannot co-exist since agents 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 agent 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 in Ely's framework. The paper by Dieckmann shows that this result is robust to imperfect information in that, even if players only have imperfect knowledge on the strategy configuration at other locations, efficiency is ensured.[1]

[1] The paper by Mailath et al. contemplates a quite different context in continuous time; in it players from two continuum populations must decide which locations to visit. They establish that if the evolutionary system is

Although the present paper shares some similarities with the papers summarized above, mainly those by Oechssler, Ely and Dieckmann, it also displays some key differences. The interaction framework is the same as in these three papers, with players distributed among a certain number of locations ("cities"). By choosing to locate in one of them, an agent selects the corresponding matching pattern induced by its population profile. Players adjust both location and strategy as in those papers; however, in our context, these adjustment opportunities never arrive simultaneously. Hence a player who receives the opportunity to migrate will not be sure that she will be able to adjust her strategy in the optimal direction in the new environment. Similarly, a player who receives the opportunity to adjust strategy can never be sure that she will be able to migrate to the appropriate location. This uncertainty has the consequence that the model no longer produces, even for the simple 2×2 case, the unqualified efficiency conclusion of Oechssler, Ely, or Dieckmann. In particular, depending on payoffs, one may encounter convention *co*-existence in the medium-run (when mutations are absent) or inefficiency in the long-run (when mutations are allowed).

In our view, these contrasting conclusions illuminate an important theoretical fragility of the customary models that study the co-evolution of strategy and location choices in a population context. In all of these models, it is of crucial importance that (at least with some positive probability) agents should be able to *correlate perfectly* both dimensions of their decision.[2] If instead one postulates the existence of *some* friction that prevents such *synchronous*adjustment, player mobility is not sufficient to ensure long-run efficiency.

In our model, a condition on off-equilibrium payoffs plays a critical role in the analysis. This condition is orthogonal to considerations of risk dominance, and is of a *qualitative* (i.e. not cardinal) nature. It is related to a distinction enunciated by Aumann (1993) while discussing the effectiveness of pre-play communication between rational players. To illustrate this condition, consider a symmetric 2×2 coordination game with two strict equilibria (H, H) and (L, L) where the former is Pareto dominant. Such a game is said to be *pure coordination game* if an L player prefers that his opponent plays L rather than H; otherwise, it is called a *stag-hunt game*. Aumann argues that cheap talk leads to efficiency if the game is one of pure coordination but not if the game is of the stag-hunt kind, since communication is credible in the former but not in the latter context.

Our results, when applied to the 2×2 case, hinge upon the same distinction but are, in a sense, the converse of Aumann's. First, in the *medium-term* (i.e. when the dynamics is unperturbed),[3] we find that convention co-existence is possible if the game is one of pure coordination, but never possible in a stag-hunt game. Once we allow for mutations (i.e. in the long-run), *simultaneous* co-existence of conventions is no longer possible. However, we still have that *both* equilibria are stochastically stable in a pure coordination game, whereas the efficient equilibrium is the unique stochastically stable convention in the alternative case. Hence, again, we observe that, even in the presence of noise, the evolutionary process

monotonic and agents have some control over their interaction pattern (e.g. they can avoid any undesired matching), then every locally stable configuration must be efficient.

[2] Even if, as in Ely (cf. his Theorem 2), this probability is allowed to become small, it is supposed to remain bounded above zero in the mutation rate.

[3] See Binmore and Samuelson (1997) for a discussion of the different time horizons that are typically considered in evolutionary models.

ensures efficiency in the long-run only when Aumann finds pre-play communication to be ineffective (i.e. in the stag-hunt game).

Note that, if we were to restrict our study to 2×2 coordination games, any of these is either pure-coordination or stag-hunt. For the general $n \times n$ case, however, matters are more complicated and such a dichotomy is no longer valid. Consequently, our analysis will be extended to this case in a natural way, but still be found to reflect essentially the same considerations illustrated for 2×2 games.

The rest of the paper is organized as follows. Section 2 presents the general model. Section 3 presents and discusses the results. Within this latter section, Section 3.1 deals with the unperturbed (medium-term) dynamics, whereas Section 3.2 addresses the analysis of the long-run dynamics where mutations are introduced. Section 4 closes the paper with a summary. For the sake of smooth discussion, all formal proofs are relegated to Appendix A.

2. The model

Consider a population of n ($\geq$3) individuals who are matched in pairs to play a finite bilateral symmetric game with strategy set $S = \{s_1, s_2, \ldots, s_Q\}$ 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 $\mathbb{L} = \{\ell_1, \ell_2, \ldots, \ell_I\}$ stand for the set of possible locations $(I \geq 2)$. To describe the state of the system, it will be enough to specify how many agents adopt each of the Q strategies in every one of the I locations at the time of play. Thus, a typical state may be identified as an $I \times Q$ matrix ω whose entries $\omega_{iq} \in \mathbb{N} \cup \{0\}$ satisfy

$$\sum_{q=1,\ldots,Q_{i=1,\ldots,I}} \omega_{iq} = n.$$

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

Fix any location, and consider the matches taking place at this location. We shall assume first that players are randomly matched within the location. Second, we shall assume that there are "constant returns" in location size so that there is no advantage or disadvantage per se in being at a larger location as long as there are at least two individuals in the location. In consequence, the expected payoff of playing strategy s_q at location ℓ_i is given by the following expression:[4]

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

[4] One specification of the matching process which gives rise to this expected payoff function is as follows. If a location has an even number of players, then pairs of them are randomly chosen to be matched *without replacement*. If a location has an odd number of players, then pairs of players are analogously selected without replacement until only one player is left. With probability one-half this player is left unmatched, and with probability one-half this player is matched again with one of the players who was already matched (i.e. one player is matched twice with probability one-half).

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

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, assuming as well that playing the game always entails positive payoffs (i.e. $\pi(s_q, s_r) > 0$ for all $q, r = 1, 2, \ldots, Q$). This assumption ensures that, in any limit state, all players will be part of some "city". It could be generalized somewhat in what follows, although it is crucial for our analysis that a player should never prefer to move to an empty location than to of play some equilibrium of the game at a populated location.[5]

We now turn to the dynamics of the model. Time t is measured discretely and indexed by $t = 1, 2, \ldots$ At each date, players may adjust both their strategy and their location in the following two-stage fashion. At stage 1, a trichotomous random variable determines which one of three mutually exclusive events occurs, each occuring with strictly positive probability:

(a) the player receives the opportunity to adjust her strategy only;
(b) the player receives an opportunity to migrate only; and
(c) the player cannot revise either strategy or location.

Given the option to adjust, the player must exercise her option immediately or not exercise it at all. Once the player makes this decision in stage 1, she proceeds to stage 2, where once more a trichotomous random variable determines which one of these three above events happens again. These adjustment opportunities are assumed to arrive independently across agents but are not necessarily independent for each given player. In particular, we shall find it convenient to simplify matters and assume that there is zero probability of receiving an adjustment opportunity in the same dimension in both stages.[6] For future reference, the probability of a strategy adjustment opportunity in stage 2 conditional upon the player receiving a migration opportunity in stage 1 is denoted by $\theta > 0$. Similarly, the probability of a migration opportunity in stage 2 conditional upon a strategy adjustment opportunity in stage 1 is denoted by $\lambda > 0$. Our primary focus will be on the case where both θ and λ are small so that, after having received a revision in the first stage, it is relatively improbable that a further revision opportunity (in the other dimension) arrives in the second stage. It is because of this uncertainty across stages (that may render it risky for agents to aim at a revision on the two dimensions of choice) that we speak of adjustment "friction" in the model.

After all migration/strategy adjustment decisions are made in stage 2, the players in any location are matched to play the game and then receive the payoffs formerly specified. As explained, the novel and key feature of our model is that, due to the contemplated asynchronicity, agents will not always be in a position to ascertain the ex post merits of either a strategy or location decision in stage 1. That is, the realization of its potential

[5] If this condition were violated for one of the equilibria of the coordination game, then it is clear that no limit state can have players choose the corresponding action. In that case, our results would have to be adapted in the obvious fashion.

[6] If a revision on the same dimension could arise in the two stages, the expression for the expected payoffs obtained at the first stage after receiving a revision opportunity (cf. for example (3) and (4)) would have to be modified by including an additional contingency. Nothing essential in our conclusions, however, is affected by this variation if the probability of a second revision opportunity is low.

benefits may require the subsequent (but uncertain) adjustment on the other dimension in stage 2. How are we then to formulate their decision problem under these circumstances? It may be useful, at the outset, to summarize the essentials of our approach.

First, let us assume that time preferences are such that, when adjusting at any date t, players are solely concerned with their payoff at (the end of) such a date. Such myopia is justified in contexts where agents are impatient and adjustment opportunities arrive relatively infrequently, so that the stream of payoffs accruing before the next possible adjustment period is typically long. Second, suppose that agents hold static expectations on what the other agents will do at the next time of play. That is, they hold the (point) beliefs that all other players will remain with the same strategy and location choices as observed previously. This assumption is customary in evolutionary game theory and, in the context of our model, is appropriate if adjustment probabilities are low, for in this case, an agent who is currently enjoying a revision opportunity should attribute a low probability to other agents being in an analogous position at the same time. Finally, let us postulate a "better-response" adjustment dynamics by assuming that agents choose, with positive probability, any location or strategy which is weakly better (in expected payoff terms) than their current choice.[7]

Consider any given t and let $\omega(t-1)$ be the state that prevailed in the preceding period. Adjustment leading to $\omega(t)$ may occur in this model in several ways. On the one hand, an agent may migrate in stage 1 (given the opportunity) with the hope that she will be able to adjust her strategy at stage 2. Alternatively, she may change strategy in stage 1, in the hope that she will be able to migrate in the second stage. Finally, one may have adjustment decisions in either stage in one dimension alone, even if the agent does not contemplate adjustments in the second dimension.

We start our analysis by considering the possibility of migration followed by strategy revision. Consider the second stage of the adjustment process, and suppose that there is an agent at location ℓ_i who is playing strategy s_q and who receives the option of adjusting her strategy. Denote by $\tilde{\omega}_i$ her expectations of the strategy configuration to prevail in ℓ_i at the end of t (i.e. the next point of play). Since expectations are assumed to be formed in a static fashion, we write

$$\tilde{\omega}_i = \omega_i(t-1) + \delta\, e_q, \tag{2}$$

where $\delta = 1$ if the player concerned has migrated in stage 1 of this period, $\delta = 0$ otherwise, and e_q stands for the Q-dimensional vector $(0, \ldots, 0, 1, 0, \ldots, 0)$ whose qth component is 1. That is, we posit that players rely on the information obtained from the preceding period alone in order to form their current (static) expectations. In particular, therefore, it is ruled out that any interim adjustment produced by some players in the first stage of the current period might impinge on the second-stage expectations held by other players. However, this is just a convenient (notation-saving) simplification that could be easily dispensed with.[8] Given (2), consider the set $\{s_r : \hat{\pi}(s_r, \tilde{\omega}_i + e_r - e_q) \geq \hat{\pi}(s_{q,} \tilde{\omega}_i\}$, where recall that $\hat{\pi}(s_q, \omega_i)$

[7] Alternatively, one may assume that an agent will change her strategy or location if this is strictly better than the status quo. We discuss the implications of this alternative assumption in Remark 3 following our main results.

[8] Specifically, note that adjustment paths where only a single agent receives a revision opportunity at each period have positive probability. Along such paths (which are always sufficient to implement any desired transition of the unperturbed dynamics) the issue of first-stage adjustments possibly influencing second-stage expectations does not arise.

stands for the expected payoff of strategy s_q facing configuration ω_i. We shall postulate that every strategy in this set (consisting of those strategies whose payoff is expected to be no smaller than that of s_q) is chosen by the player in question with positive probability.

We may now address the (first-stage) migration decision for the agent who is currently playing action s_q at location ℓ_i and is contemplating migrating to some other ℓ_j. It is assumed that the agent will choose location ℓ_j with positive probability if, given static expectations, her expected payoff from migration is at least as large as the expected payoff associated with 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. In other words, the agent decides, in stage 1 itself, both her current adjustment and the adjustment she would make in the future.[9] For the analysis, as is common in stochastic evolutionary models, it suffices to identify all migration decisions that have positive probability.

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

$$(1-\theta)\hat{\pi}(s_q, \omega_i(t-1)) + \theta\hat{\pi}(s_{q'}, \omega_i(t-1) + e_{q'} - e_q). \tag{3}$$

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 $\hat{\pi}(s_q, \omega_i(t-1))$, 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

$$(1-\theta)\hat{\pi}(s_q, \omega_j(t-1) + e_q) + \theta\hat{\pi}(s_r, \omega_j(t-1) + e_r). \tag{4}$$

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 (this probability will in general be less than one since other locations may also be better than ℓ_i) if, and only if, the maximal expected payoff from migration to ℓ_j (achieved when s_r in (4) coincides with a *best* response) is greater or equal than the minimum payoff from not moving (achieved by making $s_{q'}$ equal to s_q in (3).

To formalize matters, given s_q and ℓ_i, define

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

as the expected maximum payoff perceived at t by an agent who was formerly choosing s_q at location ℓ_i and is considering moving to location ℓ_j ($j \neq i$). Combining all of the above

[9] The alternative assumption (that we believe is less persuasive) is that the agent does not know which strategy she will adopt in stage 2 if given the option to adjust. In this case the agent faces an additional uncertainty arising from not knowing her own future actions. Under this alternative assumption, it may be verified that migration and strategy adjustment possiblities are reduced, leading to subsets of $(\mathcal{L}(s_q, \ell_i, \omega(t-1))$ and $\mathcal{S}(s_q, \ell_i, \omega(t-1))$, defined in (5) and (6). Hence the efficiency results will be even more restricted as compared to the present paper.

considerations, the set of locations ℓ_j to which the agent in question (originally in location ℓ_i) can migrate in stage 1 may be identified as follows:

$$\mathcal{L}(s_q, \ell_i, \omega(t-1)) = \{\ell_j \in \mathbb{L} : \pi^*_\theta(\ell_j; s_q, \ell_i, \omega(t-1)) \geq \hat{\pi}(s_q, \omega_i(t-1)), \ell_j \neq \ell_i\}. \quad (5)$$

Consider now the second possible case where the agent receives an opportunity to revise her strategy in stage 1. The analysis of this case parallels our earlier discussion. If the agent is currently playing strategy s_q at location ℓ_i, denote by $\pi^*_\lambda(s_r; s_q, \ell_i, \omega(t-1))$ her maximum expected payoff from choosing strategy $s_r \neq s_q$, given that she may later enjoy a migration opportunity with probability λ. On the other hand, as before, her minimum expected payoff from not changing her strategy is given by $\hat{\pi}(s_q, \omega_i(t-1))$, which corresponds to choosing the weakly "better response" of staying in location ℓ_i if given the option to migrate at stage 2. Hence the set of strategies that the player currently at location ℓ_i could switch to is given by

$$\mathcal{S}(s_q, \ell_i, \omega(t-1)) = \{s_r \in S : \pi^*_\lambda(s_r; s_q, \ell_i, \omega(t-1)) \geq \hat{\pi}(s_q, \omega_i(t-1)), s_r \neq s_q\}. \quad (6)$$

Finally, it is easy to verify that adjustment in one stage alone does not lead to any further possibilities. More precisely, the set of locations to which migration may occur in such a way is a subset of $\mathcal{L}(s_q, \ell_i, \omega(t-1))$, whereas the set of strategy adjustment that may occur is a subset of $\mathcal{S}(s_q, \ell_i, \omega(t-1))$. Hence, (5) and (6) identify *all* those transitions that display positive probability of occurrence, either via initial migration of some agent (possibly with the hope of adjusting her strategy later) or via an initial adjustment of strategy (possibly with the hope of changing her location before the next point of play). No more precise specification of matters will be required for our purposes.

Any process of adjustment with the features 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 shall analyze a perturbation of this Markov chain in which individuals occasionally experiment (or mutate).

3. Analysis

We shall assume that the underlying game G is a *strict coordination game* (i.e. a game 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. 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 divided into two parts. Firstly, we characterize the limit (or minimally absorbing) sets of the unperturbed dynamics. Secondly, we address the issue of stochastic stability, identifying those limit sets that 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) adjustment 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 $u \in \mathbb{N}$).[10] Extend now these notions to sets, so that $\Gamma(A)$ is the set of states accessible from some $\omega \in A$. Then, set A is called *absorbing* if $\Gamma(A) \subset A$. Furthermore, 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 G be a strict coordination game played at a single location. If A is a limit set of the unperturbed dynamics, then $A = \{\omega\}$ and $\omega_{1q} = n$ for some $q \in \{1, 2, \ldots, Q\}$, so that all players play the same strategy.*

Proof. See Appendix A. □

Proposition 1 shows that in the absence of migration, the process of strategy adjustment converges to a convention with all players using the same strategy. This proposition may be of some independent interest, since the result holds for a wider class of games than proved by Kandori and Rob (1995) 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.[11]

The above proposition implies that play at any location converges to a convention, so that at most one strategy is played per location. In addition, it will be shown below that drift (i.e. weakly optimal adjustments) ensures that each strategy is played in at most one location (cf. the proof of Theorem 1). The essential question then is whether conventions can *co*-exist, so that different strategies are being played at different locations. The answer to this question depends upon the adjustment probabilities, θ and λ. Our main focus in the present paper is the case that these probabilities are both small (i.e. when frictions are large). We assume therefore that $\theta = \hat{\theta}\Delta$ and $\lambda = \hat{\lambda}\Delta$, where $\hat{\theta}$, $\hat{\lambda}$ are both fixed positive numbers, and Δ is also positive but small.

Suppose that strategies s and s' are being played at two different locations. Under what conditions will s be vulnerable to disruption, due to the existence of a location playing s' to which players may migrate? Clearly s will only be vulnerable if it is a payoff dominated by s'. In this event, there are only two ways that s can be destabilized. First, if a migration opportunity arises first, s-players may migrate to the location playing s' in the hope that

[10] The notation used here is standard, $T(\omega, \omega')$ standing 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.

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

they will be able to adjust strategy subsequently. Second, if a strategy adjustment opportunity arrives first, s-players may switch to s' in the hope that a migration opportunity will arise. We shall see that the second channel of destabilization—strategy adjustment first, migration later—will always be ineffectual if Δ is sufficiently small. However, the first channel—migration first, strategy adjustment later—will still be effective in some classes of games even if Δ is small.

We now proceed to *characterize* the set of strategies which may be played in a limit set of unperturbed dynamics. We define a binary relationship $R(\hat{\theta}\Delta)$ on S as follows. For any pair of strategies $s, s' \in S$,

$$s\, R(\hat{\theta}\Delta)\, s' \Leftrightarrow \pi(s,s) > \hat{\theta}\Delta\, \pi(s',s') + (1-\hat{\theta}\Delta)\, \pi(s,s'). \tag{7}$$

Intuitively, if $sR(\hat{\theta}\Delta)s'$, then a location playing s is not vulnerable to another location playing s' through the process of first migrating and then (possibly) adjusting the strategy. Observe that a more efficient strategy can never be destabilized: if $\pi(s,s) \geq \pi(s',s')$ then $s\, R(\hat{\theta}\Delta)\, s'$, since all elements of S are strict Nash equilibria. However, this is not necessarily the case for an inefficient strategy, as we shall see shortly.

Analogously, define the binary relation $R(\hat{\lambda}\Delta)$ on S as follows:

$$s\, R(\hat{\lambda}\Delta)\, s' \Leftrightarrow \pi(s,s) > \hat{\lambda}\Delta\, \pi(s',s') + (1-\hat{\lambda}\Delta)\, \pi(s',s)\,. \tag{8}$$

Intuitively, if $s\, R(\hat{\lambda}\Delta)\, s'$, then the strategy s (at one location) is not vulnerable to strategy s' being played at another location through the process of first adjusting the strategy in the hope of later being able to migrate. Again, a more efficient strategy is not vulnerable to a less efficient strategy, although an inefficient strategy may well be.

Finally, we contemplate a binary relation $R^*(\Delta)$ on S, which captures the invulnerability of any given strategy (i.e. convention) through *any* of the two processes. It is defined as follows:

$$s\, R^*(\Delta)\, s' \Leftrightarrow sR(\hat{\lambda}\Delta)s' \wedge sR(\hat{\theta}\Delta)s'. \tag{9}$$

The relation R^* can be used to characterize those sets of conventions that may co-exist in limit states of unperturbed dynamics. Given $\Delta > 0$, this collection of subsets of S is defined as follows:

$$\Phi(\Delta) \equiv \{D \subseteq S : \forall s, s' \in D, s \neq s', s\, R^*(\Delta)s' \wedge s'\, R^*(\Delta)s\}. \tag{10}$$

Of course, $\Phi(\Delta)$ includes all the singleton subsets of S. The interesting fact is that, in general, it will also contain subsets with *more than one* strategy (this corresponds precisely to situations of convention co-existence).

In general, a subset of S with more than one element belongs to $\Phi(\Delta)$ if and only if all of the strategies included in it are maximal with respect to $R^*(\Delta)$ when restricted to this set. For future use, we are also interested in identifying those strategies which are $R^*(\Delta)$-maximal in *any* subset of S. This is simply given by the set of maximal elements of $R^*(\Delta)$ in S, defined as follows:

$$\varphi(\Delta) \equiv \{s \in S : s\, R^*(\Delta)\, s', \forall s' \in S, s' \neq s\}. \tag{11}$$

As formulated in (9), the binary relation $R^*(\Delta)$ is simply the *intersection* of the two relations $R(\hat{\theta}\Delta)$ and $R(\hat{\lambda}\Delta)$. However, it is important to understand that, for small Δ,

Table 1

	H	L
H	a, a	b, c
L	c, b	d, d

each of these constituent relations—$R(\hat{\theta}\Delta)$ and $R(\hat{\lambda}\Delta)$—displays very different properties. On the one hand, note that if Δ is small, one obtains that $s\, R(\hat{\lambda}\Delta)\, s'$ for all $s, s' \in S$ since

$$\pi(s, s) > \pi(s', s),$$

in view of the fact that every (s, s) is a strict Nash equilibrium. Thus, in this case, $R(\hat{\lambda}\Delta) = S \times S$ and, therefore, $R^*(\Delta)$ coincides with $R(\hat{\theta}\Delta)$. On the other hand, it may well happen that

$$\pi(s, s) \leq \pi(s, s')$$

in which case one will have that $R(\hat{\theta}\Delta)$ is a non-trivial (strict) subset of $S \times S$, even for arbitrarily small Δ.

To illustrate matters, consider a 2×2 coordination game with general payoff structure as indicated in Table 1. Without loss of generality; assume that $a > d$, thus H is the efficient strategy (since it is a coordination game, we must also have that $a > c$ and $d > b$).

One may distinguish the following two alternative scenarios. First, if $c \geq d$, we are in the presence of a so-called *stag-hunt game*. On the other hand, if $c < d$ we shall call it a *pure coordination game*. The essential difference between both cases does not dwell at all in the standard considerations of risk-dominance that 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 $d - c$.[12]

Consider now how the above defined constructs, $\Phi(\Delta)$ and $\phi(\Delta)$, apply to this 2×2 strategic context. Suppose that Δ is sufficiently small so that $s\, R(\hat{\lambda}\Delta)\, s'$ for any s and s'. Then, in a stag-hunt game, $\Phi(\Delta) = \{\{H\}, \{L\}\}$. In contrast, in a pure coordination game, there is a small enough value of Δ (specifically, it is enough that $\hat{\theta}\Delta < d - c/a - c$) for which $\Phi(\Delta) = \{\{H\}, \{L\}, \{H, L\}\}$, so that not only the two singletons but also the set that includes both strategies are all included in $\Phi(\Delta)$. As we shall presently see (cf. Theorem 1), this implies that two conventions may co-exist in the medium-run (i.e. as a limit state of the mutation-free dynamics) when the game is of pure coordination, but cannot co-exist if it is a stag-hunt game and the pace of play is fast enough.

On the other hand, concerning $\varphi(\Delta)$, we have that $\varphi(\Delta) = \{H\}$ for all values of Δ if the above 2×2 game is of the stag-hunt kind, whereas $\varphi(\Delta) = \{H, L\}$ if it is of pure coordination,

[12] As mentioned, the distinction between stag-hunt and pure-coordination games was first elucidated by Aumann in his discussion of pre-play communication. We discuss Aumann’s ideas and their relation to our results in the concluding section.

provided Δ is small (again, it is enough that $\hat{\theta}\Delta < d - c/a - c$). As we shall explain in Section 3.2 (cf. Theorem 2), this has similar implications in terms of "co-existence" of the two strategies (now, not simultaneously but over time) when we allow the process to be perturbed by occasional mutations.

The limit sets of the unperturbed dynamics are characterized by the following result.

Theorem 1. *Let G be a strict coordination game. 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 the following condition.*

(a) *For each location $\ell_i, i = 1, 2, \ldots, I$, there exists at most one strategy s_q such that $\omega_{iq} > 0$.*
(b) *For each strategy $s_q, q = 1, 2, \ldots, Q$, there is at most one location ℓ_i such that $\omega_{iq} > 0$.*
(c) *For all $i = 1, 2, \ldots, I, n_i = \sum_{q=1}^{Q} \omega_{iq} \geq 2$.*
(d) *The set of strategies played at ω, $\{s_q \in S : \omega_{iq} > 0, i = 1, \ldots, I\}$, is an element of $\Phi(\Delta)$.*

Proof. See Appendix A. □

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 Δ 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 Δ, 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 dis-coordination will deter migration (and thus the disruption of co-existence) if friction is sufficiently significant relative to the pace of play. These considerations would seem to suggest a possible explanation for different economic performance in neighboring regions (e.g. north versus south Italy) based on the persistence of different conventions being played in each of them despite the possibility of migration.

3.2. Stochastic stability

We now focus on a perturbation of the contemplated Markov chain where individuals are supposed to mutate at a very low rate. More precisely, it is postulated that, with some stationary and independent probability ϵ, every individual is subject to locational or strategic mutation (possibly both, with a probability ϵ^2).[13] In those events, she is simply assumed to change her prior locational or strategic decision in a way unrelated to any of the considerations described above. For concreteness, it will be assumed that every alternative option is chosen 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 gradual process of experimentation. Alternatively, it may be viewed as a test of robustness, singling out that kind of (medium-term) behavior which is resilient in the face of small noise introduced 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 reflect its dependenceexplicitly on the mutation probability. Since, as explained, we want to think of ϵ as small, we shall be interested in the limit invariant distribution resulting when $\epsilon \to 0$. Such *limit invariant distribution*

$$\mu^* \equiv \epsilon \overset{\lim}{\underset{0}{\to}} \mu_\epsilon \tag{12}$$

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

When extending 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 perturbs the dynamics of the process. This is precisely the idea captured by support of the limit invariant distribution of the process, as defined in (12). Note, of course, that in order for a state to be stochastically stable (i.e. lie in this support), it must necessarily belong to some limit set of the unperturbed dynamics.

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 (12).

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

Theorem 2. *Let G be a strict coordination game. There exists some $n_o \in \mathbb{N}$ such that, if $n \geq n_o$, a state $\omega \in \Omega$ is stochastically stable if and only if it satisfies the following condition.*

[13] It could also be postulated that when an individual is subject to mutation, both dimensions (strategy *and* location) are chosen afresh. This alternative specification would not alter our conclusions.

(i) $\exists i \in \{1, \ldots, I\}, \exists q \in \{1, \ldots, Q\}$ *s.t.* $\omega_{iq} = n$.
(ii) *The* (*unique*) *strategy* s_q *played at* ω *belongs to the set* $\varphi(\Delta)$.

Proof. See Appendix A. □

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 cardinality of the set $\varphi(\Delta)$, as defined in (11)) the process is consistent with what we could interpret as inter-temporal coexistence for when the set $\phi(\Delta)$ includes more than one strategy, it follows that there is a multiplicity of conventions in which the process spend a significant fraction of its time in the long-run.

By way of illustration, consider again the 2×2 strategic game described in Table 1. Recall that, in this context, we have that $\varphi(\Delta) = \{H\}$ for all values of Δ if the game is of the stag-hunt kind, whereas $\varphi(\Delta) = \{H, L\}$ in the case of pure coordination if Δ is small enough (specifically, if $\hat{\theta}\Delta < d - c/a - c$). Thus, we may conclude that whereas only the efficient convention H is stochastically stable if the game is of the stag-hunt kind, both monomorphic states (everyone playing H *or* everyone playing L) turn out to be stochastically stable in the pure coordination case in case of significant friction. As explained, the latter conclusion contrasts sharply with the received literature that claims that by allowing agents to adjust their locational choices as the process evolves, efficient configurations are always shown to arise as the unique long-run outcomes.

To gain some basic insight for this result, recall our former heuristic explanation of the differences between pure-coordination and stag-hunt games concerning the possible co-existence of conventions in the medium-run (cf. the discussion following Theorem 1). As explained, whereas stag-hunt games induce strong migration forces towards locations where the efficient convention is being played, these forces do not exist for the alternative kind of pure-coordination games when friction is significant. As it turns out, analogous considerations 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, through mutation, this indeed happens, a wave of migration may be immediately triggered that (with some positive probability) will lead to the eventual collapse of the inefficient convention.

We end this section with three remarks.

Remark 1. It is important to understand the role played by the parameter representing the frequency of adjustment opportunities, Δ, in the analysis. The model's novel and most interesting conclusions arise only when Δ is relatively small (i.e. when the stochastic and independent arrival of revision opportunities may be taken to induce a significant amount of "friction"). If, on the contrary, Δ is quite large, the practical implications of friction are largely removed, and our framework essentially approximates that of former literature. In particular, the set $\varphi(\Delta)$ includes only the efficient strategy, which then becomes the unique long-run convention.

Remark 2. 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 for observing any long-run (i.e. stochastically stable) outcome is of the order ϵ^{-1}, where ϵ is the mutation probability.[14] Thus, in this sense, the rate of convergence towards long-run outcomes is fast and *independent* of either the population size or the payoff structure of the game. Clearly, this is an important consideration since the relevance of any long-run prediction should be assessed in terms of the entailed rate of convergence.

Remark 3. One may allow for alternative rules for player adjustment rather than the weak better response dynamic assumed here. For instance, consider a *strict better response* dynamics where a player only adjusts if the alternative is strictly better, in expected payoff terms, as compared to the status quo. In the unperturbed dynamics, this would allow the existence of stationary states with several populated cities, all of them playing the same convention. However, these states would be destabilized by a sequence of single mutations so that concentration in a single city would result from in any stochastically stable state, and Theorem 2 would be unaffected. As a second alternative, we could postulate instead a *best response* dynamics. This is also a common formulation in the evolutionary literature but, to the best of our knowledge, its convergence to Nash equilibria has only been established for 2×2 games. Building upon this convergence, our results on stochastic stability could be readily applied to coordination games with two strategies. In particular, our substantive conclusions on the difference between pure coordination games and stag-hunt games would seem to hold also with a best response dynamic.

4. Summary and discussion

The recent literature on endogenous interactions (Oechssler, Ely Dieckmann, and Mailath, Samuelson and Shaked) has emphasized the positive role of migration in ensuring efficiency. This contrasts with the results of models with a fixed interaction structure, where the risk-dominant equilibrium is selected. Migration plays a role in destabilizing inefficient equilibria similar to the role of pre-play communication in evolutionary models (e.g. Kim and Sobel, 1995).

To see this, consider a population that is sending a particular message m and playing an inefficient equilibrium. The actions that players intend to take after an unsent message (say m') is not subject to selection pressure and may hence be replaced by the efficient action through a process of drift. At this point, it is optimal for players to send m' and play the efficient action. In Ely's model, locations play a role similar to that of messages. An inefficient action may be destabilized by a player moving to a vacant location and switching

[14] An easy way to confirm this statement is to observe that, in the language of Ellison (2000), the co-radius of any of our stochastically stable equilibria is 1.

to the efficient action. The difference, as compared to the cheap talk literature, is that this cannot be achieved entirely through drift and requires some "mutation" instead.

Our paper differs from the other literature on endogenous interactions in that it emphasizes the role of the real frictions and costs involved in the processes of migration and strategy adjustment. In contrast with the above mentioned literature, our critical assumption is that adjustment opportunities for location and strategy change almost surely never arrive simultaneously. This implies that a potential migrant is uncertain whether she will adapt to the new environment and adapt optimally. Similarly, an agent who contemplates changing her strategy is also uncertain that she will be able to change location. Despite this uncertainty, our results show that migration retains a positive role in ensuring efficiency, particularly for stag-hunt games in the long-run, irrespectively of risk considerations. However, in pure coordination games, both efficient and inefficient equilibria remain stochastically stable, provided that the pace of play is fast enough, and thus adjustment friction may display a relevant role. Under these circumstances, we even find that *simultaneous* co-existence of conventions is possible in the medium-term, possibly providing a certain basis for the explanation of geographical variation in performance across different regions or societies.

As indicated, the distinction between stag-hunt and pure-coordination games that plays a crucial role in our analysis was first introduced by Aumann, who was concerned with the credibility of cheap talk when players are fully rational. Previous literature (e.g. Farrell, 1988) had typically assumed that the promise by a player to play the efficient equilibrium was always credible. In contrast, Aumann argued that this promise should be regarded as credible in a pure coordination game, but uninformative in a stag-hunt game.

Somewhat paradoxically, our evolutionary analysis hinges upon the same distinction as Aumann's but its substantive conclusions are precisely the opposite. For, in our context, long-run efficiency is achieved in a stag-hunt game, but not ensured in a pure-coordination game. These contrasting conclusions reflect the opposite way that the payoff features of a stag-hunt game impinge on the incentives underlying "credibility" and migration. Concerning credibility, their effect is to render utterly uniformative the statement "I want to play the efficient equilibrium". However, these very same features that destroy such credibility (i.e. a player adopting an inefficient strategy would rather have her opponents play the efficient strategy than her own) are what lead to migration so "easily" disrupting an inefficient convention. Conversely, the very same payoff configuration that makes communication credible in a pure-coordination game also makes an inefficient equilibrium resilient when migration is possible.

Acknowledgements

We are grateful to an anonymous referee for the careful reading of this paper and comments. This paper was initiated during a visit of of the first author to the University of Alicante, financed by the Sabattical Program of the Spanish Ministry of Education, and completed during a simultaneous visit by both authors to the Instituto de Análisis Económico (CSIC) at Barcelona. Throughout, Fernando Vega-Redondo has received financial support from the Spanish Ministry of Education, through the CICYT project 97-0131. We are grateful to the Spanish Ministry of Education for its support.

Appendix A

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 her strategy when the system is at any given state also depends upon the player's former strategy). We may write

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

and therefore,

$$(n-1)[v(s_q, s_r, \omega_1) - v(s_q, s_q, \omega_1)] = \pi(s_q, s_q) - \pi(s_q, s_r)$$

$$(n-1)[v(s_r, s_q, \omega_1) - v(s_r, s_r, \omega_1)] = \pi(s_r, s_r) - \pi(s_r, s_q).$$

Adding the latter two expressions, one obtains

$$(n-1)[v(s_q, s_r, \omega_1) - v(s_q, s_q, \omega_1)] + (n-1)\,[v(s_r, s_q, \omega_1) - v(s_r, s_r, \omega_1]$$
$$= [\pi(s_q, s_q) - \pi(s_r, s_q)] + [\pi(s_r, s_r) - \pi(s_q, s_r)] > 0. \quad \text{(A.1)}$$

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. Consider a player who is currently playing strategy s_q at location ℓ_i, and assume that this player receives a migration opportunity in stage 1. 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

$$\pi(s_q, s_q) > \hat{\theta}\Delta\pi(s_r, s_r) + (1 - \hat{\theta}\Delta)\pi(s_q, s_r),$$

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, she will clearly not change her strategy either.

Consider now the case where the player's first revision opportunity concerns her strategy. Here again, (d) implies that

$$\pi(s_q, s_q) > \hat{\lambda}\Delta\pi(s_r, s_r) + (1 - \hat{\lambda}\Delta)\pi(s_q, s_r).$$

Thus the player will not change her strategy either. Hence $\Gamma(\omega) = \{\omega\}$.

To prove the converse, let ω be any arbitrary state. With positive probability, the player receives no adjustment opportunity in stage 1 and a strategy adjustment opportunity in stage 2 for Q successive periods. Proposition 1 implies that there is a positive-probability transition to a state where each location is homogeneous in terms of the strategy played. On the other hand, with positive probability, all loners (if any exist) can migrate to some location and also adjust to the prevailing strategy. Combining these considerations, we may assert that there exists positive probability that if state ω did not satisfy (a) or (c), a (multi-step) transition takes place to a state ω' that 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 played in each of them by all respective players. If $s_q = s_r$, then it is weakly better for players at one location to migrate to the other, and this occurs with positive probability until the location is empty. If $s_q \neq s_r$, and (d) is not satisfied, then we must have either $\neg(s_q R^*(\Delta) s_r)$ or $\neg(s_r R^*(\Delta) s_q)$. Assume the former, so that either $\neg(s_q R(\hat{\theta}\Delta) s_r)$ or $\neg(s_q R(\hat{\lambda}\Delta) s_r)$. In the former case, with positive probability the players at ℓ_i have the opportunity to migrate sufficiently early in the period so that it is optimal to migrate to ℓ_j and (eventually) adjust strategy to s_r. In the latter case, with positive probability, the players at ℓ_i (who were adopting strategy s_q) change strategy to s_r. Overall, this implies that there is positive probability that, from state ω', another state will be reached satisfying all of conditions (a)–(d). As explained, such a state is absorbing. This shows that no state violating one of the conditions stated 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 $U(\Delta)$, letting $V(\Delta)$ ($\subseteq U(\Delta)$) stand for the set of states which satisfy (i) and (ii) in the statement of the theorem.

As customary in recent stochastic evolutionary literature, our proof shall rely on the graph-theoretic framework developed by Freidlin and Wentzel (1984). 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; and
- from every $\omega' \in \Omega \setminus \{\omega\}$ there exists a *path* $\{(\omega^0, \omega^1), (\omega^1, \omega^2), \ldots, (\omega^{r-1}, \omega^r)\}$ 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 $\omega \in \Omega$,

$$q(\omega) \equiv \sum_{H \in \mathcal{H}_\omega} \prod_{(\omega', \omega'') \in H} T_\epsilon(\omega', \omega''). \tag{A.2}$$

Then, as established by Freidlin and Wentzel (Lemma 3.1, p. 177), we have:

$$\mu_\epsilon(\omega) = \frac{q(\omega)}{\sum_{\omega' \in \Omega} q(\omega')}. \tag{A.3}$$

Each $q(\omega)$ is a polynomial in ϵ. Thus, the *limit invariant distribution* defined by (12) is well defined and, therefore, unique. To compute each $q(\omega)$, it is useful to introduce a cost function on transitions:

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

that 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 $\mathrm{d}(\omega, \omega')$ denotes the number of individual choice dimensions (location or action) differing between ω and ω', then

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

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 (A.3), 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.

First, we establish that every stochastically stable state must belong to $V(\Delta)$. In view of the previous considerations, this follows from the following lemma. □

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

Proof of Lemma 1. Given any state $\hat{\omega} \in U(\Delta) \backslash V(\Delta)$, consider the following two possibilities:

(a) $S(\hat{\omega}) = \{s_q \in S : \omega_{iq} > 0\} = \{s_{\hat{q}}\}$ (i.e. is a singleton); and
(b) $|S(\hat{\omega})| \geq 2$.

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

$$c(\omega', \omega'') = 0 \quad \text{if } \omega' \notin U(\Delta) \text{ (i.e. } \omega' \text{ is not a limit state)}, \tag{A.4}$$

$$c(\omega', \omega'') = 1 \quad \text{otherwise.}$$

In case (a), consider any state ω^1 which is derived from $\hat{\omega}$ by one individual mutating to some new location $\tilde{\ell}$ (and still remain with her original strategy $s_{\hat{q}}$). Now, let ω^2 be the state where the individual who migrated to location $\tilde{\ell}$ switches to some strategy $s_{\tilde{q}} \in \varphi(\Delta)$. The transition from ω^1 to ω^2 can be conducted without any further mutation (i.e. it is costless) since the contemplated switch is weakly optimal for the agent in question. Finally, since $s_{\hat{q}} \notin \varphi(\Delta)$, it also follows that a transition from ω^1 to a state $\tilde{\omega}$ where every individual plays strategy $s_{\tilde{q}}$ in location $\tilde{\ell}$ is also costless. Therefore, the path $\{(\hat{\omega}, \omega^1), (\omega^1, \omega^2), (\omega^2, \tilde{\omega})\}$ satisfies condition (A.4), as desired.

In case (b), let $\bar{\ell}$ denote any of the locations occupied in state $\hat{\omega}$ with its (at least two) individuals adopting some given strategy $s_{\bar{q}}$. Consider a chain of transitions in which, in every one of them, *one* (and only one) individual in locations other than $\bar{\ell}$ mutates to this location, adopting action $s_{\bar{q}}$. All of these transitions involve a unit cost, except for the last one whose cost is zero. Eventually, all individuals are at location $\bar{\ell}$, adopting action $s_{\bar{q}}$. Let $\bar{\omega}$ be the resulting state. If $s_{\bar{q}} \in \varphi(\Delta)$, the desired path has been constructed since then $\bar{\omega} \in V(\Delta)$. Otherwise, simply proceed as described for case (a), connecting $\bar{\omega}$ to some other state $\tilde{\omega} \in V(\Delta)$ 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 $V(\Delta)$ to which $\hat{\omega}$ may be connected through a path $\tilde{h}$ satisfying (A.4). This can be done through the following steps.

(1) Eliminate the arrow $(\tilde{\omega}, \omega')$ in $\hat{H}$ that starts at $\tilde{\omega}$.
(2) 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 clear that, once steps (1)–(2) have been conducted, the original $\hat{\omega}$-tree $\hat{H}$ has been transformed into a well-defined $\tilde{\omega}$-tree. Choose n_o in the statement of the theorem such that, when all individuals are in a certain location choosing a common strategy, at least two individuals are required to play *any* different strategy in that location to make it worthwhile to choose this latter strategy. Formally, we require that for any $s_q, s_r \in S$:

$$(n_o - 2)\,\pi(s_q, s_q) + \pi(s_q, s_r) > (n_o - 3)\,\pi(s_r, s_q) + 2\pi(s_r, s_r).$$

The fact that the game is a strict coordination game ensures that such a n_o exists. If $n \geq n_o$, step (1) above in the construction of the $\tilde{\omega}$-tree saves at least a cost of 2 relative to the original $\hat{\omega}$-tree (here, the point to note is that if $n \geq n_o$, at least two mutations are needed in order to escape the basin of attraction of $\tilde{\omega}$). On the other hand, by adding the arrow from $\hat{\omega}$ included in path $\tilde{h}$, the cost has increased by only one unit. The combination of these considerations establish the lemma.

In order to complete the proof of the theorem, it must still be verified that for all $\hat{\omega}, \tilde{\omega} \in V(\Delta)$,

$$\min_{H \in \mathcal{H}_{\hat{\omega}}} c(H) = \min_{H \in \mathcal{H}_{\tilde{\omega}}} c(H),$$

so that *every* state in $V(\Delta)$ belongs to the support of μ^*. This is a consequence of the following final lemma. □

Lemma 2. *For all $\hat{\omega}, \tilde{\omega} \in V(\Delta)$ and every $\hat{\omega}$-tree $\hat{H} \in \mathcal{H}_{\hat{\omega}}$, there exists an $\tilde{\omega}$-tree $\tilde{H}$ such that $c(\tilde{H}) \leq c(\hat{H})$.*

Proof of the Lemma 2. Consider any $\hat{\omega}, \tilde{\omega} \in V(\Delta)$, and let $s_{\hat{q}}$ and $s_{\tilde{q}}$ be the respective (single) conventions displayed by them, with all agents located at respective locations $\hat{\ell}$ and $\tilde{\ell}$.

Assume, for simplicity,[15] that $\hat{\ell} \neq \tilde{\ell}$ and consider any $\hat{\omega}$-tree $\hat{H}$. To construct an $\tilde{\omega}$-tree $\tilde{H}$ that satisfies $c(\tilde{H}) \leq c(\hat{H})$, we define a suitable path $\tilde{h}$ joining $\hat{\omega}$ to $\tilde{\omega}$. First, let ω^1 be the state obtained from $\hat{\omega}$ when one agent mutates to location $\tilde{\ell}$ (and still chooses $s_{\hat{q}}$). Second, denote by ω^2 the state obtained from ω^1 where the single agent at location $\tilde{\ell}$ switches to strategy $s_{\tilde{q}}$. Third, consider the state ω^3 obtained from state ω^2 when an agent previously at location $\hat{\ell}$ mutates to location $\tilde{\ell}$ and is then given the unilateral option to revise her strategy (thus choosing $s_{\tilde{q}}$). Subsequently, consider the chain of additional states $\{\omega^4, \omega^5, \ldots, \omega^{n+1}\}$ where a new agent migrates from $\hat{\ell}$ to $\tilde{\ell}$ and unilaterally adjusts her strategy. Obviously, $\omega^{n+1} = \tilde{\omega}$. Furthermore, we have that $c(\hat{\omega}, \omega^1) = 1$, $c(\omega^1, \omega^2) = 0$, and $c(\omega^k, \omega^{k+1}) = 1$ for each $k = 2, 3, \ldots, n$. Therefore, $c(\tilde{h}) = n$.

Consider now a transformation of the contemplated $\hat{\omega}$-tree $\hat{H}$ into an $\tilde{\omega}$-tree as follows. Delete the arrow leaving $\tilde{\omega}$ in $\hat{H}$ and add $(\hat{\omega}, \omega^1)$. Then, replace the arrows in $\hat{H}$ that leave the states in $\{\omega^1, \omega^2, \ldots, \omega^n\}$ with the corresponding arrows in the path $\tilde{h}$. Clearly, if $\hat{h}$ denotes the path joining $\tilde{\omega}$ to $\hat{\omega}$ in the original $\hat{\omega}$-tree $\hat{H}$, we have that

$$c(\hat{H}) - c(\tilde{H}) = c(\hat{h}) - c(\tilde{h}).$$

By construction, $\tilde{h}$ is a cost-minimal path linking two states in $V(\Delta)$ where the population is concentrated at different locations. Thus, $c(\hat{h}) - c(\tilde{h}) \geq 0$ and the desired conclusion follows in this case. The proof of the theorem is thus complete. □

References

Aumann, R., 1993. Nash equilibria are not self-enforcing. In: Gabszewicz, J., Thisse, J.-F., Wolsey, L. (Eds.), Economic Decision Making: Games, Econometrics and Optimization. Elsevier, Amsterdam, pp. 201–206.

Binmore, K., Samuelson, L., 1997. Muddling through: noisy equilibrium selection. Journal of Economic Theory 74, 235–265.

Dieckmann, T., 1998. The evolution of conventions with mobile players. Journal of Economic Behavior and Organization 38, 93–111.

Ellison, G., 1993. Learning, local interaction, and coordination. Econometrica 4, 1047–1073.

Ellison, G., 2000. Basins of attraction, long-run stochastic stability and the speed of step-by-step evolution. Review of Economic Studies 67, 17–46.

Ely, J., 1995. Local conventions. Advances in Theoretical Economics 2 (1), Article 1, http://www.bepress.com/bejte/advances/vol2/iss1/art1.

Farrell, J., 1988. Communication, coordination and Nash equilibrium. Economics Letters 27, 209–214.

Foster, D., Young, P., 1990. Stochastic evolutionary game dynamics. Theoretical Population Biology 38, 219–232.

Freidlin, M.I., Wentzel, A.D., 1984. Random Perturbations of Dynamical Systems. Springer-Verlag, New York.

Kandori, M., Mailath, G., Rob, R., 1993. Learning, mutations, and long-run equilibria in games. Econometrica 61, 29–56.

[15] Observe that all those states where a *given* convention is played at different locations must be jointly included in, or excluded from, the support of μ^*. This implies that the assumption $\hat{\ell} \neq \tilde{\ell}$ can be made without loss of generality.

Kandori, M., Rob, R., 1995. Evolution of equilibria in the long-run: a general theory and applications. Journal of Economic Theory 65, 383–414.
Kim, Y.-G., Sobel, J., 1995. An evolutionary approch to pre-play communication. Econometrica 63, 1181–1193.
Mailath, G., Samuelson, L., Shaked, A., 1995. Evolution and endogenous interactions. Social Systems Research Institute, Working Paper No. 9426, University of Wisconsin–Madison.
Oechssler, J., 1997. Decentralization and the coordination problem. Journal of Economic Behavior and Organization 32, 119–135.
Young, P., 1993. The evolution of conventions. Econometrica 61, 57–84.