EXPECTATIONS, DRIFT AND VOLATILITY
IN EVOLUTIONARY GAMES*

Fernando Vega-Redondo**

WP-AD 94-02

    ** Instituto Valenciano de Investigaciones Económicas and University of Alicante.

EXPECTATIONS, DRIFT AND VOLATILITY
IN EVOLUTIONARY GAMES

Fernando Vega-Redondo

A B S T R A C T

This paper introduces expectations into the framework of evolutionary games. On the one hand, (myopic) players are assumed to behave optimally according to the expectations they hold at each point of the process. On the other hand, expectations themselves are continuously updated according to the players' latest experience. The possibility of random drift on expectations (i.e., arbitrary variation on them not opposed by selection forces) produces sharp volatility across equilibria. Specifically, all Nash equilibria (but only these) have positive weight in the limit stationary distribution, independently of risk -or payoff-dominance considerations.

KEYWORDS: Drift, Volatility, Evolution.

1.- INTRODUCTION

Recent evolutionary models have provided substantial insight into the important issue of equilibrium selection in games. (See, for example, Foster & Young (1990), Canning (1992), Kandori, Mailath & Rob (1993), or Young (1993).) Most of this evolutionary literature has focused on a theoretical framework which incorporates the following two components:

(a) Selection (or sometime called "Darwinian") dynamics, whereby players adjust their strategies in the direction of what is best response to the current situation.

(b) Experimentation (or "mutation") dynamics, by which players occasionally explore strategies not prescribed by (a).

Formally, the combination of the above two dynamics induces an ergodic Markov chain, on which the following question is asked: What is the long-run behavior of the system when experimentation becomes a relatively rare phenomenon ? In some interesting contexts (for example, a population which is randomly matched to play a 2×2 coordination game) this literature is able to provide a clearcut answer: In the long run, it is the risk-dominant equilibrium which will be played by the population "most of the time" (cf. Kandori, Mailath & Rob (1993)).

Recently, Robson (1993) and Vega-Redondo (1993) have pointed out that the former conclusion must be modified if, instead of (a) above, an alternative selection dynamics based on imitation is contemplated. Specifically, they propose a formulation by which players simply shift (perhaps only gradually) towards those strategies that currently enjoy a higher average payoff in the population. In this case, the Pareto efficient equilibrium (possibly risk-dominated) ends up being played most of the time in the long run, as the experimentation probability becomes small. Thus, somewhat paradoxically, coordination on Pareto efficient outcomes turns out to be more easily achieved by agents who, in a certain sense, are more "boundedly rational" than in the Kandori *et al* (1993) context (i.e., agents who can just "imitate" rather than "best-respond" to the current situation).

This paper deviates from (a) above in a direction somewhat opposite to that of Robson (1993) and Vega-Redondo (1993). It also focuses on a context where agents are randomly matched every period to play a simple coordination game. Players, however, are assumed substantially more sophisticated than in Kandori *et al* (1993). Specifically, they are assumed able to hold expectations on the future state of the system which, unlike what is implicitly considered by (a), need not be "static". That is, in the present context, agents' expectations may allow for the possibility that the strategy profile prevailing next period might be different from the current one.

Formally, each player's "expectations" is identified with a mapping that, to every possible strategy configuration currently observed, associates some probability measure over the next strategies adopted by the remaining players. At any point in the process, each player is taken to react optimally to her current expectations. On the other hand, she is also assumed to update her expectations in order to make them consistent with the latest (observed) realization of the process (see Section 2 for details). The combination of optimal responses and expectation updating just described defines the "Darwinian" part of the process. As standard, this dynamics is then complemented by a "mutation" component which, for each agent, induces independent and arbitrary transitions.

In this context, the following customary question is posed: What is the long-run behavior of the system as the mutation probability becomes small ? Now, the answer turns out to be substantially different from any of those described before. Specifically, it will be shown that, in the present expectation-based context, all Nash equilibria (but only these) are played some significant fraction of time in the long run. Thus, in contrast with the received evolutionary literature, we may speak of "equilibrium volatility" rather than "equilibrium selection", as a description of the long-run dynamics of the process.

It may be worthwhile at this point to advance, only informally, which is the basic theoretical idea underlying such conclusion of volatility. In a nutshell, gradual accumulation of drift, as applied to the expectation side of the process, can be singled out to be the key factor inducing such state of affairs.

6

To explain matters simply, consider a stationary realization of the process in which a given (equilibrium) strategy of the coordination game, say A, is both continuously played and continuously expected by all players. In the absence of any experimentation with alternative strategies, this situation will remain stationary. But then, most of the <u>components</u> of the expectations prevailing along the process will be essentially "redundant" (i.e., will refer to the merely hypothetical and unobserved contingencies where some player adopts a strategy different from A). This leaves open the possibility that "mutation" on these redundant components may <u>accumulate</u> without experimenting any adverse selection pressure. When enough such mutations have taken place in a suitable direction, transitions to other equilibrium strategy profiles (say, to one where all players adopt a new strategy B) can be simply triggered by a small strategy deviation from the original situation. The only requirement for this to unfold is that prevailing expectations happen to "interpret" such small deviation as a "signal" for the contemplated transition.

The considerations just outlined are quite reminiscent of the role attributed by some biologists (e.g. Kimura & Otha (1971)) to the phenomenon of random drift in processes of natural selection. Although there is neither a full understanding nor agreement among biologists about the implications of genetic drift, it is often considered to be one of the main forces inducing the wide variability observed in biological contexts. The reasons proposed in this respect (see, for example, Futuyama (1979, pp. 272-78) bear some resemblance to those outlined above: random drift may allow for the evolution of certain characteristics which, under the "right" future circumstances, could prove highly successful and become predominant.

In a quite different vein, the model proposed here may also be seen as a very stylized formalization of a traditional black box in economics: the so-called "animal spirits" (see Weil (1989)). Here, those spirits, animal or otherwise, are conceived as the outcome of a random process which proceeds in the background, thus remaining unnoticed by "external observers". When this process is already "mature" (i.e., the right number and type of mutations have gradually accumulated), its effects may become abruptly apparent in response to relatively minor disturbances. Of course, this may only represent a very incomplete account of such a complex phenomenon. A more satisfactory approach

should also include, *inter alia*, a richer specification of the sources of "mutation" at both the strategy and expectation levels.

The remainder of the paper is organized as follows. Next, Section 2 presents the model. In Section 3, the formal analysis is carried out. Finally, Section 4 concludes with a discussion of the results, as well as some generalizations. For the sake of smooth presentation, proofs are contained in the Appendix.

2.- THE MODEL

Consider a finite population of players, indexed by $i = 1,2,...,N$, with N even. Every period $t = 1,2,...$, they are assumed randomly matched in pairs to play a bilateral symmetric game, with strategy set $S = \{A,B\}$ and payoff function $\pi: S \times S \longrightarrow \mathbb{R}$ summarized by the following table:[1]

	A	B
A	a,a	d,c
B	c,d	b,b

Figure 1

We shall focus on the case of a coordination game with two pure-strategy, symmetric, and strict Nash equilibria: (A,A) and (B,B). The other two possibilities (where the game has either a weakly dominant strategy or a unique mixed-strategy equilibrium) yield a much less interesting and relatively trivial analysis. Thus, in terms of the payoffs specified in Figure 1, we shall consider the case where the following inequalities hold: $a > d$, $b > c$. Moreover, in order to abstract from complications that are peripheral to our main concerns, we shall also rule out the degenerate case where both equilibria have exactly the same "basins of attraction". That is, we shall assume that $a+d \neq c+b$.

---

(1)Here, $\pi( )$ stands for the payoff for the row player.

Within each period, we may either assume that players are matched "many times", or that there is just one round of play. Both scenarios are consistent with our formulation. The <u>state</u> of the system at any given time t is identified with the 2N-tuple $\omega(t) \equiv \left(s_i(t),e_i(t)\right)_{i=1}^{N}$ which specifies the strategy $s_i(t)$ and expectations $e_i(t)$ adopted by each player $i = 1,2,...,N$ at t. The set of all such possible states will be denoted by Ω.

Only pure strategies will be allowed. Thus, $s_i(t) \in S = \{A,B\}$ for each player $i=1,2,...,N$. On the other hand, the <u>expectations</u> $e_i(t)$ of each player i (in fact, a "pattern" of expectations) are assumed mappings of the form:

$$e_i(t): S^N \longrightarrow \Delta(S^{N-1}) \tag{1}$$

which, for any current strategy profile $s(t) \equiv \left(s_1(t),...,s_N(t)\right) \in S^N$ assigns a corresponding probability measure $e_i(t)\left(s(t)\right) \in \Delta(S^{N-1})$ over the future strategy profile $s_{-i}(t+1) \in S^{N-1}$ played by the other players $j \neq i$ at t+1. This probability measure $e_i(t)\left(s(t)\right) \equiv p_i(t)$ will be labelled the <u>beliefs</u> of player i at t.

As advanced, the strategy choice $s_i(t+1)$ of each player i at t+1 will be required to be a best response to the <u>beliefs</u> $p_i(t)$ which this player holds at t about $s_{-i}(t+1)$. Define:

$$\beta(s_i,p_i) \equiv \sum_{\hat{s}_{-i} \in S^{N-1}} \left\{ p_i(\hat{s}_{-i}) \left[\sum_{j \neq i} \frac{1}{N-1} \pi(s_i,\hat{s}_j)\right] \right\} \tag{2}$$

as the expected payoff for player i of playing s_i if she is randomly matched against the other N-1 players and she attributes *ex-ante* probability $p_i(\hat{s}_{-i})$ that her opponents play each strategy profile $\hat{s}_{-i} \equiv (\hat{s}_j)_{j \neq i}$.

We shall require that, for all $t \in \mathbb{N}$ and every $i = 1,2,...,N$,

$$s_i(t+1) \in BR\left(p_i(t)\right) \equiv \left\{ s_i^* \in S: \beta\left(s_i^*,p_i(t)\right) \geq \beta\left(s_i,p_i(t)\right), \forall s_i \in S \right\}. \tag{3}$$

The above expression formalizes the first part of the selection (or Darwinian) dynamics. When the set $BR\left(p_i(t)\right)$ is not a singleton (i.e., both strategies are best responses) the player will be assumed to use some arbitrary rule to

select her strategy. For simplicity, we shall suppose that such rule is "anonymous" (i.e., symmetric on players' indices) and deterministic.[2]

The second part of the selection dynamics involves the updating of expectations. Given the state $\omega(t) \equiv \left(s_i(t), e_i(t)\right)_{i=1}^{N}$ prevailing at some given t and the strategy profile $s(t+1) = \left(s_i(t+1)\right)_{i=1}^{N}$ played at t+1, the following updating rule is postulated for the expectations $e_i(t+1)$ of each player i = 1,2,...N:

$$e_i(t+1)(s) = e_i(t)(s), \qquad \text{for } s \neq s(t)$$
$$= \delta[s_{-i}(t+1)], \text{ for } s = s(t), \tag{4}$$

where $\delta[s_{-i}] \in \Delta(S^{n-1})$ stands for the beliefs concentrated on s_{-i}, i.e., $\delta[s_{-i}](s_{-i}) = 1$.

Thus, we simply assume that the players "match" the latest experience observed. Specifically, for any strategy profile s different from the previous s(t), e(t+1) induces the same beliefs as in e(t). (The interpretation here is that the latest observation "s(t+1) has followed s(t)" is not relevant for beliefs associated to a configuration $s \neq s(t)$.) For s = s(t), on the other hand, e(t+1) simply postulates that the latest observation will be fully matched. Admittedly, this formulation is somewhat extreme in that expectation revision is quite drastic. It has the advantage, however, of being especially tractable. It seems quite plausible that more flexible rules (such as, for example, "increase the probability attributed to the observed transition") would still yield similar conclusions under an appropriate reformulation.

The combination of expressions (3) and (4) define the selection dynamics. Its transitions can be described through a function:

$$6$$
$$D: \Omega \longrightarrow \Omega . \tag{5}$$

As explained, such selection dynamics will be complemented by "mutation", which is formulated as follows. Every t, each player i=1,2,...,N has some positive and independent probability $\varepsilon > 0$ of making a transition for her

---

(2) Allowing for general stochastic rules which may depend on players' indices simply makes the theoretical framework more cumbersome, without affecting the nature of our conclusions.

corresponding part of the state $\left(s_i(t+1),e_i(t+1)\right)$ which differs from that prescribed by $D(\cdot)$.[3] In that event, it will be assumed that she chooses every other (s_i',e_i') with positive probability, bounded above zero for all t.

Let

$$\mathcal{E} : \Omega \longrightarrow \Delta(\Omega) \tag{6}$$

formalize the (stochastic) transition due to mutation just described. The full description of the process ("selection plus mutation") may be compactly described by the function

$$F : \Omega \longrightarrow \Delta(\Omega) \tag{7}$$

defined by $F(\omega) = \mathcal{E}\left(D(\omega)\right) \in \Delta(\Omega)$ for all $\omega \in \Omega$. Since the set Ω is finite, $F(\cdot)$ defines a Markov chain. Its transition matrix will be denoted by $P(\varepsilon)$, in order to express explicitly its dependence on ε, a key parameter of our ensuing analysis.

3.- ANALYSIS

Because of players' mutation, the Markov chain induced by $F(\cdot)$ is obviously aperiodic with a single recurrent class. Thus, by standard results (c.f. Karlin & Taylor (1975, Theorem 1.3, p. 35)), it has a unique stationary distribution. Let us denote this distribution by $\mu(\varepsilon) \in \Delta(\Omega)$.

Conceiving of mutation as a quite infrequent phenomenon, it is natural to think of ε as relatively small. Specifically, we shall focus on the limit case where ε converges to zero, as a a especially convenient benchmark of analysis. The long-run behavior of the system under these circumstances is summarized by the limit stationary distribution $\mu^* \in \Delta(\Omega)$ defined by:

$$\mu^*(\omega) = \lim_{\varepsilon \to 0} \mu(\varepsilon)(\omega), \quad \forall \omega \in \Omega. \tag{8}$$

---

(3) It is inessential that all players have the same probability of experimentation. They could well have different such probabilities, as long as they all converge to zero at the same rate in the limit exercise contemplated below for the limit stationary distribution.

11

We can also think of μ^* as embodying a certain robustness criterion on what may be conceived as the main component of the evolutionary process: the selection dynamics. More specifically, the limit stationary distribution can be regarded as singling out those long-run predictions of the selection dynamics which are robust to the introduction of some arbitrarily small perturbation (i.e., mutation) by which players occasionally behave "erratically". (See Kandori *et al* (1993) for an elaboration on this motivation.)

Given μ^*, let $\lambda^* \in \Delta(S^N)$ denote the (marginal) probability measure on strategy profiles induced by μ^*, i.e.,

$$\lambda^*(s) = \sum_{\omega \in \Omega(s)} \mu^*(\omega), \tag{9}$$

where $\Omega(s) \equiv \{\omega = (s_i, e_i)_{i=1}^N : (s_1, ..., s_N) = s\}$. We are now in a position to state the main result of the paper.

Theorem: *There exists some* $\underline{N}$ *such that if* $N \geq \underline{N}$, *then* **supp**$(\lambda^*) = \{(A,A,...A),(B,B,...,B)\}$.

By the previous theorem, the long-run behavior of the system is (almost) fully concentrated (if ε is small) on the two "monomorphic" strategy profiles. Each of these strategy profiles will be observed a significant fraction of time (bounded above zero, as ε converges to zero) in the long run. At each of them, the whole population plays either one or the other equilibrium of the game in every one of their bilateral encounters.

The theorem, however, does not provide us information on what are the particular expectations and dynamic paths (possibly cyclic) which underlie its conclusion. To gain further insight into this point, define the following set:

$$Q = \left\{ \omega = (s_i, e_i)_{i=1}^N \in \Omega : \exists h, h' \in \{A, B\} \text{ s.t. } \forall i = 1, 2, .., N, \right.$$
$$\left. s_i = h, \ e_i(h, ..., h) = \delta[(h', ..., h')] \right\} \tag{10}$$

12

The set $Q \subset \Omega$ consists of all those states which satisfy:

(a) each player adopts the same strategy $s_i = h \in \{A,B\}$;

(b) each player has the same prediction that a certain strategy $h' \in \{A,B\}$ (possibly different from the one currently played h) will be adopted by all players next period.

Thus, the set Q consists of all those states which, in the absence of mutation, have the population always play and rightly predict one of the two equilibria, possibly in alternation. These states, moreover, allow for any expectations off the "realized path". (It is precisely such degree of freedom off the realized path which permits the drift on expectations, as explained below.) The next result, establishes that all states in the set Q will be visited a significant fraction of time in the long run.

Proposition: *There exists some* $\underline{N}$ *such that if* $N \geq \underline{N}$, *then* $\mathrm{supp}(\mu^*) = Q$.

Proof: See the Appendix.

The preceding Proposition indicates that the equilibrium volatility established across monomorphic strategy profiles will materialize in two different ways. On the one hand, there will be long stretches of time (for ε small) in which one of the two monomorphic profiles will be continuously adopted by the population. On the other hand, 2-period cycles will also occur a significant fraction of the time, with the population alternating between the two monomorphic strategy profiles.

4.- DISCUSSION

A useful way of understanding our results is by relying on the notion of *expectation components* (see Samuelson (1993) for a related, but somewhat different, concept). An expectation component is defined as a subset $X \subset \Omega$ which is maximal with respect to the following two properties::

13

(a) $X = \bigcup_{k=1}^{r} X_k$ where each $X_k \subset \Omega$ is a <u>limit</u> <u>set</u> of the selection

dynamics,[4] and every X_k induces the <u>same</u> set of long-run strategy profiles. That is, there is some $Y \subseteq S^N$ such that, for all $k = 1,2,...,r,$

$$\{s \in S^N: \exists \omega = (s_i,e_i)_{i=1}^N \in X_k \text{ s.t. } s = (s_1,...,s_n)\} = Y. \tag{11}$$

(b) For every $k,k' = 1,2,...,r,$ $k \neq k'$, there exist a sequence of distinct indices $\{k_1,k_2,...,k_s\}$ with $k_1=k$ and $k_s=k'$ such that, for every k_j, $j=1,2,...,s-1$, there are some $\omega_j \in X_j$, $\omega_{j+1} \in X_{j+1}$ which differ in (exactly) <u>one</u> single mutation (i.e., there is only one player whose state component differs between ω_j and ω_{j+1}.)

It is not difficult to see (cf. the Appendix for details) that the process has only the following three expectation components:

$$C_A \equiv \left\{ \omega = (s_i,e_i)_{i=1}^N \in \Omega: \forall i=1,2,..,N, \ s_i = A, \ e_i(A,...,A) = [(A,...,A)] \right\}$$

$$C_B \equiv \left\{ \omega = (s_i,e_i)_{i=1}^N \in \Omega: \forall i=1,2,..,N, \ s_i = B, \ e_i(B,...,B) = \delta[(B,...,B)] \right\}$$

$$C_{AB} \equiv \left\{ \omega = (s_i,e_i)_{i=1}^N \in \Omega: \forall i=1,2,..,N, \ h,h' \in \{A,B\}, \ h \neq h', \right.$$
$$\left. s_i = h, \ e_i(h,...,h) = \delta[(h',...,h')] \right\}.$$

Note that the three above components jointly define a partition of the set Q. <u>Within</u> each of them, transitions across <u>any</u> two states simply require the operation of the selection dynamics and sufficient accumulation of <u>single</u> mutations, the same number of them in each direction. Thus, if any state in a given component belongs to the support of the limit stationary distribution, so must be the case for any other state of the component.

---

[4] Given a certain Markov chain, a limit set is defined as a set of states Z which satisfy: (i) if the process lies in Z at some point in time, it remains in it forever after (with probability one); (ii) every state in Z is reached with positive probability from any other state in Z after a pre-specified finite number of transitions. See the Appendix for a formal definition of this standard concept.

As suggested in our previous discussion, these transitions within components capture the intuitive idea of "drift". All states of the same component have identical expectations along the (common) path of strategy profiles that they induce, but differ away from this path. Since contingencies off the realized path are irrelevant as long as players' do not deviate from this path, "mutation" on expectations associated to those off-the-path contingencies can arise and accumulate without experimenting any selection forces (for or against).

At some point, however, some mutation on the realized path must arise if the process is to ever abandon a given component. The key step in the proof of our results is to show that, for any two components, the same minimum number of mutations is needed on the realized path in order to make the transition feasible in either direction. Such "easiest" transitions may only be implementable, however, via some pair of appropriately selected states in each respective component. Thus, in this perspective, drift performs a crucial task in making any transition away from a given component possible: it directs the system towards some state within this component where the contemplated transition is easiest.

I conclude with some brief comments on possible generalizations of the model as well as possible avenues for further research. It is easy to check that the essential gist of our analysis, i.e., the volatility of equilibrium behavior under expectation drift, remains valid for any context in which players are randomly matched to play a general (symmetric) coordination game. In the long run, all equilibria will be observed a significant fraction of the time, with non-equilibrium behavior being a ephemeral state of affairs which only occurs in the processes of transition across equilibrium components.

For games which are not of coordination, the situation is bound to be substantially more complex. In general, we must expect that non-equilibrium behavior will be observed in the long run, even for infinitesimal mutation probability. This is also the state of affairs prevailing in received evolutionary models, where only restrictive conditions on the game being played allows one to ensure the long-run convergence to equilibrium behavior. (See Kandori & Rob (1992) or Young (1993).) In any case, the volatility across "components" (not necessarily equilibrium components) would still seem to

15

remain essentially applicable in such general contexts, again as a consequence of the interaction between expectations and drift.

As suggested by the previous comments, to extend the realm of application to more general games and matching contexts should certainly be one of the important topics of future research in evolutionary game theory. There is, however, another line of generalization of present evolutionary models which is also naturally suggested by the present paper: the consideration of alternative behavioral paradigms.

In the Introduction of the paper we referred to three such paradigms:

- imitation (Robson (1993), Vega-Redondo (1993)),
- "static" best response (Kandori, Mailath & Rob (1993), Young (1993)),
- "expectation-based" best response, defined by (3) and (4).

As discussed, the implications of each of these behavioral rules are quite different. Consequently, the following natural question arises: What would be the long-run outcome of an evolutionary process where, say, the three of them were integrated into a single framework. In such an evolutionary process, it would seem reasonable to allow for the possibility that, occasionally, each agent should be able to switch across behavioral rules (not only strategies) depending on their relative performance. (Of course, this performance could include different "implementation costs" associated to the diverse complexity of the different rules.)

Ideally, the long-run analysis of the process would provide some insight on the relative strengths of alternative behavioral rules, thus yielding some criterion for both behavioral and equilibrium "selection". The recent work of Banerjee & Weibull (1992) and Stahl (1993) may be interpreted as a first step along this direction.

A P P E N D I X

<u>Proof of the Theorem</u>:

It is easy to check that every state in Ω which belongs to the support of the limit stationary distribution μ^* must also belong to some limit set of the <u>selection</u> dynamics induced by $D(\cdot)$. (The argument is a direct adaptation a of similar one used, for example, by Kandori *et al* (1993). It simply relies on the upper hemi-continuity of $\mu(\varepsilon)$ with respect to ε.) Thus, a first step in the argument is to find the limit sets of $D(\cdot)$. For the sake of completeness, a formal definition of the standard concept of limit (or absorbing) set follows, as particularized to the Markov chain (induced by) $D(\cdot)$.

<u>Definition</u>: *A set* $\Theta \subseteq \Omega$ *is a limit set of the Markov chain* $D(\cdot)$ *if:*
(1) $\forall$ $\omega \in \Theta$, $D(\omega) \in \Theta$.
(2) $\forall$ $\omega,\omega' \in \Theta$, $\exists m \in \mathbb{N}$ *s.t.* $D^{(m)}(\omega) = \omega'$.

Denote by $\mathscr{L}$ the set of limit sets of $D(\cdot)$. (Note that, obviously, any two different limit sets must be disjoint.) The required information on the limit sets of $D(\cdot)$ is contained in the following Lemma.

<u>Lemma 1</u>: *There exists some* $\underline{N}$ *such that if* $N \geq \underline{N}$, *then* $\bigcup_{\Theta \in \mathscr{L}} \Theta = Q$, *where* Q *is defined in (10)*.

<u>Proof</u>: The inclusion $Q \subseteq \bigcup_{\Theta \in \mathscr{L}} \Theta$ is obvious (i.e., every state in Q clearly belongs to some limit set of $D(\cdot)$). Therefore, only the converse inclusion needs to be established. Specifically, it is enough to show that, for any given limit set Θ of $D(\cdot)$, we have $\Theta \subset Q$.

First, we introduce some additional notation. Given any $s \in S^N$, denote:

$$\gamma_i(s) \equiv \{\tilde{s}_i \in S: \sum_{\substack{j \neq i}} \frac{1}{N-1} \pi(\tilde{s}_i, s_j) \geq \sum_{\substack{j \neq i}} \frac{1}{N-1} \pi(s_i', s_j), \forall s_i' \in S\}, \tag{12}$$

17

i.e., the strategies which are best response for player i to the "residual" strategy profile s_{-i} induced by s. Relying on the correspondences $(\gamma_i)_{i=1}^N$ we then define the following set

$$\Lambda \equiv \{s \in S^N: \bigcup_{i=1}^N \gamma_i(s) = S\}, \tag{13}$$

which consists of all those strategy profiles for which the union of the induced best-response correspondences $(\gamma_i(s))_{i=1}^N$ spans the whole strategy space S.

We may now proceed to showing the desired inclusion. Since Ω is finite, the <u>whole</u> limit set considered, Θ, must define a <u>single</u> finite cycle $(\omega^1, \omega^2, ..., \omega^m)$ which satisfies:

$$D(\omega^r) = \omega^{[r+1]}, \quad r=1,2,...,m, \tag{14}$$

where [·] stands for "modulo m". Let s^r and e^r, r=1,2,...,m, denote the strategy and expectation projections induced by each respective ω^r. By the expectation updating rule (4), expectations e_i^r must satisfy that:

$$e_i^r(s^r) = \delta[\hat{s}_{-i}^r], \quad i =1,2..., \ r = 1,2,...m, \tag{15}$$

where $\hat{s}^r \in S^N$ is some <u>common</u> strategy profile for every player. Specifically, $\hat{s}^r$ simply coincides with the profile $s^{[r'+1]}$, for the latest preceding state $\omega^{r'}$ for which $s^{r'}= s^r$. (Of course, it may occur that r' = r. Moreover, for stationary states ω^r, $\hat{s}^r = s^r$.)

Combining (15) and (3), we may conclude that, for all r=1,2,...,m, we must have one of the following three possibilities:[5]

(i) $s^r = (A,A,...,A)$;

(ii) $s^r = (B,B,...,B)$;

(iii) $\hat{s}^{r-1} \in \Lambda$.

_____

[5] The third possibility does not exclude (i) or (ii).

However, from our genericity assumption on the game payoffs (i.e., a+d ≠ c+b), (iii) above can be excluded if N is large enough. For, in this case, it is immediate to check that:[6]

$$s \in \Lambda \implies \forall s' \in \left(\gamma_i(s)\right)_{i=1}^{N}, \ s' \notin \Lambda, \tag{16}$$

which precludes that any $\hat{s}^r$ along a cycle of states induced by $D(\cdot)$ may belong to Λ. This completes the proof of the Lemma.

As noted in the text, the set Q may be partitioned into the three expectation components C_A, C_B, and C_{AB}. To complete the proof of the Theorem, we need to establish the following two additional lemmas with respect to these components.

Lemma 2: *For every two states* ω, ω' *belonging to any* <u>one</u> *expectation component,* C_A, C_B, *or* C_{AB}, $\mu^*(\omega) > 0 \iff \mu^*(\omega') > 0$.

Lemma 3: *For any two states* ω, ω' *such that* $\omega \in C_A$, $\omega' \in C_B$, $\mu^*(\omega) > 0 \iff \mu^*(\omega') > 0$.

Proof of Lemma 2 (sketch): The proof of Lemma 2 relies on an argument of Samuelson (1993, Theorem 2). He proposes the notion of "adjacent limit sets": two limit sets are adjacent if a transition across them can be implemented with only one mutation (cf. requirement (b) in our definition of expectation components in Section 4). He then shows that a limit set belongs to the

---

support of the limit stationary distribution if, and only if, all the adjacent limit sets also do. I shall not prove formally this conclusion, although the intuition underlying it should be clear: the transitions across two limit sets that are adjacent are of the order ε, as $\varepsilon \longrightarrow 0$, in _either_ direction. Thus, both limit sets should have weights of comparable order in the limit stationary distribution. By an appropriate construction of a chain of single mutations joining any two limit sets that belong to a certain expectation component, the same conclusion is extendable to any two such sets.

Proof _of_ _Lemma_ $\underline{3}$: We may rely on an idea similar to that of the previous proof. It is enough to show that, for each $h = A,B$, there is some states $\tilde{\omega} \in C_h$, $\hat{\omega} \in C_{h'}$, $h' \neq h$, such that $\hat{\omega}$ can be reached from $\tilde{\omega}$ through the operation of $D(\cdot)$ and just _one_ mutation. For concreteness, let $h = A$ and $h' = B$ in the previous statement. Consider the state $\tilde{\omega} = \left(\tilde{s}_i, \tilde{e}_i\right)_{i=1}^N$, defined as follows, for all $i = 1, 2, ..., N$:

$$\tilde{s}_i = A,$$
$$\tilde{e}_i(A,...,A) = \delta[(A,...,A)], \tag{17}$$
$$\tilde{e}_i(s) = \delta[(B,...,B)], \quad \forall s \neq (A,...,A).$$

From $\tilde{\omega}$, it only takes one mutation on the strategy chosen by any player for the selection dynamics $D(\cdot)$ to perform a transition to the state $\hat{\omega} = \left(\hat{s}_i, \hat{e}_i\right)_{i=1}^N \in C_B$ given by:

$$\hat{s}_i = B,$$
$$\hat{e}_i(A,...,A) = \delta[(A,...,A)], \tag{18}$$
$$\hat{e}_i(s) = \delta[(B,...,B)], \quad \forall s \neq (A,...,A).$$

This completes the proof of the Lemma.

From the previous three lemmas, two possibilities can arise. One is that $C_{AB} \subset \text{supp}(\mu^*)$. In that case, the conclusion of the Theorem follows directly. Otherwise, if $C_{AB} \not\subset \text{supp}(\mu^*)$, then $C_{AB} \cap \text{supp}(\mu^*) = \varnothing$ and $C_h \subset \text{supp}(\mu^*)$ for some $h \in \{A,B\}$. But by Lemma 3, we must then have

$$\text{supp}(\mu^*) = C_A \cup C_B, \tag{19}$$

which again proves the desired conclusion. ∎

20

Proof of the Proposition:

Denote $C_{AVB} \equiv C_A \cup C_B$. From our previous considerations, it is enough to show that the transitions across C_{AVB} and C_{AB} require the same minimum number of mutations (together with the operation of the selection dynamics) in either direction. In fact, this common number of mutations turns out to coincide with those required to "escape the "basin of attraction" of the risk-dominated equilibrium.

Let us focus on the transition from C_{AVB} to C_{AB} and assume, without loss of generality, that a+d < c+b, i.e., the equilibrium (A,A) is risk-dominated by (B,B). Consider some state $\tilde{\omega} \equiv \left(\tilde{s}_i, \tilde{e}_i\right)_{i=1}^{N} \in C_{AVB}$ satisfying:

$$
\begin{aligned}
\tilde{s}_i &= A, \\
\tilde{e}_i(A,...,A) &= \delta[(A,...,A), \\
\tilde{e}_i(B,...,B) &= \delta[(A,...,A)].
\end{aligned}
\tag{20}
$$

From such a state $\tilde{\omega}$, let $\eta \in \mathbb{N}$ be the minimum number of mutations on players' strategy strategy choices which transforms the profile $\tilde{s} = (A,...,A)$ into some strategy profile $\hat{s} \in \Lambda$, as defined in (13). By our assumption on payoffs, $\eta < N/2$ if N is large enough.

Assume η individuals simultaneously mutate into strategy B from state $\tilde{\omega}$ at some t_0. Then, the expectation updating rule (4) implies that the next time $t_1 > t_0$ at which the selection dynamics induces a strategy profile $s(t_1) = (A,...,A)$ the number of players adopting strategy B at $s(t_1+1)$ will be larger than N/2. (Note that such a t_1 is bound to occur in finite time because of Lemma 1 and the expectations $\tilde{e}_i(B,...,B) = \delta[(A,...,A)]$ which are postulated in (20).) But then, the next $t_2 > t_1$ such that $s(t_2) = (A,...,A)$, we shall have $s(t_2+1) = (B,...,B)$ and a limit set in C_{AB} will have been reached. (Of course, the same comment on the finiteness of t_2 is applicable here as before for t_1.)

The previous considerations show that only η mutations are needed to perform the transition from some state in C_{AVB} to a state in C_{AB}. By the definition of Λ, no smaller number of mutations will be sufficient. This

21

completes the first part of the argument, pertaining one of the directions of the transition: from $C_{A \lor B}$ to C_{AB}. Applying a similar kind of reasoning, the converse transition is seen to require exactly the same number η of mutations. This completes the proof. $\blacksquare$

REFERENCES

Banerjee, A. & J. Weibull (1992): "Evolutionary Selection and Rational Behavior", in Rationality and Economics, A. Kirman & M. Salmon (eds.), Blackwell.

Canning, D. (1992): "Average Behavior in Learning Models", Journal of Economic Theory 57, 442-72.

Foster, D & P. Young (1990): "Stochastic Evolutionary Game Dynamics", Theoretical Population Biology 38, 219-32.

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

Futuyama, D. J. (1979): Evolutionary Biology, Sinauer Associates: Sunderland, Mass.

Kandori, M., G. Mailath, & R. Rob (1993): "Learning, Mutations, and Long Run Equilibrium in Games", Econometrica 61, 29-56.

Kandori, M. & R. Rob (1992): "Evolution of Equilibria in the Long Run: A General Theory and Applications", mimeo, Princeton University and University of Pennsylvania.

Karlin, S. & H.M. Taylor (1975): A First Course in Stochastic Processes, London: Academic Press.

Kimura, M. & T. Otha (1971): Theoretical Aspects of Population Genetics, Princeton University Press, Princeton, NJ.

Robson, A. (1993): "Random Pairing and Fast Convergence to Efficient Equilibria in Coordination Games", mimeo, University of Western Ontario.

Samuelson, L. (1993): "Stochastic Stability in games with Alternative Best Replies", Journal of Economic, forthcoming.

Stahl D. (1993): "Evolution of Smart$_n$ Players", <u>Games</u> <u>and</u> <u>Economic</u> <u>Behavior</u> 5, 604-17.

Vega-Redondo, F. (1993): "On the Evolution of Cooperation in General Games of Common Interest", <u>A</u> <u>Discusión</u> W.P. no. 93-14. University of Alicante.

Weil, P. (1989): "Increasing Returns and Animal Spirits", <u>American</u> <u>Economic</u> <u>Review</u> 79 (4), 889-94.

Young, P. (1993): "The Evolution of Conventions", <u>Econometrica</u> 61, 57-84.

PUBLISHED ISSUES

WP-AD 90-01   "Vector Mappings with Diagonal Images"
C. Herrero, A.Villar. December 1990.

WP-AD 90-02   "Langrangean Conditions for General Optimization Problems with Applications to Consumer Problems"
J.M. Gutierrez, C. Herrero. December 1990.

WP-AD 90-03   "Doubly Implementing the Ratio Correspondence with a 'Natural' Mechanism"
L.C. Corchón, S. Wilkie. December 1990.

WP-AD 90-04   "Monopoly Experimentation"
L. Samuelson, L.S. Mirman, A. Urbano. December 1990.

WP-AD 90-05   "Monopolistic Competition: Equilibrium and Optimality"
L.C. Corchón. December 1990.

WP-AD 91-01   "A Characterization of Acyclic Preferences on Countable Sets"
C. Herrero, B. Subiza. May 1991.

WP-AD 91-02   "First-Best, Second-Best and Principal-Agent Problems"
J. Lopez-Cuñat, J.A. Silva. May 1991.

WP-AD 91-03   "Market Equilibrium with Nonconvex Technologies"
A. Villar. May 1991.

WP-AD 91-04   "A Note on Tax Evasion"
L.C. Corchón. June 1991.

WP-AD 91-05   "Oligopolistic Competition Among Groups"
L.C. Corchón. June 1991.

WP-AD 91-06   "Mixed Pricing in Oligopoly with Consumer Switching Costs"
A.J. Padilla. June 1991.

WP-AD 91-07   "Duopoly Experimentation: Cournot and Bertrand Competition"
M.D. Alepuz, A. Urbano. December 1991.

WP-AD 91-08   "Competition and Culture in the Evolution of Economic Behavior: A Simple Example"
F. Vega-Redondo. December 1991.

WP-AD 91-09   "Fixed Price and Quality Signals"
L.C. Corchón. December 1991.

WP-AD 91-10   "Technological Change and Market Structure: An Evolutionary Approach"
F. Vega-Redondo. December 1991.

WP-AD 91-11   "A 'Classical' General Equilibrium Model"
A. Villar. December 1991.

WP-AD 91-12   "Robust Implementation under Alternative Information Structures"
               L.C. Corchón, I. Ortuño. December 1991.

WP-AD 92-01   "Inspections in Models of Adverse Selection"
               I. Ortuño. May 1992.

WP-AD 92-02   "A Note on the Equal-Loss Principle for Bargaining Problems"
               C. Herrero, M.C. Marco. May 1992.

WP-AD 92-03   "Numerical Representation of Partial Orderings"
               C. Herrero, B. Subiza. July 1992.

WP-AD 92-04   "Differentiability of the Value Function in Stochastic Models"
               A.M. Gallego. July 1992.

WP-AD 92-05   "Individually Rational Equal Loss Principle for Bargaining Problems"
               C. Herrero, M.C. Marco. November 1992.

WP-AD 92-06   "On the Non-Cooperative Foundations of Cooperative Bargaining"
               L.C. Corchón, K. Ritzberger. November 1992.

WP-AD 92-07   "Maximal Elements of Non Necessarily Acyclic Binary Relations"
               J.E. Peris, B. Subiza. December 1992.

WP-AD 92-08   "Non-Bayesian Learning Under Imprecise Perceptions"
               F. Vega-Redondo. December 1992.

WP-AD 92-09   "Distribution of Income and Aggregation of Demand"
               F. Marhuenda. December 1992.

WP-AD 92-10   "Multilevel Evolution in Games"
               J. Canals, F. Vega-Redondo. December 1992.

WP-AD 93-01   "Introspection and Equilibrium Selection in 2x2 Matrix Games"
               G. Olcina, A. Urbano. May 1993.

WP-AD 93-02   "Credible Implementation"
               B. Chakravorti, L. Corchón, S. Wilkie. May 1993.

WP-AD 93-03   "A Characterization of the Extended Claim-Egalitarian Solution"
               M.C. Marco. May 1993.

WP-AD 93-04   "Industrial Dynamics, Path-Dependence and Technological Change"
               F. Vega-Redondo. July 1993.

WP-AD 93-05   "Shaping Long-Run Expectations in Problems of Coordination"
               F. Vega-Redondo. July 1993.

WP-AD 93-06   "On the Generic Impossibility of Truthful Behavior: A Simple Approach"
               C. Beviá, L.C. Corchón. July 1993.

WP-AD 93-07   "Cournot Oligopoly with 'Almost' Identical Convex Costs"
               N.S. Kukushkin. July 1993.

WP-AD 93-08   "Comparative Statics for Market Games: The Strong Concavity Case"
L.C. Corchón. July 1993.

WP-AD 93-09   "Numerical Representation of Acyclic Preferences"
B. Subiza. October 1993.

WP-AD 93-10   "Dual Approaches to Utility"
M. Browning. October 1993.

WP-AD 93-11   "On the Evolution of Cooperation in General Games of Common Interest"
F. Vega-Redondo. December 1993.

WP-AD 93-12   "Divisionalization in Markets with Heterogeneous Goods"
M. González-Maestre. December 1993.

WP-AD 93-13   "Endogenous Reference Points and the Adjusted Proportional Solution for Bargaining Problems with Claims"
C. Herrero. December 1993.

WP-AD 94-01   "Equal Split Guarantee Solution in Economies with Indivisible Goods Consistency and Population Monotonicity"
C. Beviá. March 1994.

WP-AD 94-02   "Expectations, Drift and Volatility in Evolutionary Games"
F. Vega-Redondo. March 1994.

WP-AD 94-03   "Expectations, Institutions and Growth"
F. Vega-Redondo. March 1994.

WP-AD 94-04   "A Demand Function for Pseudotransitive Preferences"
J.E. Peris, B. Subiza. March 1994.