ON THE EVOLUTION OF COOPERATION
IN GENERAL GAMES OF COMMON INTEREST*

Fernando Vega-Redondo**

WP-AD 93-11

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

ON THE EVOLUTION OF COOPERATION
IN GENERAL GAMES OF COMMON INTEREST
Fernando Vega-Redondo

A B S T R A C T

This paper investigates an evolutionary model of equilibrium selection in which agents are randomly paired every period to play some general symmetric game of common interest (i.e., a game where some strategy profile Pareto-dominates all other configurations). Along the process, players tend to imitate the strategy (or strategies) which currently yield the highest average payoff. Occasionally, they also experiment with arbitrarily chosen strategies. In the long run, if the rate of experimentation is small, the population ends up playing the uniquely efficient strategy profile "most of the time". Moreover, the convergence to such state of affairs is relatively fast, independently of population size.

1.- INTRODUCTION

This paper proposes a very simple model for the evolution of cooperation which is characterized by the following features:

(i) A finite population is randomly paired every period to play some (bilateral, finite, and symmetric) game of common interest;[1]

(ii) Players occasionally enjoy the opportunity to revise their strategy;

(iii) When revising their strategies, players imitate that strategy (or strategies) which induced the highest average payoff in the preceding period.

(iv) Quite "infrequently", they also experiment with arbitrary strategies different from those prescribed by (iii).

This context differs from those recently explored in the evolutionary literature on equilibrium selection in the following main respects (see, for example, Kandori, Mailath, and Rob (1993) and Young (1993)):

(a) Players are randomly recombined after <u>every</u> round of play, the particular payoff they receive in each period being a random variable. This contrasts, for example, with the approach of Kandori *et al* (1993) where the payoff accruing to each strategy every period is identified with its average payoff after many encounters. (In this latter case, one relies on the Law of Large Numbers to identify this average payoff with the ex-ante expected payoff.)

(b) In revising their strategies, players imitate the best former strategy rather than play myopic best response to the previous population profile (as in Kandori *et al* (1993)) or to a sample of past actions (as in Young (1993)).

As it turns out, the alternative evolutionary context proposed in this paper has implications quite different from those obtained by most of the previous research in this area. Specifically, the (unique) Pareto efficient

---

[1]   As in Aumann & Sorin (1989), symmetric games of common interest are defined to be those which have a symmetric (equilibrium) profile which Pareto dominates all others. See below for a formal definition.

Being of common interest, the game would seem to induce a quite clearcut prediction: the unique Pareto efficient and symmetric pair of strategies, itself an equilibrium, "should" be played by all agents in every encounter. This, however, presumes some degree of coordination on the part of the players which should not be <u>exogenously</u> taken for granted, at least in an evolutionary framework. Indeed, as stressed by Aumann & Sorin (1989) there might be risk-related considerations which would make it highly unlikely that players, in fact, coordinate on such an equilibrium. Similar considerations also underlie many of the evolutionary models of equilibrium selection existing in the literature, most of which single out the risk-dominant equilibrium as their long-run prediction.[2]

2.2. The Evolutionary Process

Time is modelled discretely. At each period t = 1,2,..., let $\sigma_i(t)$ denote the strategy adopted by player $i \in N$ in t. We shall assume that only pure strategies can be chosen, i.e., $\sigma_i(t) \in S$ for each i and t.

The mechanism by which players adjust their strategies over time consists of two parts: the imitation and the experimentation components of the process.

<u>Imitation</u>

Given the profile $[\sigma_i(t)]_{i \in N}$ of strategies chosen at t, and the particular associated realization of the matching mechanism, there results a corresponding profile of players' payoffs $[\pi_i(t)]_{i \in N}$. Ex-ante, of course, this profile is a random variable. However, once materialized, it induces a corresponding vector of average payoffs for each of the strategies played by

---

[2] See, for example, the papers by Kandori, Mailath, & Rob (1993) and Young (1993) mentioned above. Two exceptions to this state of affairs may be found in Binmore & Samuelson (1993) and Canals & Vega-Redondo (1993) for 2×2 games.

In the first paper, a Pareto efficient equilibrium may still be selected over the risk-dominant one if higher differential payoffs within the basin of attraction of the former equilibrium offsets the effect of a larger basin of attraction for the latter. Thus, even in this case, Pareto efficient equilibria are not always selected.

An unambiguous selection of the efficient equilibrium occurs in the second paper. In it, however, the framework is quite different from the present one in that hierarchical, multilevel selection operates along the process.

Experimentation

As usual in the recent evolutionary literature, the above described selection process is complemented with some notion of experimentation. This involves a small and independent probability that, in every period, each player $i \in N$ adjusts her strategy in some arbitrary direction. More specifically, each player will be assumed to adopt, with some small probability ε every period, some strategy different from the one prescribed by imitation. For simplicity,[3] it will be supposed that all other strategies have an equal probability of being adopted through this mechanism. In other words, experimentation (often labelled "mutation") is unbiased.

3.- ANALYSIS

The process described in the previous section can be formalized as an stochastic process. Its states consist of the specification of a tuple of the form $[(\sigma_i)_{i=1,2,\ldots,n} , (\rho_q)_{q=1,2,\ldots,m}]$, where each σ_i is identified with the current strategy choice of agent i, and each ρ_q represents the average payoff of strategy s_q induced by the current particular realization of the matching mechanism. Of course, not all such tuples are possible states of the process. Given a particular strategy profile $\sigma \equiv (\sigma_i)_{i=1,2,\ldots,n}$, only those vectors of average payoffs $\rho \equiv (\rho_q)_{q=1,2,\ldots,m}$ which are consistent with some possible realization of the random matching mechanism are well-defined states. The set of all of them will be denoted by Ω.

Due to experimentation, the stochastic process described is non-periodic and positively recurrent, i.e., there is positive probability of visiting every state from every other state. Therefore, the process is ergodic, having a unique invariant probability distribution that summarizes the long-run behavior of the system, independently of initial conditions. Let $\mu(\theta,\varepsilon)$ denote such distribution. As advanced, given θ, we shall be interested in the *limit stationary distribution* μ^*, defined as follows:

---

[3] Strictly, it is only required that the experimentation probability have full support over the full set of alternative strategies.

<u>Remark</u> <u>1</u>: *Fast convergence to the limit invariant distribution.*

As emphasized by some authors (see, for example, Ellison (1991)) the practical relevance of the long-run convergence results of evolutionary models hinges upon the speed of convergence one may expect for the realizations of the stochastic process. If such expected convergence is slow, it will generally be necessary to "wait for a long time" before the realized dynamics of the process approximates its theoretical limit distribution. In view of these considerations, it is quite important to evaluate the speed of convergence of our contemplated stochastic process and compare it to that of other models explored in this literature. For the sake of concreteness, I will focus on the model analyzed by Kandori, Mailath, & Rob (1993).

Rather than addressing formally issues of speeds of convergence, it is perhaps more transparent to illustrate matters by focusing on the related concept of "waiting time". More specifically, we shall concern ourselves with the maximum expected waiting time for observing a shift to the predicted long-run equilibrium from any other state of the process (i.e., the inverse of the minimum transition probability). In the context studied by Kandori *et al*, such expected waiting time is of the order $\varepsilon^{-\lambda n}$, where $\lambda \in (0,1)$ is a parameter associated to the data of the game being played. Thus, for large populations (i.e., large n) and small ε, the corresponding waiting time will be very large.

However, in our context, it is easy to check (see, for example, Claim 2 in the Appendix) that such expected waiting time will be of the order ε^{-2}, independently of the population size. Thus, even for large populations, the convergence towards the long-run prediction of the model will in average be relatively fast, as compared with the model proposed by Kandori *et al*.

<u>Remark</u> <u>2</u>: *Extensions to general games.*

The validity of our Theorem crucially depends on its restriction to common-interest games. For example, not even if all equilibria of a symmetric game are both pure and symmetric (i.e., lie in the "diagonal") can one assure that efficiency will be the unique long-run behavior of the process.

11

adjustment. In particular, it would also yield identical conclusions as to the preeminence of risk-dominance considerations in processes of long-run equilibrium selection for 2×2-games. In order for our logic to apply in selecting for the efficient equilibrium, strategy adjustment has to be sufficiently fast relative to the operation of the matching mechanism. Specifically, it cannot be so slow as to eliminate the randomness in payoffs induced by this mechanism.

4.- SUMMARY

This paper has proposed a new evolutionary model for equilibrium selection in games which:

(i) contemplates imitation rather than best response as the basis of strategy selection, and

(ii) assumes that this process of strategy selection takes place after each round of play, on the basis of the particular outcome produced by the matching mechanism.

The combination of these two features of the model, together with the usual postulate of infrequent experimentation, has been shown to yield the long-run selection of the unique efficient equilibrium in games of common interest. We have also argued that the convergence of the process to such long-run outcome is relatively fast (and independent of population size), as compared to some other evolutionary models studied in the literature.

concept of limit stationary distribution) is well defined. As ε converges to zero, the evaluation of (5) hinges upon the rate $\mathcal{O}\left(q_\omega(\varepsilon)\right)$ at which, for each $\omega \in \Omega$, $q_\omega(\varepsilon)$ converges to zero. In order to determine these rates of convergence we shall rely on the following concepts.

Firstly, given any state $\omega \in \Omega$, denote by $T(\omega) \subset \Omega$ the set of states which may succeed ω with positive probability when <u>no</u> agent experiments. (Of course, if experimentation were considered, this set would coincide with the whole state space.) Next, for any two states $\omega \equiv [(\sigma_i),(\rho_q)]$ and $\omega' \equiv [(\sigma_i'),(\rho_q')]$, define a distance function $d(\cdot)$ on Ω by:

$$d(\omega,\omega') = \operatorname{card} \{i \in N: \sigma_i \neq \sigma_i'\}. \tag{6}$$

This function simply counts the number of agents who play a different strategy in each of the two states.

Combining the two former constructs, we introduce the following central concept. Given any two states, ω and ω', define

$$c(\omega,\omega') = \min \{ d(\tilde{\omega},\omega'): \tilde{\omega} \in T(\omega) \} \tag{7}$$

as the *cost of transition* from ω to ω'. This cost represents the minimum number of experimentations needed to complement the imitation component of the dynamics in order to complete the transition from ω to ω' with positive probability. (Naturally, such cost of transition will be equal to zero if the imitation dynamics is alone sufficient to carry out the contemplated transition with positive probability.)

Correspondingly, we may define the *cost of a tree* H (or, analogously, of a path h) as follows:

$$c(H) = \sum_{(\omega,\omega') \in H} c(\omega,\omega'). \tag{8}$$

When $\varepsilon \longrightarrow 0$, it follows from (4) that, for any $\omega \in \Omega$, we may compute the order of convergence of its corresponding $q_\omega(\varepsilon)$ as follows:

15

Associated to state $\tilde{\omega}$, we shall also need to consider the (sub)path of h, $\check{h} = \{(\hat{\omega},\omega^1),(\omega^1,\omega^2),...,\ (\omega^{k-1},\tilde{\omega})\}$, which links $\hat{\omega}$ and $\tilde{\omega}$. Of course, $\tilde{\omega}$ may coincide with ω, or even with $\hat{\omega}$. In the former case, $\check{h} = h$; in the latter, $\check{h}$ will be identified with the "empty path", rendering redundant part of the ensuing argument.

Clearly, state $\hat{\omega}$ and its predecessors in H induce a corresponding $\hat{\omega}$-tree restricted to this set. (Here, we generalize the concept introduced above in order to allow for trees whose vertices span a _proper_ subset of Ω.) Represent such tree by $\tau(\hat{\omega},H)$, where the general notation $\tau(\omega',H')$ stands for the "induced tree" obtained from the tree H' by selecting the new root $\omega' \in H'$. Using this notation, we now define a series of induced (sub)trees as follows:

$$
\begin{aligned}
\tau^0 &= \tau(\hat{\omega},H), & H^0 &= H\backslash\tau^0 ; \\[4pt]
\tau^1 &= \tau(\omega^1,H^0), & H^1 &= H^0\backslash\tau^1; \\[4pt]
\tau^2 &= \tau(\omega^2,H^1), & H^2 &= H^1\backslash\tau^2; \\
&\;\;\vdots & &\;\;\vdots \\
\tau^{k-1} &= \tau(\tilde{\omega},H^{k-2}), & H^{k-1} &= H^{k-2}\backslash\tau^{k-2}; \\[4pt]
\tau^k &= \tau(\omega,H^{k-1}), & H^k &= H^{k-1}\backslash\tau^k = \varnothing
\end{aligned}
\tag{12}
$$

where recall that ω^1, ω^2,..., ω^{k-1} stand for the states in the above path $\check{h}$ which links $\hat{\omega}$ and $\tilde{\omega}$.

Note that, as reflected by (12) (in particular, by the fact that $H^k = \varnothing$), the trees in the set $T \equiv \{\tau^0,\tau^1,...,\tau^k\}$ induce a partition of the full set of states Ω. We next need to verify the following claim on their "aggregate cost".

Claim 1: $\sum\limits_{i=0}^{k} c(\tau^i) \equiv c^* \leq c(H)\text{-}1$. _Moreover, if_ $\omega \notin B_{\hat{\omega}}$ _there exists some_ $\underline{n} \in \mathbb{N}$ _such that, if_ $n \geq \underline{n}$, _then_ $c^* \leq c(H)\text{-}3$.

Proof: If $\omega \neq \tilde{\omega}$, denote by $\check{h}^+$ the sub-path of h which links $\hat{\omega}$ to the immediate successor of $\tilde{\omega}$. Otherwise (i.e, if if $\omega = \tilde{\omega}$), $\check{h}^+$ is identified with $\check{h}$, as defined above.

17

previous ω^{1j} in the tree $\hat{\tau}$ (i.e, on the second level of the tree $\hat{\tau}$).[5] Proceeding in this fashion for as many stages as necessary, we shall eventually come to one of the following two configurations:

(i) If $\omega \in B_{\hat{\omega}}$, only one single $\hat{\omega}$-tree will remain spanning the whole Ω;

(ii) If $\omega \notin B_{\hat{\omega}}$, two different trees will remain: an $\hat{\omega}$-tree and an ω-tree, jointly spanning the whole state space Ω. Moreover, the former tree includes all states in $B_{\hat{\omega}}$.

Note that the above described procedure has not added any links between states which could increase the total original cost c^* induced by the trees in T. Thus, if (i) applies, the desired conclusion follows immediately from Claim 1. If instead (ii) applies, this conclusion follows from the following final claim.

Claim 2: For all $\omega' \in \Omega$, $\omega' \neq \hat{\omega}$, $\exists \check{\omega} \in B_{\hat{\omega}} : c(\omega', \check{\omega}) \leq 2$.

Proof: Consider any strategy profile $(\sigma_i)_{i=1,2,...,n}$ where some positive even number of agents play strategy $s_{\hat{q}}$. Associated to the particular outcome of the matching process where all of these latter players are matched among themselves, there is a feasible state compatible with such strategy profile where the average payoff of strategy $s_{\hat{q}}$, $\rho_{\hat{q}}$, equals $a_{\hat{q}\hat{q}}$. Let $\check{\Omega}$ be the set of all such states. Clearly, $\check{\Omega} \subset B_{\hat{\omega}}$. Moreover, for any state $\omega' \in \Omega$, there is always some state $\check{\omega} \in \check{\Omega}$ such that, after no more than two agents experiment with $s_{\hat{q}}$, there is positive probability (independent of ε) that the process makes the transition from ω' to $\check{\omega}$. Thus, $c(\omega', \check{\omega}) \leq 2$, establishing the desired conclusion. ∎

By the claim just proven, the ω-tree in (ii) can be connected to some state $\check{\omega}$ in the complementary $\hat{\omega}$-tree such that a full $\hat{\omega}$-tree arises (i.e., one spanning all states in Ω) whose cost is no larger than c^*+2. Recall that, by claim 1, $c^*+2 < c(H)$ when $\omega \notin B_{\hat{\omega}}$. Therefore, such consolidated $\hat{\omega}$-tree has a cost strictly lower than the original ω-tree H. This establishes the Lemma, thus completing the proof of the Theorem.

---

5 The procedure is operated sequentially across "levels" of the tree τ in order to prevent that, at some stage, a certain state ω and its predecessors could be required linked to one of the predecessors of ω itself.

R E F E R E N C E S

Aumann, R. & S. Sorin (1989): "Cooperation and Bounded Recall", *Games and Economic Behavior*, 1, 5-39.

Binmore, K. & L. Samuelson (1993): "Muddling Through: Noisy Equilibrium Selection", University College & University of Wisconsin, mimeo.

Canals J. & F. Vega-Redondo (1993): "Multilevel Evolution in Games", A Discusión, Working Paper no. 93-10, University of Alicante.

Ellison, G. (1991): "Learning, Local Interaction, and Coordination", Massachusetts Institute of Technology, mimeo.

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

Kandori, M., G. Mailath, & R. Rob (1993): "learning, Mutation, and long-Run Equilibrium Selection in Games", *Econometrica*, 61, 29-56.

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

PUBLISHED ISSUES

33   "Consequences of the Manipulation of Lindahl Correspondence: An Example"
            C. Bevía, J.V. LLinares, V. Romero, T. Rubio. 1990.

34   "Egalitarian Allocations in the Presence of Consumption Externalities"
            C. Herrero, A. Villar. 1990.

SECOND PERIOD

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 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. December 1993.