Industriens Utredningsinstitut

THE INDUSTRIAL INSTITUTE FOR ECONOMIC AND SOCIAL RESEARCH

A list of Working Papers on the last pages

No. 345, 1992

EVOLUTION AND RATIONALITY: SOME RECENT GAME–THEORETIC RESULTS

by

Abhijit Banerjee and Jörgen W. Weibull

Paper presented at the Tenth World Congress of the International Economic Association, Moscow, August 24–28, 1992.

October 1992

EVOLUTION AND RATIONALITY:

SOME RECENT GAME–THEORETIC RESULTS*

Abhijit Banerjee   and   Jörgen W. Weibull

Department of Economics
Harvard University
Cambridge, MA. 02138
U.S.A.

Department of Economics, and Institute
for International Economic Studies
Stockholm University
106 91 Stockholm
Sweden

First draft: 15 March 1992. Current version: 15 October 1992.

ABSTRACT

This paper surveys the recent literature on evolutionary game theory from the viewpoint of economic theory. The aim is to understand to what extent evolutionary arguments can substitute rationality–based arguments as a foundation for Nash equilibrium and other non–cooperative solution concepts. We conclude that there is an intimate connection between the attractors of evolutionary processes and Nash equilibrium even in environments which are much more general than the simple setting of the standard biological model. Furthermore, even when the evolutionary process does not converge, it's long–run evolution will, in continuous time, necessarily eliminate dominated and iteratively dominated strategies. However, we also find that the current evolutionary set–up needs substantial generalization if it is to incorporate plausible adaptive behavior by intelligent but boundedly rational players. We suggest some directions for such generalizations and show by example that in more general environments the connection between evolutionary outcomes and Nash equilibrium may be weaker.

I. INTRODUCTION

A large part of economics, and economic theory in particular, relies on such solution concepts as Nash equilibrium and its refinements. Unfortunately, it is difficult to provide a solid theoretical or empirical justification for Nash equilibrium behavior. "Rationality," or even "common knowledge of rationality," is not enough to generate such behavior. Among other things, one also needs to assume that the players coordinate their beliefs about each others' actions (Armbruster and Boege, 1979; Johansen, 1982; Bernheim, 1984; Pearce, 1984; Bernheim, 1986; Binmore, 1987; Aumann, 1987; Tan and Werlang, 1988; Aumann and Brandenburger, 1991). Moreover, in games with a dynamic structure, the very notion of rationality becomes problematic, and common knowledge of rationality may even lead to logical contradictions (Rosenthal, 1981; Binmore, 1987; Bicchieri, 1989; Basu, 1988, 1990).

A completely different approach is to leave the epistemology of rationality aside and instead ask whether or not economic agents, and human decision–makers in general, behave *as if* they met the stringent rationality and coordination conditions inherent in Nash equilibrium and other non–cooperative solution concepts. One process which may justify this "*as if*" approach is "natural" selection.[1] More precisely, one may ask whether evolutionary selection among more or less boundedly "rational" behaviors in strategic interaction situations leads to (aggregate and/or long–run) Nash equilibrium behavior. One can then think of "players" as "hosts" of a variety of "competing" behaviors — including potentially "rational" behaviors — and ask which of these behaviors survive in an evolutionary selection process (Binmore, 1988). If the behaviors selected by evolutionary processes result in a Nash equilibrium outcome, then one can claim that, whether or not players are genuinely "rational" and "coordinate their

---

[1] For pioneering suggestions in this direction, in the context of competitive markets, see Alchian (1950), Friedman (1953) and Winter (1964, 1971).

beliefs," in the long run they will behave as if they did indeed meet the underlying rationalistic assumptions (see Selten (1991b) for a discussion of some relevant paradigms).

The most promising methodology for such an evolutionary justification of Nash equilibrium play is provided by recent advances in evolutionary game theory, a paradigm pioneered by the British biologist John Maynard Smith. This paradigm has been developed along two lines, one static approach using evolutionary stability as its key concept (Maynard Smith and Price, 1973; Maynard Smith, 1974, 1982), and one dynamic approach based on explicit Darwinian (or Malthusian) selection of behaviors. In the dynamic approach, one imagines pairwise and randomly matched interactions in a large population of individuals (Taylor and Jonker, 1978; Hofbauer et al., 1979; Zeeman, 1980, 1981; Schuster et al., 1981a; Schuster and Sigmund, 1983; Bomze, 1986, Hofbauer and Sigmund, 1988). The players are confined to a limited menu of particularly simple behaviors, viz. those corresponding to always playing the same pure strategy in the two—player game in question. Hence, one may think of the players as being "programmed" to always play a certain (more or less involved) strategy when encountering another individual. In the biologists' models, payoffs represent reproductive fitness, usually measured as (expected) number of offspring, and the offspring inherits its (single) parent's behavior. Hence, the growth rate of a population share of individuals programmed to a certain behavior is proportional to its payoff in the current population composition. Using a term coined by the British socio—biologist Richard Dawkins, this reproductive selection dynamics is usually called the *replicator dynamics*.[2] The focus of the biologists' analysis is on population states which are attractors in this dynamics, such states being the natural candidates for long—run

---

[2] Dawkins used the term "replicator" for entities which can get copied and which are such that (a) their properties can affect their probability of being copied, and (b) the line of descent copies is potentially unlimited.

aggreagate population behavior.

The relevance of this literature for the social sciences in general, and economics in particular, is limited in several ways. First, in economics, the payoffs in a game do usually not represent reproductive fitness, but rather firm's profits or households' utility. Moreover, in most such applications, the selection mechanism is not biological but the aggregate result of more or less conscious choices made by the individuals. Accordingly, economists have generalized some of the analysis to fairly wide classes of selection dynamics containing the replicator dynamics as well as a variety of more or less rational individual adjustment processes (Nachbar, 1990; Friedman, 1991; Kandori et al., 1991; Samuelson and Zhang, 1992; Swinkels, 1992). A related limitation in the standard set—up of evolutionary game theory is its focus on behaviors which are not conditioned on any information that a player may have, such as the current aggregate population behavior or the type of the opponent player. One way to relax this rigidity, indicated above, is to let players over time revise their choice of strategy in the light of the current population state (Banerjee and Weibull, 1991; Dekel and Scotchmer, 1992). An alternative way is to expand the menu of behaviors available at each moment (Robson, 1990; Banerjee and Weibull, 1991; Stahl, 1992; Banerjee and Weibull, 1992). There is also no a priori reason why, in all relevant situations, all the players should be modelled as coming from the same population. Indeed, a number of papers treat interactions between distinct populations, occupying different player roles in the game in question (Schuster et al., 1981b; Schuster and Sigmund, 1985; Hofbauer and Sigmund, 1988; Friedman, 1991; Ritzberger and Vogelsberger, 1991; Samuelson and Zhang, 1992; Swinkels, 1992).

The present survey discusses a number of papers on explicitly dynamic models of evolutionary selection of behaviors in games. This literature is currently growing fast and is closely related to at least three strands in economic theory, none of which is discussed here. The obviously most closely related of these strands is the important

literature on the static criterion of evolutionary stability.[3] The two other omitted fields are learning models and models in which decision makers are represented as automata.[4] Moreover, we are aware that we have not even covered all current and important contributions to explicitly dynamic evolutionary analyses. In particular, we do not discuss multi–population dynamics. We feel that many important aspects of the relationship between rationality and evolution — the focus of the present survey — come out already in symmetric two–person games. For the same reason, we only briefly touch upon explicitly stochastic dynamic models.

The material is organized as follows. Basic notation and elements of non–cooperative game theory are given in Section II. The explicitly dynamic version of standard evolutionary game theory, in particular its implications for "rational" and "coordinated" behavior, is discussed in Section III. Section IV discusses some extensions to a fairly wide class of dynamics of relevance for models of boundedly rational individuals who revise their strategy choice over time. Section V discusses some models of evolutionary selection among more general boundedly rational behaviors. Our main conclusions are summarized in Section VI.

II.   NOTATION AND PRELIMINARIES

The analysis in the present paper is restricted to finite and symmetric two–player games in normal form. Let $I=\{1,2,...,k\}$ be the set of *pure* strategies. Accordingly, a *mixed* strategy is a point x on the unit simplex $\Delta=\{x\in\mathbb{R}^k_+: \Sigma_i x_i=1\}$ in k–dimensional Euclidean

---

[3] Every evolutionary stable strategy $x\in\Delta$ is asymptotically stable (see footnote 25) in the replicator dynamics (Hofbauer et al., 1979), so the connection between the static and dynamic approaches is fairly tight.

[4] Unfortunately, a discussion of connections between these three research paradigms falls outside the scope of present paper. Two relevant papers are Blume and Easley (1991), for a connection between learning behaviors and evolution, and Binmore and Samuelson (1992), for the connection between automata and evolution.

space. The *support* of a mixed strategy $x \in \Delta$ is the subset $C(x) = \{i \in I : x_i > 0\}$ of pure strategies which are assigned positive probabilities. A strategy x is called *interior* (or *completely mixed*) if $C(x) = I$.

Let a_{ij} be the *payoff* of strategy $i \in I$ when played against strategy $j \in I$, and let A be the associated k×k payoff matrix. Accordingly, the (expected) payoff of a mixed strategy x, when played against a mixed strategy y, is $u(x,y) = x \cdot Ay = \Sigma_{ij} x_i a_{ij} y_j$. The payoff function $u : \Delta^2 \rightarrow \mathbb{R}$ is clearly bi–linear, and the payoff of a pure strategy $i \in I$, when played against a mixed strategy y, is $u(e^i, y)$, where e^i is the i'th unit vector in $\mathbb{R}^k$, etc. We will frequently identify a pure strategy $i \in I$ with its mixed–strategy counterpart $e^i \in \Delta$.

A pure strategy $i \in I$ is *weakly dominated* if there exists a strategy $x \in \Delta$ which never earns a lower payoff and sometimes a higher payoff (i.e., $u(x,y) \geq u(e^i, y)$ $\forall y \in \Delta$, with strict inequality for some y). A pure strategy $i \in I$ is *strictly dominated* if there exists a strategy $x \in \Delta$ which always earns a higher payoff (i.e., $u(x,y) > u(e^i, y)$ $\forall y \in \Delta$). A *best reply* to a strategy $y \in \Delta$ is a strategy $x \in \Delta$ such that $u(x,y) \geq u(x',y)$ $\forall x' \in \Delta$. For each $y \in \Delta$, let $\beta(y) \subset \Delta$ be its set of best replies. A *Nash equilibrium* is a pair (x,y) of mutually best replies, a Nash equilibrium is *strict* if each strategy is the unique best reply to the other, and a Nash equilibrium (x,y) is *symmetric* if x=y. By Kakutani's Fixed Point Theorem, every finite and symmetric game has at least one symmetric Nash equilibrium.

One solution concept which is weaker than Nash equilibrium is iterative strict dominance. A pure strategy $i \in I$ is said to be *iteratively strictly undominated* if it is not strictly dominated in the original game G, nor in the game G' obtained from G by removal of all strictly dominated strategies, nor in the game G" obtained from G' by removal of all strategies which are strictly dominated in G', etc. A related, but different, solution concept is rationalizability. A pure strategy $i \in I$ is *never a best reply* if there exists no mixed strategy $y \in \Delta$ against which $i \in I$ is a best reply, and $i \in I$ is

rationalizable (Bernheim, 1984; Pearce, 1984) if it is not a "never best reply" in the original game G, nor in the game G' obtained from G by removal of all "never best replies," nor in the game G" obtained from G' by removal of all "never best replies," etc. Each of these two methods of iterated elimination of pure strategies stops in a finite number of steps. Pearce (1984) has shown that, while the two remaining sets may differ in games with more than two players, they coincide in all two–player games.[5] Hence, in the present setting, a strategy i∈I is rationalizable if and only if it is iteratively undominated.

III. NASH EQUILIBRIUM AS THE RESULT OF EVOLUTIONARY ADAPTATION

In the basic dynamic setting of evolutionary game theory, one imagines pairwise and randomly matched interactions in a large population of individuals, each interaction taking the form of play of a symmetric and finite two–player game. As in Section II, let I denote the strategy set, A the payoff matrix, and $u:\Delta^2 \to \mathbb{R}$ the associated payoff function. Interactions take place contiuously over time, the current *aggregate behavior* of the population, or its *state*, is a vector $x=(x_1,...,x_k)\in\Delta$, where each component x_i is the population share currently using strategy i∈I. Accordingly, it is immaterial for a player whether he plays against a player using a mixed strategy $x\in\Delta$ or against a randomly drawn individual from a population of pure–strategy players in state $x\in\Delta$. In both cases, the (expected) payoff of a pure strategy i∈I is $u(e^i,x)$. In the population setting, the *average payoff* in the population is simply $u(x,x) = \Sigma x_i u(e^i,x)$. In the biologists' models, payoffs represent "fitness", which is usually taken to be the

---

[5] A strictly dominated strategy is never a best reply, and hence the set of rationalizable strategies is always a subset of the set of strategies surviving the iterated elimination of strictly dominated strategies.

(expected) number of offspring, and each offspring inherits its (single) parent's strategy, so (implicitly presuming the law of large numbers) the *replicator dynamics* in continuous time is

$$\dot{x}_i = u(e^i - x, x) \cdot x_i \qquad [\forall i \in I, \forall x \in \Delta], \qquad (1)$$

with dots for time derivatives.[6]

In other words, the growth rate $\dot{x}_i / x_i$ of each sub–population $i \in I$ equals the difference between its current payoff and the current population average. Evidently, a sub–population whose strategy $i \in I$ is a best reply to the current aggregate behavior $x \in \Delta$ has the highest growth rate in this dynamics, but also other sub–populations may grow, viz. those who use strategies which do better than average.

In the special case $k=2$, i.e. a (symmetric) 2×2 game, it is easily established that this implies convergence to some (symmetric) Nash equilibrium, from any interior initial state. More exactly, if initially all strategies in the game are used by some individuals, i.e. $x(0)$ is interior, then the state $x(t)$ converges over time to some state $x^* \in \Delta$ such that (x^*, x^*) constitutes a Nash equilibrium (Hofbauer and Sigmund, 1988). (In certain cases, the limit state x^* depends on the initial state $x(0)$.) Hence, the evolutionary approach lends strong support to the Nash equilibrium paradigm in such low–dimensional settings.

Unfortunately, for general symmetric $k \times k$ games the connection between evolutionary dynamics and Nash equilibrium is somewhat less tight. The problem is that the evolutionary selection process in higher–dimensional spaces need not converge.

---

[6] This system of ordinary differential equations is polynomial. Hence, it has a unique solution through every intial state $x(0) \in \Delta$. Moreover, it leaves the unit simplex Δ invariant, $x(0) \in \Delta \Rightarrow x(t) \in \Delta \ \forall t \geq 0$, and the system (1) is unaffected by affine trans–formations of payoffs, modulo a change of time scale. For an analysis of the validity of the implicit use of the law of large numbers, see Boylan (1992).

However, just as in the 2×2 case any population state which is (Lyapunov) stable in the replicator dynamics corresponds to a symmetric Nash equilibrium, and a strategy which is in Nash equilibrium with itself is stationary in this dynamics.[7] In other words, if evolutionary selection induces no movement in the composition of the population's aggregate behavior, and that behavior is dynamically robust with respect to small perturbations, then it is compatible with the stringent rationality and coordination hypotheses in the rationalistic justification of Nash equilibrium behavior:

Proposition A (Bomze, 1986): If the population state x is (Lyapunov) stable in the replicator dynamics (1), then (x,x) is a Nash equilibrium. If (x,x) is a Nash equilibrium, then x is stationary in the replicator dynamics.

The intuition behind these claims is fairly straight–forward. The second claim is simplest to see. For if (x,x) is a Nash equilibrium, and hence x∈Δ is a best reply to itself, then all pure strategies in its support C(x) earn the same (maximal) payoff, which also is the average payoff, since strategies not in C(x) are absent from the population. Thus each sub–population is either extinct or earns the average payoff, establishing $\dot{x}=0$ in (1). The intuition for the stability claim is that if (x,x) is not a Nash equilibrium, and hence x is not a best reply to itself, then there exist some pure strategy i which earns more against x than some pure strategy j in the support of x. Hence, along all solution paths of (1) in a neighbourhood of x, sub–population j decreases towards zero over time, implying that the population state leaves the neighborhood (since by hypothesis j∈C(x), and hence $x_j>0$). Consequently, x then is not stable. In fact, by similar arguments one can show that, just as in the 2×2 case, if the population state

---

[7] A state is called *stationary* if, starting in x, the system remains at x. A state x∈Δ is said to be (*Lyapunov*) *stable* if there for every neighborhood V⊂Δ of x exists some neighborhood U⊂Δ of x such that, starting in U, x(t) never leaves V, see e.g. Hirsch and Smale (1974).

converges from an initial state in which all strategies are used, then the limiting state has to be Nash equilibrium behavior:[8]

> **Proposition B (Nachbar, 1990):** If an interior dynamic path in the replicator dynamics (1) converges to some $x \in \Delta$, then (x,x) is a Nash equilibrium.

This simple observation implies that every strictly dominated strategy is wiped out from the population, granted all strategies are represented in the initial population and that the induced dynamic path converges. But what if aggregate behavior does not converge? It turns out that all strictly dominated strategies are nevertheless wiped out. In fact, as shown by Samuelson and Zhang (1992), only strategies which are rationalizable can survive evolutionary selection, given all pure strategies are initially present in the population:[9]

> **Proposition C (Samuelson and Zhang, 1992):** If a pure strategy is not rationalizable, then its population share converges to zero along any interior dynamic path in the replicator dynamics (1).

Consequently, even if the evolutionary selection process fails to converge, in the long run virtually no individual will behave irrationally in the sense of playing strategies which are never best replies, nor will they play strategies which are never best replies when others avoid strategies which are never best replies etc. In fact, since all non–rationalizable strategies are wiped out in the long run, all players eventually behave

---

[8] If the initial state $x(0)$ is interior, then so are all future states $x(t)$. In other words, the interior of the unit simplex Δ is invariant in the replicator dynamics (1), and so one may unambiguously speak about interior (solution) *paths*.

[9] Samuelson and Zhang establish this for interactions between two populations, implying the present claim as a corollary. For a direct proof see Weibull (1992).

almost as if rationality (in the sense of best–reply behavior) were common knowledge. What is missing from the kind of behavior presumed in Nash equilibrium play is the coordination of beliefs.[10]

However, the intuition for this result is not so straight–forward. Indeed, by way of a cleverly construed counter–example, Dekel and Scotchmer (1992) show that the result is *not* valid in the standard discrete–time version of the replicator dynamics,

$$x_i(t+1) = \frac{u[e^i, x(t)]}{u[x(t), x(t)]} \cdot x_i(t) .$$

(2)

In this version, each generation lives for one period, and all individuals reproduce simultaneously. If an individual using strategy $i\epsilon I$ has $u[e^i, x(t)]$ offspring when the population state is $x(t)$, and all offspring inherit their "parent's" strategy, then (2) results.[11] It turns out that a strategy which is strictly dominated by a mixed strategy, but not by any pure strategy, may, along non–convergent solution paths to (2), periodically do sufficiently well to avoid extinction.

The example in Dekel and Scotchmer (1992) is the special case $\alpha=0.35$ and $\beta=0.1$ of the following extension of the so called "Rock–Scissors–Paper" game:

$$A = \begin{bmatrix} 1 & 2+\alpha & 0 & \beta \\ 0 & 1 & 2+\alpha & \beta \\ 2+\alpha & 0 & 1 & \beta \\ 1+\beta & 1+\beta & 1+\beta & 0 \end{bmatrix} .$$

---

[10] Even in certain non–convergent cases does Nash equilibrium appear; Schuster et al. (1981a) show that if a game has a unique interior Nash equilibrium $p\epsilon\Delta$, then p is also the time average of every periodic dynamic path in Δ (whenever such exist).

[11] Unlike the continuous–time version of the replicator dynamics, (2) requires all payoffs to be positive. With this restriction, the dynamics (2) leaves the unit simplex, as well as its interior, invariant. However, while the continuous–time version (1) is invariant under affine transformations of payoffs, (2) is not. Indeed, addition of a positive constant γ to all payoffs in (2) does affect its solution paths in Δ (Hofbauer and Sigmund, 1988).

It is easily verified that the fourth strategy is strictly dominated by the unique Nash equilibrium strategy $m=(\frac{1}{3},\frac{1}{3},\frac{1}{3},0)\in\Delta$ if $\alpha>3\beta>0$, a condition met by Dekel's and Scotchmer's numerical example. Moreover, one can show that near the boundary face $x_4=0$ of the unit simplex Δ, i.e., in states where only a small fraction of the population uses the dominated strategy 4, this sub–population grows whenever aggregate behavior $x\in\Delta$ is at some distance from m. In fact, when $\alpha<4\beta$, as in Dekel's and Scotchmer's example, x_4 grows outside a circular disk as shown in Figure 1.[12] Hence, if, along a dynamic path near this face, the population state drifts along the edges, then the population share x_4 grows. In the continuous–time replicator dynamics (1) this does not happen; for the system converges towards m from any initial state on (the relative interior of) that face. However, in the discrete time dynamics (2), the system diverges on the face $x_4=0$ towards the three edges of that face. Dekel and Scotchmer (1992) prove that $x_4(t)$ converges to zero in the dynamics (2) if and only if initially all three undominated strategies appear in identical shares.


Figure 1

---

[12] Near the face $x_4=0$, $\dot{x}_4>0$ iff $x_1^2+x_2^2+x_3^2 > 1-2\beta/\alpha$, see Weibull (1992) for details.

However, the discrete–time version (2) does not seem entirely compelling for general evolutionary analyses since it assumes that the whole population reproduces at the same time. As a first approximation, it appears more natural to assume that agents reproduce continuously, or else discretely but in smaller batches, at a more or less uniform rate over time. Indeed, Cabrales and Sobel (1992) show that if one uses discrete time but lets only a small fraction, λ, of the whole population reproduce each time, then the discrete–time dynamics differs from (2) and in fact becomes more and more similar to the continuous–time dynamics as λ decreases. In this fashion, they establish the validity of Proposition C also for discrete–time versions of the replicator dynamics with λ sufficiently small. As a result, the anomaly raised by Dekel and Scotchmer is not too damaging to the qualitative results discussed above.

In the current formulation, the population share of each strategy changes in a deterministic way. A more general formulation would allow for a stochastic evolution of the population shares. A pioneering contribution in this vein is Foster and Young (1990), who add white noise to the (continuous–time) replicator dynamics (1). In this fashion they obtain an ergodic stochastic process with a unique limiting distribution. (Hence, unlike the deterministic replicator dynamics, its long–run behavior is independent of its initial state.) In this framework, and using powerful analytical techniques due to Freidlin and Wentzell (1984), they are able to derive interesting limiting results for the case when the white noise terms is reduced towards zero, for a number of numerical examples.[13]

---

[13] See Fudenberg and Harris (1992) for an alternative stochastic specification.

IV.   NASH EQUILIBRIUM AS THE RESULT OF BOUNDEDLY RATIONAL ADAPTATION

In applications to economics and other social sciences, payoffs usually represent utility or profits, and, moreover, one is interested in the dynamics of cognitive and social adaption processes, such as learning or imitation of successful behavior, rather than in biological reproduction as such.[14]   Hence, we now ask whether the results for the replicator dynamics carry over to such dynamics.   In fact, economists have recently studied properties of certain classes of evolutionary dynamics which are intended to be wide enough to contain a variety of plausible social evolutionary processes.   All of these dynamics are, in one way or another, monotone with respect to payoff differences between current strategies (Nachbar, 1990; Friedman, 1991; Samuelson and Zhang, 1992; Swinkels, 1992).

Here, we formalize one particular such class, viz. dynamics with the property that if some pure strategy earns more than another, then the first sub–population grows at a higher rate than the second.   Clearly the replicator dynamics belongs to this class, but note the new interpretations that such a generalization admits; one may now think in terms of infinitely lived, boundedly rational individuals who consciously choose their strategy, or, more precisely, revise their choice of strategy over time.   The dynamics will simply represent the aggregate effect of the rules for revising strategies that individuals use.[15]   A wide range of dynamics are admissible under this specification, including

---

[14] An exception being evolutionary theories for the formation of preferences such as altruism etc.

[15] To this date, only very special and simplistic social evolutionary adaptation mechanisms have been studied (Friedman, 1991). In this respect, the learning literature is more advanced than the evolutionary literature, see discussion in Section V.

arbitrarily fast revisions towards the current best replies.[16]

Formally, the class of dynamics under consideration, to be called *monotone*, is given by

$$\dot{x}_i = \varphi_i(x) \cdot x_i \qquad [\forall i \in I, \forall x \in \Delta], \tag{3}$$

for some Lipschitz continuous function $\varphi : \Delta \rightarrow \mathbb{R}^k$ satisfying the *orthogonality* condition $x \cdot \varphi(x) = 0$ at all states $x \in \Delta$, and the *monotonicity* condition $\varphi_i(x) \lesseqgtr \varphi_j(x)$ if $u(e^i, x) \lesseqgtr u(e^j, x)$, conditions clearly met by the replicator dynamics.[17] The orthogonality condition is necessary and sufficient to leave the unit simplex Δ invariant.

It is easily verified that, in the special case of (symmetric) 2×2 games, the qualititative properties of the replicator dynamics are shared by all monotone dynamics (Weibull, 1992). Hence, irrespective of how close the dynamic adjustment is to instant best–reply behavior, aggregate behavior always converges to Nash equilibrium in these low–dimensional settings.

Moreover, it is not difficult to show that, even for arbitrary (symmetric) k×k games, many properties of the replicator dynamics are valid for any monotone dynamics. In particular, Propositions A and B generalize directly:

---

[16] Note that the limiting case of individuals who instantly switch to the best replies in general does not induce a well–defined dynamics in continuous time, since the best–reply correspondence need not be lower hemi–continous, and hence not always permit a continuous selection, see Gilboa and Matsui (1991) for an approach to this issue.

[17] The latter property is called *relative monotonicity* in Nachbar (1990), *order compatibility* (of predynamics) in Friedman (1991), and *monotonicity* in Samuelson and Zhang (1992).

Proposition A' (Friedman, 1991): If the population state x is (Lyapunov) stable in some monotone dynamics (3), then (x,x) is a Nash equilibrium. If (x,x) is a Nash equilibrium, then x is stationary in any monotone dynamics (3).[18]

Proposition B' (Nachbar, 1990): If an interior dynamic path in a monotone dynamics (3) converges to some state x, then (x,x) is a Nash equilibrium.

The more subtle Proposition C does not appear to be fully generalizable to the present class of dynamics; instead of asserting the extinction of all strategies which are strictly (iteratively) dominated by some pure *or mixed* strategy, Samuelson and Zhang only assert, for monotone dynamics, the extinction of all strategies which are strictly (iteratively) dominated by some *pure* strategy:

Proposition C' (Samuelson and Zhang, 1991): If a pure strategy is (iteratively) strictly dominated by a pure strategy, then its population share converges to zero along any interior dynamic path in any monotone dynamics (3).[19]

The stochastic version of monotone dynamics (3) requires that if one strategy performs better than another, the share of the former is likely to grow faster than that of the latter. A model along these lines is offered by Kandori, Mailath and Rob (1991), who consider stochastic evolution under monotone dynamics in a class of 2×2

---

[18] Friedman (1991) makes the slightly weaker claim (for a wider class of dynamics) that *asymptotic* stability implies Nash equilibrium play, see his Proposition 3.3. A simple proof of the present claim is given in Weibull (1992).

[19] Samuelson and Zhang (1991) in fact show that the statement in Prop.C' is valid for two—population dynamics, implying Proposition C' as a corollary. They also show, again for two—population dynamics, that the slightly stronger statement in Proposition C is valid for a subclass of monotone dynamics which they call *aggregate monotone*. Nachbar (1990) established Proposition C' in the special case of games in which only one strategy survives the iterated elimination of pure strategies which are strictly dominated by other pure strategies.

coordination games. In contrast to Foster and Young (1990), who add a stochastic term to the continuous–time replicator dynamics (1), Kandori, Mailath and Rob assume deterministic montone adjustment towards the more profitable strategy, and introduce the stochastic element by assuming that the population is finite, time is discrete, and each player has a positive (i.i.d) probability of switching, or 'mutating,' to the other strategy. Given these assumptions, the evolution of each individual's strategy is governed by an irreducible Markov chain, and therefore each individual's play will converge to the same ergodic distribution over the two strategies. Furthermore, Kandori, Mailath and Rob show that, as the (exogneous) mutation probability tends to zero, the probability that, in the ergodic limiting distribution, the individual will play the risk–dominant strategy, goes to 1. Hence, not only does this approach, just as the deterministic monotone dynamics approach, lend support for Nash equilibrium; in simple games the stochastic approach even selects among Nash equilibria which, taken individually, are stable in the corresponding deterministic monotone dynamics.[20]

V.   MORE GENERAL BOUNDEDLY RATIONAL BEHAVIORS

The results discussed in the preceding section suggest that the close connection between evolution and Nash equilibrium that was obtained under the replicator dynamics extends to a wide range of adaptive dynamics.[21] Nevertheless, this can only be among the very first steps in studying the long–run evolution of boundedly rational behaviors. In the present section, we try to identify important limitations of the model of bounded

---

[20] See Kandori and Rob (1992) for an extension to k×k games. See Nöldeke and Samuelson (1992) for a model of stochastic evolutionary dynamics in a class of sequential–move games in extensive form.

[21] Actually, the extension is not quite complete since the stronger Proposition C holds only for aggregate monotone dynamics, see Samuelson and Zhang (1991). It is still not clear whether there are reasonable social adaptation procedures which are monotone but not aggregate monotone, see Friedman (1991) and Swinkels (1992) for discussions of wide classes of adaptation dynamics.

rationality underlying the evolutionary models discussed above, and we investigate some of the consequences of generalizing that model.

The only element of the *history* of play that goes into the above models of strategy adaptation is the current payoff. In reality, individuals use their memory, and the success of different strategies over (at least part of) the past count in their decision making. A variety of models of behavioral adaptation based on the history of play is supplied by the learning literature, a research field falling outside the scope of this survey. However, an interesting environment for evolutionary selection over history–contingent behaviors arises when interacting individuals use differing learning or search rules. These rules may be more or less ad hoc, some may use Bayesian updating etc. Within the context of a stochastic asset market with this flavor, Blume and Easley (1991) find that the link between evolutionary selection and rationality is weak. Although their analysis is quite interesting, its setting falls outside the scope of the present survey. Another element of the history that appears often to be significant in social adaptation is the pattern of choice made by others. People may, for example, be prone to choose strategies that others are using, or phrased differently, adopt prevalent behaviors. Weibull (1992) suggests that if players can sample other players' current payoffs and move from less to more profitable strategies on the basis of such sampling, then the resulting selections dynamics is monotone, and hence the results of the previous section apply.

Another important aspect which is not captured in the monotone adaptive dynamics discussed in the previous section is that, in the real world, people are to some extent *forward looking*. There are really two independent issues here. First, in the standard evolutionary models discussed in the previous sections, individuals are myopic in the sense of only paying attention to current payoffs. In models with a finite number of (long–lived) individuals, it may not always be optimal to choose the myopically best strategy. By choosing some other strategy, a player may influence other players' future

behavior to his advantage. Secondly, in choosing their strategies, agents (in finite or infinite populations) may want to anticipate the future behavior of others. After all, they know that as they themselves change strategy, everybody else is also changing theirs, so that what is currently a good strategy may not be quite so good by the time it gets played. This kind of anticipatory behavior is beyond the scope of the models outlined above (see Selten, 1991a). On the other hand, continuous–time monotone dynamics does permit strategy adaptation arbitrarily close to the instantaneous adoption of a best reply (IABR henceforth). It may be argued that as one approaches IABR, the individuals' strategy choices may not be very different from those which incorporate a forward–looking element. (For, after all, the limit case of IABR is one in which agents perfectly anticipate what the others will choose.) However, this is not true; as each individual's play approaches IABR, the whole population may change strategies faster and faster. As a result, even though everybody is only a little late in responding to what the others are doing, they may be way off in strategy space.

Lack of anticipatory behavior is one aspect of the broader issue that standard evolutionary models are based on a very mechanical model of boundedly rational play. Bounded rationality (BR henceforth) should not be equated with using a mechanical rule to move towards a strategy that is currently doing well. Boundedly rational players *think*, albeit imperfectly. The imperfection in their thinking can take many forms and there is no a priori reason why some of these are more appropriate for evolutionary analysis than others. Therefore, the right way to do evolutionary analysis appears to be to allow for the widest possible variety of boundedly rational behaviors.

The paper which comes closest to this project is Stahl (1992), who, in an otherwise standard set–up of evolutionary game theory, introduces a hierarchy of increasingly sophisticated players who he calls *smart*$_n$ players, for n=0,1,2,..[22] The

---

[22] Also Blume and Easley (1991) analyze evolutionary selection of more or less sophisticated behaviors, but do this in the context of a stochastic asset market rather

usual pre—programmed players in evolutionary game theory are called $smart_0$, and, for $n>0$, a $smart_n$ player is one who knows the current aggregate behavior of all players at all lower smartness levels. This limits the possible beliefs of a $smart_n$ player about what his potential opponents could be playing, and he plays a (pure) strategy which is optimal under some belief satisfying this restriction.[23] The total population, within which the usual random matchings takes place, is decomposed into n^* smartness categories, and the distribution accross strategies within each such category n can be represented as a point x^n on the (usual) unit simplex Δ, for $n=0,1,2,...,n^*$.

If the share of the total population in smartness category n is λ_n, the aggregate population behavior is the convex combination $p = \Sigma\lambda_n x^n$, again a point on the unit simplex Δ. Stahl uses a version of the discrete—time replicator dynamics (2) which induces dynamic paths close to those of the continuous—time version (1) (c.f. discussion of Dekel and Scotchmer (1992) and Cabrales and Sobel (1992) in Section III above). In Stahl's setting, this selection dynamics specifies the growth rates of each of the $k\cdot(n^*+1)$ population shares.

Despite this much more general environment, Stahl is able to show that the counterparts to Propositions B and C hold. He also shows that, whether or not the process converges, the subpopulation of $smart_0$ players will never be wiped out, granted, of course, that initially all types of player are present. Intuitively, the $smart_0$ players survive despite the presence of more rational players because some of them happen to be programmed to the "right" strategies, and hence these earn the same payoff in a stationary state as smarter players. Put in the author's words: "being right is just as good as being smart." However, one should remember that in general the aggregate population behavior changes over time. Therefore, smarter players, who can adjust

---

than for pairwise game playing.

[23] In order to specify selections within best reply sets, all players of smartness $n\geq 1$ are assumed to have some lexicographic preference over strategies.

their play to their environment, will often be in a better position than $smart_0$ players who always play the same strategy, and, consequently, the survival of $smart_0$ players is non—trivial:

Proposition D (Stahl, 1992): If an interior path in the discrete—time replicator dynamics converges to some $p \in \Delta$, then (p,p) is a Nash equilibrium. The population shares of $smart_0$ players playing non—rationalizable strategies converge to zero along any interior dynamic path. Moreover, for any interior path in the dynamics, the limit superior of the population share $\lambda_0(t)$ of $smart_0$ players is positive.

In sum: as long as the evolutionary process in Stahl's setting converges, the long—run outcome remains a Nash equilibrium. In this sense, the "as if" justification of Nash equilibrium generalizes beyond the standard environment of evolutionary game theory. On the other hand, whether the process will converge and what its long—run behavior will be like if it does not converge, does depend on the composition of the population in terms of levels of rationality. The result concerning the survival of $smart_0$ players is important because it says that if the initial population contains $smart_0$ players, the long—run population will also contain $smart_0$ players and therefore (as long as the process does not converge) the presence of $smart_0$ players may have a real effect on the long—run behavior of the system. In other words, one cannot simply ignore the actual levels of rationality of the players in the game, as the proponents of the "as if" methodology would suggest we do.

Another generalization of the basic evolutionary environment, different from the ones considered by Stahl (1992), is introduced in Banerjee and Weibull (1992). We define the set of types to be some finite set $T=\{1,2,...,r\}$, and consider any (symmetric) two—player game with pure strategy set $I=\{1,2,...,k\}$. One can then think of a possible

behavior as a map from T to I, i.e., a map b which for each type $\tau \in$ T of opponent prescribes a strategy i$\in$I. Denote the set of all behaviors $\mathscr{B}$. A *player–cell* then is an element of the product set T$\times \mathscr{B}$, i.e., each player in cell (τ,b) is himself of type $\tau \in$ T and uses strategy b$(\tau')\in$I when meeting a player of type $\tau' \in$ T.

In this terminology, the standard environment of evolutionary game theory is the special case of k types, i.e., T=I, and every player of type i$\in$T is constrained to the constant behavior b(j)=i $\forall$j$\in$T. Banerjee and Weibull (1991) corresponds to the somewhat less special case of k+1 types, i.e., T=I$\cup\{$k+1$\}$, each of the k first types being constrained to constant behaviors, just as in the standard setting of evolutionary game theory, and the k+1'st type being constrained to the best–reply mapping.[24] More specifically, we there study the replicator dynamics (1) in an environment with k types of smart_0 players, one for each pure strategy i$\in$I, and one additional type, k+1, who plays a best response against each of the other types, and some rationalizable strategy (or strategies) against its own type. Let α denote the payoff that players of the *optimizing* type k+1 obtain when meeting each other.

In this mixed set–up, and in the correspondingly augmented continuous–time replicator dynamics (1), we provide sufficient conditions for smart_0 players to persist, or vanish, in the long run. Let I^+c I be the (possibly empty) subset of pure strategies which earn more than α when meeting their best replies, and let I^-c I be the (possibly empty) subset of strategies which earn less than α when meeting their best replies. One may refer to strategies in I^+ as *aggressive* and those in I^- as *yielding*. It is intuitively clear that if the game has some aggressive strategy, then the stationary state in which all players are of the optimizing type k+1 is unstable in the replicator dynamics. For suppose there are virtually only such players around. Then these obtain a lower payoff

---

[24] In a similar spirit, Dekel and Scotchmer (1992, Example 2) introduce players who play best responses to last period's population state. Banerjee and Weibull (1991) also study the case of players who play best responses to the current state, in continuous time.

than the few aggressive smart$_0$ players who are present, since the latter meet virtually only optimizing players. Conversely, all yielding smart$_0$ players vanish in the long run:

Proposition E (Banerjee and Weibull, 1991): Consider any interior path in the replicator dynamics (1) for k types of smart$_0$ players and one optimizing type, k+1. If $I^+ \neq \phi$, the share of optimizing players does not converge to 1 over time. For each $i \in I^-$, the population share x_i of players of type i converges to zero.

An immediate consequence of this result is that if the game lacks aggressive strategies, which for example is the case with every constant–sum game which has no symmetric Nash equilibrium in pure strategies, then only optimizing players survive in the long run, along any interior dynamic path. In contrast, if a game has some aggressive strategy (such as the "hawk" strategy in the famous Hawk–Dove game), then some aggressive smart$_0$ players may survive in the long run, along with optimizing players.

Moreover, in striking contrast with the results for the replicator dynamics in the standard setting of evolutionary game theory, (Lyapunov) stable states in environments containing optimizing players need not correspond to Nash equilibrium behavior at all (cf. Proposition A). In fact, even *asymptotically* stable states may involve the playing of a strictly dominated strategy.[25] This disturbing phenomon arises in the following dominance–solvable game:

$$A = \begin{bmatrix} 3 & 1 & 6 \\ 0 & 0 & 4 \\ 1 & 2 & 5 \end{bmatrix} .$$

---

[25] A state $p \in \Delta$ is called *asymptotically stable* if there is a neighborhood $U \subset \Delta$ of p such that solution paths in U approach x as $t \to \infty$, see e.g. Hirsch and Smale (1974).

Here strategy 2 is strictly dominated by strategy 1 (and by strategy 3), and, once strategy 2 has been deleted, strategy 1 strictly dominates strategy 3. In the absence of $smart_0$ players, the outcome would hence be the standard non—cooperative solution that all optimizing players use the unique rationalizable strategy 1. Likewise, if only $smart_0$ players were present in the population, as in the standard setting of evolutionary game theory, the long—run dynamic outcome would, once again, be that only strategy 1 would be used (Proposition C). But the presence of both types of player allows for an asymptotically stable outcome in which all three strategies are used.

To see how this is possible, first note that the strictly dominated strategy 2 is the unique aggressive strategy of the game, and strategy 3 is its unique yielding strategy. By Proposition E, the share x_4 of optimizing players will not converge to 1, and the share x_3 of players of type 3 will converge to zero. Figure 2 shows the qualitative properties of the replicator dynamics near the face $x_3=0$ of the unit simplex. The diagram shows one basin of attraction for a continuum of stationary states, all of which contain a mix of $smart_0$ players of type 1 and optimizing players. In these (Lyapunov) stable states, all players use the unique rationalizable strategy 1. However, one also sees a basin of attraction for an asymptotically stable state in which 2/3 of the players are optimizers and 1/3 are $smart_0$ of the aggressive type 2. In this stationary state, the latter type of player of course use strategy 2, since by definition they have no choice, and the optimizing players use the unique rationalizable strategy 1 against each other and the iteratively dominated strategy 3 against the aggressive players — the latter strategy being the best reply to the strictly dominated strategy 2. As a result, the aggregate population behavior is the completely mixed strategy m=(4,3,2)/9. The reason why $smart_0$ players of type 1 are selected against in a neighborhood of this state is that they perform worse than the optimizing players in states with aggressive (type 2) players present.

Figure 2

Robson (1990) enriches the standard set—up of evolutionary game theory in a similar way. In his model, players may be of (a few different) types, some players can distinguish the type of their opponent and condition their behavior accordingly. He shows that this richer set—up enables evolutionary selection for the Pareto efficient Nash equilibrium in simple coordination games (just as the model of Kandori, Mailath and Rob (1991) would predict), and that it also enables temporary cooperation in the Prisoners' Dilemma game. However, he does not allow for all possible behaviors, and hence, as he notes, cooperation in the Prisoners' Dilemma may break down if some players of the "cooperative type" in fact plays the dominant strategy.

Both these models (Robson, 1990; Banerjee and Weibull, 1991) contain restrictions on the set of behaviors. As the above example shows, such restrictions can lead to outcomes which are sharply at odds with Nash equilibrium. In principle, one

should allow for all possible behaviors (i.e. the full set $\mathscr{B}$) when testing whether or not evolution in such richer behavioral frameworks can lead away from the Nash paradigm (Banerjee and Weibull, 1992).

VI. CONCLUSIONS

In broad outlines, the current stage of research in the field of game–theoretic approaches to evolution and rationality provides some support for the rationalistic approach of non–cooperative game theory. One conclusion which appears to be fairly robust is that if aggregate behavior is stationary and (Lyapunov) stable with respect to evolutionary dynamic forces, then it corresponds to some Nash equilibrium of the underlying game. In this sense, evolutionary selection makes individuals behave as if they obeyed the stringent rationality and coordination requirements underlying the rationalistic approach to Nash equilibrium. Likewise, if the evolutionary selection process converges, then its limiting state (even if it is not stable) again corresponds to some Nash equilibrium. There is nothing inherent in the evolutionary process, however, which guarantees its convergence. If the process does not converge, aggregetate behavior is nevertheless rational, in the sense of not being strictly (iteratively) dominated. Thus, when evolutionary convergence is lost, so is the coordination but not the rationality of the induced aggregate behavior.

Moreover, the above qualitative results remain valid even if the biological reproduction dynamics is replaced by some boundedly rational process of individual behavior adaptation, and irrespective how close the adaptive behavior is to instant switching to optimal strategies. However, the possibility remains open, in games with more than two pure strategies, that fewer trajectories are convergent as adaptive behavior approaches instant switching to optimal strategies. Using somewhat metaphorical language, let us imagine that we identify the form of "rational and

coordinated" behavior, represented by some Nash equilibrium in a given game, by a "point of full rationality", p, in some "rationality space." It is then possible that there is a neighborhood of this point outside which long–run aggregate behavior conforms to the Nash equilibrium but inside which it does not, except at the point p itself. If this is the case, the Nash equilibrium paradigm is problematic, representing only an isolated "ideal point" in "rationality space." It is clearly important for the validity of both the rationalistic and the standard evolutionary approach to determine whether this is the case. More generally, it is an interesting research task to understand whether the introduction of more sophisticated players might make the evolutionary process less likely to converge.

The qualitative results mentioned above all concern individuals who are constrained to particularly simple behaviors in the game played. Judging from some recent research, evolutionary selection in richer behavior spaces, containing more or less boundley rational, information–conditioned behaviors, evolutionary selection, even when convergent, need not result in Nash equilibrium behavior. However, it still remains an open question whether results which are negative in this sense are robust with respect to expansions of the space of behaviors. Maybe we will see some "rationalistic" equilibrium behavior result from evolutionary selection processes in sufficiently rich behavior spaces (Banerjee and Weibull, 1992).

There is also the further issue of whether we are justified in rationalizing Nash equilibrium play in terms of arguments which refer only to the long–run outcomes. After all, the long run could be very long indeed. In terms of interpreting the above results, it would hence be extremely useful if environments and adaptation rules could be identified for which the convergence to Nash equilibrium is relatively rapid/slow.

Finally, a specific laguna in the literature can be located; there is a serious shortage in explicit selection mechanisms with plausible and general boundedly rational micro foundations. Also, it seems desirable to complement the current standard

framework of pairwise interactions towards multi—player interactions (see references mentioned in the introduction) and towards more market—like settings. Viewing this whole enterprise as bridge—building between evolutionary (biological) and rationalistic (economic) approaches, one could say that while most current research builds from the biology end of the bridge, borrowing much of the general setting and technical machinery from there, we now need more work on its economics end, using the settings of standard economic institutions.

REFERENCES

Alchian, A., 1950, "Uncertainty, Evolution and Economic Theory", **Journal of Political Economy** 58, 211–222.

Armbruster, W. and W. Boege, 1979, "Baayesian game theory", in O. Moeschlin and D. Pallaschke (eds), **Game Theory and Related Topics**, North–Holland.

Aumann, R., 1987, "Correlated equilibrium as an expression of Bayesian rationality", **Econometrica** 55, 1–18.

Aumann, R. and A. Brandenburger, 1991, "Epistemic conditions for Nash equilibrium", Hebrew University and Harvard Business School.

Banerjee A. and J.W. Weibull, 1991, "Evolutionary selection and rational behavior", forthcoming in Kirman A. and M. Salmon (eds.), **Rationality and Learning in Economics**, Blackwell.

Banerjee A. and J.W. Weibull, 1992, "Evolution of behaviors contingent upon opponents' identities", work in progress.

Basu, K., 1988, "Strategic irrationality in extensive games", **Mathematical Social Sciences** 15, 247–260.

Basu, K., 1990, "On the non–existence of a rationality definition for extensive games", **International Journal of Game Theory** 19, 33–44.

Bernheim, D., 1984, "Rationalizable strategic behavior", **Econometrica** 52, 1007–1028.

Bernheim, D., 1986, "Axiomatic characterizations of rational choice in strategic environments", **Scandinavian Journal of Economics** 88, 473–488.

Bicchieri, C., 1989, "Self–refuting theories of strategic interaction: a paradox of common knowledge", **Erkenntnis** 30, 69–85.

Binmore K., 1987, "Modelling Rational Players: Part I", **Economics and Philosophy** 3, 179–214.

Binmore K., 1988, "Modelling Rational Players: Part II", **Economics and Philosophy** 4, 9–55.

Binmore K. and L. Samuelson, 1992, "Evolutionary Stability in Repeated Games Played by Finite Automata", **Journal of Economic Theory** 57, 278–305.

Blume, L. and D. Easley, 1991, "Economic Natural Selection and Adaptive Behavior", mimeo, Dept. of Economics, Cornell University, forthcoming in Kirman A. and M. Salmon (eds.), **Rationality and Learning in Economics**, Blackwell.

Bomze, I. M., 1986, "Non–Cooperative Two–Person Games in Biology: A Classification", **International Journal of Game Theory** 15, 31–57.

Boylan, R., 1992, "Laws of large numbers for dynamical systems with randomly matched individuals", **Journal of Economic Theory** 57, 473–504.

Cabrales A. and J. Sobel, 1992, "On the limit points of discrete selection dynamics", **Journal of Economic Theory** 57, 407–419.

Canning, D., 1989, "Convergence to Equilibrium in a Sequence of Games with Learning", LSE Disc. Paper TE/89/190.

Dekel, E. and S. Scotchmer, 1992, "On the Evolution of Optimizing Behavior", **Journal of Economic Theory** 57, 392–406.

Friedman, D., 1991, "Evolutionary Games in Economics", **Econometrica** 59, 637–666.

Friedman, M., 1953, "The Methodology of Positive Economics", in Friedmann M., **Essays in Positive Economics**, University of Chicago Press, Chicago.

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

Freidlin, M. and A. Wentzell, 1984, **Random Perturbations of Dynamical Systems**, Springer Verlag.

Fudenberg, D. and C. Harris, 1992, "Evolutionary dynamics with aggregate shocks", **Journal of Economic Theory** 57, 420–441.

Gilboa, I. and A. Matsui, 1991, "Social stability and equilibrium", **Econometrica** 59, 859–867.

Hirsch, M. and S. Smale, 1974, **Differential Equations, Dynamical Systems, and Linear Algebra**, Academic Press.

Hofbauer, J., P. Schuster and K. Sigmund, 1979, "A note on evolutionary stable strategies and game dynamics", **Journal of Theoretical Biology** 81, 609–612.

Hofbauer, J. and K. Sigmund, 1988, **The Theory of Evolution and Dynamical Systems**, London Mathematical Society Students Texts, Vol. 7, Cambridge University Press, Cambridge.

Johansen, L., 1982, "The status of the Nash type of noncooperative equilibrium in economic theory", **Scandinavian Journal of Economics** 84, 421–441.

Kandori, M., G. Mailath and R. Rob,, 1991, "Learning, Mutation and Long Run Equilibria in Games", mimeo, University of Pennsylvania.

Kandori, M. and R. Rob, 1992, "Evolution of equilibria in the long run: a general theory and applications", University of Pennsylvania, CARESS WP #92–06R.

Maynard Smith, J. and G.R. Price, 1973, "The Logic of Animal Conflict", **Nature** 246, 15–18.

Maynard Smith, J., 1974, "The Theory of Games and the Evolution of Animal Conflicts", **Journal of Theoretical Biology** 47, 209–221.

Maynard Smith, J., 1982, **Evolution and the Theory of Games**, Cambridge University Press, Cambridge.

Nachbar, J., 1990, "'Evolutionary' Selection Dynamics in Games: Convergence and Limit Properties", **International Journal of Game Theory** 19, 59–89.

Nöldeke, G. and L. Samuelson, 1992, "The evolutionary foundations of backward and forward induction", University of Bonn and University of Wisconsin, mimeo.

Pearce, D., 1984, "Rationalizable strategic behavior and the problem of perfection", **Econometrica** 52, pp.1029–1050.

Ritzberger, K. and K. Vogelsberger, 1991, "The Nash field", Institute for Advanced Studies, Vienna, mimeo.

Robson, A.J., 1990, "Efficiency in evolutionary games: Darwin, Nash and the Secret Handshake, **Journal of Theoretical Biology** 144, 379–396.

Rosenthal, R., 1981, "Games of perfect information, predatory pricing and the chain–store paradox, **Journal of Economic Theory** 25, 92–100.

Samuelson, L., 1991, "Does evolution eliminate dominated strategies?", University of Wisconsin, mimeo.

Samuelson, L. and J. Zhang, 1992, "Evolutionary Stability in Asymmetric games", **Journal of Economic Theory** 57, 363–391.

Schuster, P., Sigmund K., Hofbauer J. and R. Wolff, 1981a, "Selfregulation of behaviour in animal societies I: Symmetric contests", **Biological Cybernetics** 40, 1–8.

Schuster, P., Sigmund K., Hofbauer J. and R. Wolff, 1981b, "Selfregulation of behaviour in animal societies II: Games between two populations without selfinteraction", **Biological Cybernetics** 40, 9–15.

Schuster, P. and K. Sigmund, 1983, "Replicator dynamics", **Journal of Theoretical Biology** 100, 533–538.

Schuster, P. and K. Sigmund, 1985, "Towards a dynamics of social behaviour: strategic and genetic models for the evolution of animal conflicts", **Journal of Social Biological Structure** 8, 255–277.

Selten,R., 1991a, "Anticipatory learning in two–person games", in R. Selten (ed), **Game Equilibrium Models I**, Springer Verlag.

Selten, R, 1991b, "Evolution, Learning and Economic Behavior", **Games and Economic Behavior** 3, 3–24.

Stahl, D., 1992, "Evolution of Smart$_n$ Players", University of Texas, mimeo.

Swinkels, J., 1992, "Adjustment dynamics and rational play in games", Stanford University.

Tan T. and S.R. Werlang, 1988, "The Bayesian foundations of solution concepts of games", **Journal of Economic Theory** 45, 370–391.

Taylor, P. and L. Jonker, 1978, "Evolutionary Stable Strategies and Game Dynamics", **Mathematical Biosciences** 40, 145–156.

Weibull, J.W., 1992, "An introduction to evolutionary game theory", Department of Economics, Stockholm University, mimeo.

Winter, S., 1964, "Economic 'Natural' Selection and the Theory of the Firm", **Yale Economics Essays** 4, 225–272.

Winter, S., 1971, "Satisficing, selection, and the innovating remnant", **Quarterly Journal of Economics** 85, 237–261.

Zeeman, E.C., 1980, **Population Dynamics from Game Theory**, Lecture Notes in Mathematics 819, Springer Verlag.

Zeeman, E.C., 1981, "Dynamics of the Evolution of Animal Conflicts", **Journal of Theoretical Biology** 89, 249–270.