

Weibull, Jörgen W.

Working Paper

What have we learned from Evolutionary Game Theory so far?

IUI Working Paper, No. 487

Provided in Cooperation with:

Research Institute of Industrial Economics (IFN), Stockholm

Suggested Citation: Weibull, Jörgen W. (1997) : What have we learned from Evolutionary Game Theory so far?, IUI Working Paper, No. 487, The Research Institute of Industrial Economics (IUI), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/94851>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WHAT HAVE WE LEARNED FROM EVOLUTIONARY GAME THEORY SO FAR?

JÖRGEN W. WEIBULL*

STOCKHOLM SCHOOL OF ECONOMICS AND I.U.I.

May 7, 1998

ABSTRACT. Evolutionary theorizing has a long tradition in economics. Only recently has this approach been brought into the framework of non-cooperative game theory. Evolutionary game theory studies the robustness of strategic behavior with respect to evolutionary forces in the context of games played many times in large populations of boundedly rational agents. This new strand in economic theory has led to new predictions and opened up doors to other social sciences. The discussion will be focused on the following questions: What distinguishes the evolutionary approach from the rationalistic? What are the most important findings in evolutionary game theory so far? What are the next challenges for evolutionary game theory in economics? Doc: *esem.tex*

1. INTRODUCTION

In the past ten years or so economists have become interested in evolutionary game theory. New models have been developed and results from the biology literature have been generalized. The economics literature now contains a large number of journal articles and a few monographs in this field, see the bibliography at the end of the essay.

The rationalistic approach to game theory assumes one perfectly rational player for each position in the game, and these players play the game exactly once, the game and the equilibrium being common or mutual knowledge, see e.g. Tan and Werlang (1988) and Aumann and Brandenburger (1995). By contrast, the evolutionary approach assumes that the game in question is played many times by boundedly rational players who are randomly drawn from large populations and who have little or no information about the game. These two approaches are complementary, and each has its merit. In some games the predictions are the same, in others they differ. Also, there are differences within each of the two approaches. This essay tries to outline the state of art of evolutionary game theory in economics and compare its predictions with those in standard non-cooperative game theory.

* This essay is composed of notes for an invited talk held 30 August 1997 at the Econometric Society European Meeting in Toulouse. I am grateful for comments from Eric van Damme and Fernando Vega-Redondo to an early draft.

1.1. The economist's research routine. Much of current economic research of the applied theory variant goes as follows:

1. Study an economic issue.
2. Construct a model, solve for its equilibria.
3. Use the equilibria, or a selection of these, as predictions for the economic issue.

Suppose that the researcher goes out of his way and shows his work to a non-economist friend. *She* asks: Why do you model us as being able to solve any complex mathematical problem, to form sophisticated expectations about each other and even coordinate these expectations between us? Why don't you model people as we are and as we behave, boundedly rational and not perfectly coordinated? Why not learn from psychology, sociology, observations from field studies and laboratory experiments? *You*, the economist-researcher, answer that you expect that indeed in many situations people will tend to behave as if they were as smart and well-coordinated as the imagined agents are in your model, granted, of course, that their environment is sufficiently stationary and that the economic agents are familiar with the interaction, for example by repeatedly acting in this or similar situations. Suboptimal behaviors will be selected against in the long run by a trial-and-error process, and market forces will select against firms and firm practices that perform poorly. *She* says: If this is your justification, why don't you prove this as if claim? Why don't you identify conditions under which boundedly rational economic agents can be expected to behave according to your rationalistic economic models?

This is ultimately what the economics branch of evolutionary game theory strives to do. And it is not a Panglossian exercise. The aim is not to justify conventional equilibrium predictions. The predictions may potentially be, and indeed turns sometimes out to be, sets of non-equilibrium behaviors.

1.2. Limitations to the discussion. Nothing will be said about *evolutionary economics*. This is a field by itself within the profession, usually associated with the work of Schumpeter. Evolutionary economics and evolutionary game theory have so far developed completely independently from each other. Nothing will be said about the *evolution of preferences*. There is a literature on evolutionary selection of preferences, in particular on the survival value of altruism and risk attitudes. However, evolutionary game theory takes preferences as fixed and given. Nothing will be said about *market survival*. There is an emerging and potentially important literature, starting with Winter (1964), followed by Winter (1971), Nelson and Winter (1982), and recently taken up by Blume and Easley (1992, 1995, 1996), Dutta (1992), Dutta and Radner (1993), Radner (1995), Vega-Redondo (1997), and Nöldeke and Samuelson (1997). Although the topic is closely related to that of evolutionary game theory, these two literatures so far differ methodologically. Hopefully elements of both will be combined in the future.

Maybe needless to say, the discussion is also limited by idiosyncratic limitations to my memory, knowledge and understanding. I apologize for omissions and misrepresentations.

2. THE EVOLUTIONARY APPROACH

2.1. General characteristics of an evolutionary model. An evolutionary model combines two processes: one *selection process* that favors some varieties over others, and one process that creates this variety, to be called the *mutation process*. In evolutionary game theory, the varieties in question are strategies in a game.

In nature, the basic selection mechanism is biological survival and reproduction, and the mutation process is basically genetic. *In the market place*, the basic selection mechanism is economic survival, and the mutation process is experimentation, innovation and mistakes. In both cases there is also an element of individual and social learning. *In other social and economic interactions*, the selection mechanism can in extreme cases be that of biological or social survival, but under more normal circumstances individuals and households adapt by way of individual and social learning.

A qualitative difference between evolutionary and rationalistic approaches is that while the second focuses on individuals and what goes on in their minds, the evolutionary approach usually instead analyzes the population distribution of behaviors (decision rules, strategies). One could say that the selection process replaces the mental process of choice made by rational players in classical non-cooperative game theory, while the mutation process replaces the mental process of exploring the strategy set and strategies' payoff consequences.¹

The evolutionary approach potentially closes the open-endedness of the bounded-rationality approach: which out of the plethora of boundedly rational behaviors will survive in a population of strategically interacting agents? This requires, of course, that we have a clear picture of the nature of the relevant selection and mutation processes. There is still a long way to go.

See Selten (1991) for a discussion of rationalistic and evolutionary theories of economic behavior.

2.2. Evolutionary theorizing in economics. Evolutionary explanations are not new to economics. Indeed, they can be found in the social sciences much before Darwin. For example, Adam Smith writes: The division of labor, from which so many advantages are derived, is not originally the effect of any human wisdom, which foresees and intends that general opulence to which it gives occasion. It is the necessary, though very slow and gradual consequence of a certain propensity in human nature which has in view no such extensive utility; the propensity of truck, barter, and exchange one thing for another. (Wealth of Nations, p.25)

¹Easley and Rustichini (1996) develop an axiomatic model of decision processes in the mind of an economic agent, including an evolutionary mechanism of how decision rules are mentally selected in an uncertain environment.

Evolutionary thinking can also be found in the writings of *Malthus*, *Marshall*, *Schumpeter* and *Hayek*.

Turning to game theorists, we note that John Nash, already when suggesting his equilibrium concept, had in mind two interpretations, one rationalistic and one that he called the mass action interpretation : It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal. To be more detailed, we assume that there is a population (in the sense of statistics) of participants for each position of the game. Let us also assume that the average playing of the game involves n participants selected at random from the n populations, and that there is a stable average frequency with which each pure strategy is employed by the average member of the appropriate population. Since there is to be no collaboration between individuals playing in different positions of the game, the probability that a particular n -tuple of pure strategies will be employed in a playing of the game should be the product of the probabilities indicating the chance of each of the n pure strategies to be employed in a random playing. (John Nash, Ph.D. thesis, pp. 21-22.) Nash claims that this leads to play of a Nash equilibrium: Thus the assumption we made in the mass action interpretation led to the conclusion that the mixed strategies representing the average behavior in each of the populations form an equilibrium point. (*op.cit.*, p. 22)

Current research on learning and evolution is done much in the spirit of this interpretation.

3. EVOLUTIONARY GAME THEORY

3.1. The evolutionary stability concept. A key concept in evolutionary game theory is that of an evolutionarily stable strategy (ESS), a concept due to Maynard Smith and Price (1973), see also Maynard Smith (1974,1982). Such a strategy is robust to evolutionary pressures in a specific sense: a population playing such a strategy is *uninvadable* by any other strategy. Suppose pairs of individuals are repeatedly drawn at random from a large population to play a symmetric (and finite) two-person game, and suppose that initially all individuals play a certain pure or mixed strategy x in the game. Now let in a small population share of individuals who all play some other pure or mixed strategy y . The incumbent strategy x is evolutionarily stable if, for each such mutant strategy y , there exists a positive invasion barrier such that if the population share of individuals playing the mutant strategy y falls below this barrier, then x earns a higher (expected) payoff than y in the post-entry population.

Formally: for each strategy $y \neq x$ the following inequality should hold for all sufficiently small $\varepsilon > 0$

$$u[x, (1 - \varepsilon)x + \varepsilon y] > u[y, (1 - \varepsilon)x + \varepsilon y] , \quad (1)$$

where the expression on the left-hand side denotes the (expected) payoff to strategy x when played against the mixed strategy $(1 - \varepsilon)x + \varepsilon y$ that represents the post-entry population mix, and the expression on the right-hand side likewise denotes the (expected) payoff to strategy y .

A useful characterization of evolutionary stability, indeed its original definition, is that a strategy x is evolutionarily stable if and only if (a) it is a best reply to itself, (b) it is a better reply to all other best replies than these are to themselves. To see that (a) is necessary it is sufficient to observe that otherwise there would exist a better reply y to x . Represented in a sufficiently small population share this mutant strategy would almost always meet strategy x and hence earn more than x on average. Likewise, (b) is necessary because otherwise there would be some alternative best reply y to x which would earn the same payoff as x when meeting x and at least as much as x when meeting y , and hence y would on average earn more in the post-entry population.

Note that the evolutionary stability criterion does not explain *how* a population arrives at such a strategy. Instead, it asks whether, once reached, a strategy is robust to evolutionary pressures. Note also that evolutionary stability is a robustness test against a single mutation at a time. It does not deal with situations in which two or more mutant strategies are simultaneously present in the population. It thus implicitly treats mutations as sufficiently rare events so that the population has time to return to *status quo* before another mutation occurs.

Despite its biological stance, evolutionary stability also provides a relevant robustness criterion for human behaviors in a wide range of situations. Evolutionary stability then requires that any small group of individuals who try some alternative strategy do less well than the individuals who stick to the incumbent strategy. Consequently, individuals who use the incumbent strategy have no incentive to change their strategy and the latter have an incentive to switch to the incumbent strategy. An evolutionarily stable strategy in such a social setting may be thought of as a *convention* or an established *code* of conduct.

Unfortunately, many games have no evolutionarily stable strategies. Accordingly, researchers have investigated weakenings and set-valued versions of evolutionary stability, see Maynard Smith (1982), Thomas (1985a,b) and Swinkels (1992a,b). Moreover, the ESS concept does not fruitfully generalize to n -player games. Essentially, it then requires *strict* Nash equilibrium play, i.e., that each strategy should be the *unique* best reply to the strategy profile, see Selten (1980,1983,1988).

3.2. The replicator dynamics. The replicator dynamics is an explicit model of a selection process, specifying how population shares associated with different pure strategies in a game evolve over time. The mathematical formulation of the replicator dynamics is due to Taylor and Jonker (1978). They imagine a large population of individuals who are randomly matched over time to play a finite symmetric two-player game, just as in the setting for evolutionary stability. However, here individuals only play pure strategies. A *population state* is a distribution x over pure strategies. Such

a state is mathematically equivalent with a mixed strategy in the game.

If the payoffs in the game represent biological fitness, i.e., the number of offspring, and each child inherits its single parent's strategy, then the number of individuals using pure strategy i will (in a large population) grow exponentially at a rate that equals the (expected) payoff $u(e^i, x)$ to pure strategy i when played against the mixed strategy x that represents the current strategy distribution in the population. It follows that the growth rate of the population *share* using any pure strategy i equals the difference between the strategy's payoff and the average payoff in the population. The latter is identical with the (expected) payoff $u(x, x)$ to mixed strategy x when played against itself. This is the single-population *replicator dynamics* for symmetric two-player games:

$$\dot{x}_i = [u(e^i, x) - u(x, x)] x_i \quad (2)$$

Note that the best pure replies to the current population state x have the highest growth rate in the population, the second-best pure replies have the second-highest growth rate etc. However, although more successful pure strategies grow faster than less successful ones, the average payoff in the population need not grow over time. The reason for this possibility is that if an individual is replaced by an individual using a better strategy, the opponents meeting this new individual may receive lower payoffs. This is for example the case in a Prisoner's Dilemma game. If initially virtually all individuals play cooperate then they will gradually switch to defect, and the average payoff will fall. However, if the game is *doubly* symmetric in the sense that the two players always receive equal payoffs, then the *Fundamental Law of Natural Selection* does hold: the average payoff in the population grows over time, although not necessarily to a global maximum (Losert and Akin, 1983). This is for example the case in a coordination game, where all individuals gradually move toward the same pure strategy.

The replicator dynamics can be readily generalized to n -player games, played by individuals randomly matched in n -tuples from n population, one for each player position, just as in Nash's mass action interpretation. There are two versions of the n -population replicator dynamics, one due to Taylor (1979), and another due to Maynard Smith (1982).

More generally, we will refer to (single- and multiple-population) selection dynamics in the following two categories. *Payoff-positive* selection dynamics, where *all* pure strategies that perform above average have positive growth rates, and *all* pure strategies that perform below average have negative growth rates. *Weakly payoff-positive* selection dynamics, where at least *some* pure strategy that performs above average (granted such a strategy exists) has a positive growth rate. Both versions of the n -population replicator dynamics are payoff-positive, and all payoff-positive selection dynamics are weakly payoff-positive.

4. LEARNING MODELS AND SELECTION DYNAMICS

One can divide learning models into three broad categories, belief-based learning, reinforcement learning, and learning by imitation. It has recently been shown that the replicator dynamics results from certain models in the last two categories.

4.1. Reinforcement learning. A central model in the psychology literature on individual learning is the so-called *reinforcement* model due to Bush and Mosteller (1951). The basic idea goes back to *The Law of Effect* (Thorndike, 1898): Choices that have led to good outcomes in the past are more likely to be repeated in the future.²

The Bush-Mosteller reinforcement learning model, and generalizations of it, have been used in a number of laboratory experiments where human subjects play games, see Roth and Erev (1995), Erev and Roth (1997), and Camerer and Ho (1997). Unfortunately, little is known about the general mathematical properties of these models.

However, Börgers and Sarin (1997) make a theoretical comparison of Cross (1973) version of Bush and Mosteller's (1951) learning model with the (Taylor 2-populations) replicator dynamics. While this learning process is stochastic and evolves in discrete time, the replicator dynamics is deterministic and evolves in continuous time. Börgers and Sarin show that, in an appropriately constructed continuous time limit, their learning process is approximated, over bounded time intervals, by the replicator dynamics.

More exactly, they study a finite two-player game played repeatedly in rounds $n = 1, 2, \dots$, between a fixed pair of players who use mixed strategies. Each player updates the probabilities with which she uses the pure strategies at her disposal as follows. If player 1 (likewise for player 2) in round n of the game uses pure strategy k and obtains a positive payoff $V_k(n)$, a stochastic variable that depends on the random choice made by player 2, then she increases her future probability for using this strategy, the more so the higher is the payoff. Conditional on having played strategy k in round n , her probabilities in round $n + 1$ are updated as follows:

$$\begin{cases} X_k(n+1) = X_k(n) + \delta V_k(n) [1 - X_k(n)] \\ X_h(n+1) = X_h(n) - \delta V_k(n) X_h(n) \quad \forall h \neq k \end{cases} \quad (3)$$

Player 2 updates her choice probability vector Y in the same way. All payoffs are assumed to lie in the open unit interval.³ Hence all choices are reinforcing; once a strategy has been used its probability is increased.

Starting from any initial probability vectors $X(0) = x^o$ and $Y(0) = y^o$, equations (3) define a Markov chain $\{X(n), Y(n)\}_{n=1}^{\infty}$ in the game's mixed-strategy space. The parameter $\delta > 0$ represents the time between two rounds of the game: $t = \delta n$ is the real time at which round n is played. Börgers and Sarin obtain their continuous-time limit of this process, evaluated at any finite real time t , by letting $n \rightarrow \infty$

²Note that choice is implicitly taken to be probabilistic in the citation.

³The payoffs in this model cannot be interpreted as von Neumann-Morgenstern utilities.

and $\delta \rightarrow 0$ in such a way that $\delta n = t$. Hence, rounds are played at shorter and shorter time intervals and probabilities are adapted in proportionately smaller amounts. They show that in this limit the process places unit probability on the state $(x(t), y(t))$ that the replicator dynamics would have reached at time t if its initial state at time zero had been (x^o, y^o) . In this sense, the replicator dynamics approximates the reinforcement dynamics (3) over bounded time intervals.

However, the asymptotic properties of the two models are quite distinct. To see the intuition for this discrepancy, suppose, for instance, that the payoff to player 1 is constant - independent both of her strategy choice and that of player 2. Let the initial state of player 1 in the reinforcement dynamics assign equal probabilities to all pure strategies available to player 1. Likewise, let the initial state of the replicator dynamics assign equal population shares to all pure strategies. The solution to the replicator dynamics is clearly constant: all population shares remain equal forever. However, a realization of the reinforcement dynamics may easily converge over time to any one of the pure strategies available to player 1. For the pure strategy that happens to be drawn in the first round will have a higher probability to be drawn in the next round etc., a property of the reinforcement dynamics that makes it possible for player 1 to lock in on anyone of her pure strategies over time (with *a priori* equal probability for each pure strategy to be so selected). Börgers and Sarin show that for any finite two-player game the reinforcement dynamics converges with probability one to a pure strategy profile - unlike the replicator dynamics.

4.2. Learning by way of imitation.

A. Aspiration levels. Gale, Binmore and Samuelson (1995) provide a simple model of social learning in a finite but large population of individuals playing pure strategies. Each individual maintains an aspiration level for the payoff to be earned in the game. Again at discrete times $0, \delta, 2\delta, \dots$, where $\delta > 0$ is small, the population share δ of individuals, drawn at random from the population, compare their current payoffs with their aspiration levels. If an individual's realized payoff falls below her aspiration level then she imitates a randomly drawn individual in the population, with equal probability for all other individuals in the same player population.

It follows that if aspiration levels have a rectangular distribution (over some interval containing all possible payoff values) in the population, then the probability for imitation is linearly decreasing in the expected payoff to the individual's current strategy. The authors show that for small δ this process can be approximated by the replicator dynamics over bounded time intervals.

B. Review impulses. Alternatively one may view individual strategy adaptation as a stochastic process in continuous time. Suppose that every now and then each individual in a finite population gets an impulse to revise her (pure) strategy choice. If these impulses arrive according to i.i.d. Poisson processes, then the probability of simultaneous impulses is zero, and the aggregate process is also a Poisson process. Moreover, the intensity of the aggregate process is just the sum of the intensities of

the individual processes. If the population is large, then one may approximate the aggregate process by deterministic flows given by the expected values.

Björnerstedt and Weibull (1996) study a number of such models, where revising individuals imitate other individuals in their own player population, and show that a number of payoff-positive selection dynamics, including all three versions of the replicator dynamics, may be so derived.⁴ In particular, if the individual revision rate is linearly decreasing in the expected payoff to her strategy (or to the individual's latest payoff realization), then the intensity of each pure strategy's Poisson process will be proportional to its population share, and the proportionality factor will be linearly decreasing in its expected payoff. If every revising individual selects her future strategy by imitating a randomly drawn individual in their own player population, then the resulting flow approximation is again the replicator dynamics.

C. Why imitate. Schlag (1997) analyses the question what imitation rules an individual *should* choose, when she now and then has the opportunity to imitate another individual in the same player position but is otherwise constrained by severe restrictions on information and memory. He finds that if the individual wants a learning rule that is *payoff increasing* in all stationary environments, then the individual should (a) always imitate (not experiment) when changing strategy, (b) never imitate an individual whose payoff realization was worse than her own, and (c) imitate individuals whose payoff realizations are better than her own, with a probability that is proportional to this payoff difference.

This model is developed for a variety of environments. In the context of a finite two-player game, Schlag imagines pairwise matchings between individuals randomly drawn from two equally large finite populations, one for each player position. Individuals always play pure strategies. After every payoff realization each individual samples at random another individual in her player population and compares the two payoff realizations. A *behavior rule* is a function that maps the pair of payoff realizations and strategies used to a probability distribution over the set of pure strategies available in the player position of the game, the new strategy to adopt. In other words, the only data the individual is allowed to use is such a pair of payoff realizations and pure strategies. In particular, payoff realizations from earlier rounds are neglected. Moreover, it is assumed that the individual uses the same behavior rule in all games (decision situations) with the same number of pure strategies (decision alternatives); i.e., individuals need not know what game they play, it is sufficient that they know the number of pure strategies at their disposal.

A behavior rule that results in a weakly increasing expected payoff in any such game, and for any fixed strategy distribution in the opponent population, is called *improving*. The main result in the paper is a characterization of all such rules. One feature of improving rules is that they are *imitative*: an individual either sticks to her original strategy or adopts the sampled individual's strategy; she never switches

⁴The approximations in Björnerstedt and Weibull (1996) are heuristic. Rigorous approximation techniques are adapted to this setting in Benaïm and Weibull (1997).

to a third strategy.

A behavior rule is said to *dominate* the improving rules if there is no improving rule that yields a higher expected rate of payoff improvement in some game and against some strategy distribution in the opponent population. Schlag identifies a certain behavior rule that dominates the improving rules, and which has properties (a)-(c) given above. This *proportional imitation* rule is itself an improving rule and is shown to have a number of other attractive properties. Schlag shows that a discrete-time version of the (Taylor two-population) replicator dynamics approximates the induced stochastic process over any given finite time horizon, granted the populations are sufficiently large.

Remark: In an earlier version of his paper, Schlag showed that the two above-mentioned learning rules (subsections A and B), which do not rely on the imitated individual's payoff, dominate those improving rules that do not depend on the imitated individual's payoff.

5. SOME RESULTS IN EVOLUTIONARY GAME THEORY

The focus will here be on explicitly dynamic population models for finite n -player games, where individuals play pure strategies. First deterministic selection dynamics are studied, then stochastic evolutionary models will be discussed where a deterministic or stochastic selection process is combined with a stochastic mutation process.

5.1. In the long-run. An immediate way to study the properties of a deterministic dynamic model is to select an initial population state and go for a cup of coffee while the computer runs. One should then just remember to initially let *all* pure strategies be present in the population, since initially extinct strategies will remain extinct forever in a *selection* process. Such solution trajectories will be called *interior*. The solution trajectory may or may not settle down over time, i.e., may *converge* or *diverge*.

A. Convergence. If the population state does converge, what is the nature of the long-run limit state? It turns out that in any weakly payoff-positive selection dynamics, and along any convergent interior solution trajectory, the limit state necessarily constitutes a Nash equilibrium. The researcher, when he or she returns after the coffee break, will find that the population plays (an approximation of) some Nash equilibrium. This result was first proved for the (single-population) replicator dynamics by Nachbar (1990). See Weibull (1995) for the more general result.

In sum: if the selection process meets the relatively mild condition of weak payoff-positivity, and if aggregate behavior settles down over time, then individuals in the long-run population state will behave as if they all expected a particular Nash equilibrium played a best reply against this - just as claimed in Nash's mass action interpretation.

Remark: Use of selection dynamics of the present type presume that whole strategies are selected (get imitated, learned). By contrast, if, in an extensive-form game,

only the local strategies, or actions, at *reached* information sets are selected (get imitated, learned) then convergence does not necessarily imply Nash equilibrium in the limit, but some correspondingly weaker equilibrium property such as *self-enforcing equilibrium* (see Fudenberg and Levine (1993a,b, 1997), Kalai and Lehrer (1993), Fudenberg and Kreps (1995), Nöldeke and Samuelson (1995)).⁵

B. Divergence. If an interior solution to a weakly payoff-positive selection dynamics converges over time, then we have seen that the surviving strategies are rational in the sense of being best replies to the resulting mixed-strategy profile. The question hence is what happens if the solution trajectory *does not* converge. When there is no hope of equilibrium play in the long run we are led to the question whether play is *rational*.

A basic rationality postulate in non-cooperative game theory is that players never use pure strategies that are strictly dominated.⁶ This postulate requires no knowledge of other players' preferences or behavior. A more stringent rationality-*cum*-knowledge postulate is that players never use pure strategies that are *iteratively* strictly dominated. In addition to avoidance of strictly dominated strategies, this postulate requires that all players know each others payoffs, that they know that they know each others payoffs, etc. up to a some finite level of mutual knowledge sufficient to halt the procedure of iterated elimination of strictly dominated pure strategies (see e.g. Tan and Werlang (1988)).

A fundamental question in evolutionary game theory thus is whether evolutionary selection processes do eliminate all strictly dominated pure strategies or even all iteratively strictly dominated pure strategies. If all iteratively strictly dominated strategies do vanish, this provides an evolutionary justification for the presumption that strategically interacting agents behave as if it were mutual knowledge that they are rational. Clearly, this justification is more compelling the wider is the class of evolutionary selection processes for which this result is valid.

Akin (1980) shows that all strictly dominated pure strategies vanish along any interior solution trajectory to the (single-population) replicator dynamics in any (finite) symmetric two-player game. Samuelson and Zhang (1992) extend this conclusion to all iteratively strictly dominated pure strategies in a certain subclass of payoff-

⁵Self-enforcing equilibrium essentially requires that each player's strategy be optimal under her beliefs and that each player's beliefs are correct on the equilibrium path through the game tree. In general, this criterion is weaker than that of a Nash equilibrium. However, under certain conditions these two criteria are *outcome equivalent* in the sense that to each self-enforcing equilibrium there exists a Nash equilibrium which induces the same probability distribution over the end-nodes of the game tree, see Fudenberg and Levine (1993a,b).

⁶A pure strategy is strictly dominated if there is some pure or mixed strategy that yields a higher payoff against all strategy profiles in the game. A pure strategy is iteratively strictly dominated if it is strictly dominated in the original game G , or in the reduced game G' obtained by elimination from G of all strictly dominated pure strategies in G , or in the further reduced game G'' obtained by elimination from G' of all strictly dominated pure strategies in G' , etc. The set S of pure-strategy profiles being finite, this procedure stops after a finite number of iterations (and the result is independent of the details of the elimination procedure).

positive two-population selection dynamics that they call *aggregate monotonic*. Dekel and Scotchmer (1992) show, by way of a counter-example, that a strictly dominated strategy may survive forever in a discrete-time version of the replicator dynamics. Björnerstedt (1995) provide a weakly payoff-positive selection dynamics in continuous time and a game in which a strictly dominated strategy survives forever. Hofbauer and Weibull (1996) introduce a class of weakly-payoff positive selection dynamics that contain the aggregate monotonic selection dynamics. This property, *convex monotonicity*, is sufficient, and essentially also necessary, for the elimination of iteratively strictly dominated pure strategies in continuous time in any (finite) n -player game. Cabrales (1996) develops a stochastic version of the (Taylor, n -population) replicator dynamics by adding perpetual random shocks to the population state and to the payoffs. He provides conditions under which iteratively strictly dominated strategies still get wiped out in the long run.

Remark: As a by-product we obtain that selection dynamics that wipe out iteratively strictly dominated strategies lend some support to *forward induction* reasoning in extensive-form games. Forward induction relies in part on players assuming other players to avoid strictly dominated strategies (see van Damme (1989)). This question is studied in some detail in Nöldeke and Samuelson (1993).

5.2. Stability. A complementary approach to long-run analysis is to study *stability properties* of population states, i.e., how population states react to small perturbations.

Bomze (1986) shows that if a population state is (weakly) dynamically stable in the (single-population) replicator dynamics, then this state, viewed as a mixed strategy, necessarily is a best reply to itself.⁷ Hence, not only the static stability criterion of evolutionary stability, but also dynamic stability in the replicator dynamics implies Nash equilibrium play. This result can be generalized to any weakly payoff-positive selection dynamics in any (finite) n -player game (Weibull (1995)). In sum: stability against evolutionary pressures, formalized in various ways, require Nash equilibrium play. However, not all Nash equilibria are stable in this respect. Hence, these evolutionary stability criteria are *refinements* of the Nash equilibrium concept. A comparison with non-cooperative refinements follows.

Remark: If in extensive-form games not whole strategies are selected, only the local strategies at reached information sets, then the conclusion is accordingly weakened to self-enforcing equilibria. Moreover, stability is harder to obtain since there is more scope for drift at unreached information sets, see Nöldeke and Samuelson (1993).

A. More cutting power against mixed-strategy Nash equilibria. Mixed-strategy Nash equilibria are thought by many game theorists and practitioners of

⁷The stability criterion used here is that of Lyapunov stability: a state x is *Lyapunov stable* if every neighborhood B of x contains a neighborhood A of x such that every solution starting in A remains forever in B . A state that is not Lyapunov stable is called *unstable*.

game theory to be inherently unstable since players do not have an incentive to randomize according to the equilibrium prescription. Despite this, completely mixed strategy equilibria easily meet such requirements of the Nash equilibrium concept as Selten's (1975) trembling hand perfection. By contrast, stability with respect to evolutionary pressures does reject certain completely mixed-strategy equilibria.

Consider the 2×2 *Coordination game* with payoff bi-matrix

	<i>A</i>	<i>B</i>
<i>A</i>	4, 4	0, 2
<i>B</i>	2, 0	3, 3

There are two strict Nash equilibria, (A,A) and (B,B), and one mixed-strategy equilibrium. In the latter both players assign probability 0.6 to pure strategy A, rendering each player an equilibrium payoff of 2.4. This mixed equilibrium meets the requirements in the non-cooperative game-theory literature: it is robust to *all* strategy trembles and to *all* small perturbations of payoffs. However, it is not evolutionarily stable, and it is unstable in the replicator dynamics. To see the latter property, note that if the population share playing A is slightly above (below) its equilibrium value, 0.6, then pure strategy A (B) earns a higher payoff than the other, and so the replicator dynamics brings the population state towards pure strategy A (B).

B. Less cutting power against weakly dominated strategies. Weakly dominated strategies are ruled out by non-cooperative requirements, such as trembling-hand perfection. They are also ruled out by the evolutionary stability criterion.⁸ However, several researchers have pointed out that weakly dominated strategies may be weakly *dynamically* stable in such deterministic selection dynamics as the replicator dynamics. Thus, the requirement of weak dynamic stability in such dynamics lends little support to the practice of discarding weakly dominated strategies, see Binmore (1990), Samuelson (1988,1993,1994), Samuelson and Zhang (1992), and Weibull (1995).

Consider the two-player *Entry-Deterrence game* represented by the payoff bi-matrix.

	<i>yield</i>	<i>fight</i>
<i>enter</i>	2, 2	0, 0
<i>stay out</i>	1, 4	1, 4

This game can be interpreted as the stage game in Selten's (1978) *Chain-Store game*, where player 1 is a potential entrant into 2's monopoly market. The set of Nash equilibria consists of two components, the singleton set where player 1 enters and player 2 yields, and the continuum set where player 1 stays out and player 2 fights.

⁸For suppose a (pure or mixed) strategy x is evolutionarily stable and weakly dominated by some (pure or mixed) strategy y . Then y is an alternative best reply to x , and x doesn't do better against *any* strategy than y does, by weak dominance. In particular, x does not do better than y against y itself, contradicting the hypothesis that x was evolutionarily stable.

with a probability of at least 0.5. The unique subgame perfect equilibrium of the well-known extensive form representation of this game (see Figure 1 below) is that the entrant does enter and the monopolist yields. The Nash equilibrium threat of the latter to *ght* is not credible. In terms of dominance relations: the pure strategy *ght* is weakly dominated by the pure strategy *yield*.

Applying the (Taylor 2-population) replicator dynamics to this game one obtains a flow diagram as shown in Figure 1 (see also Gale, Binmore and Samuelson (1995), Weibull (1995)). The diagram shows that the evolutionary selection pressure against the strategy *ght* is weak when only a small fraction of individuals in the potential-entrant population try to enter. The latter learn that this is costly so their population fraction shrinks over time. Consequently, the selection pressure against *ght* is decreased. As shown in the diagram potential entrants may learn to stay out before all monopolists have learned not to *ght*.

Figure 1: The (Taylor 2-population) replicator dynamics in the Entry-Deterrence Game.

As a by-product we conclude that dynamic evolutionary stability does not lend strong support to backward induction. As pointed out in Gale, Binmore and Samuelson (1995) this suggests a resolution of Selten's (1978) famous chain-store paradox even in the case of a single store, and without invoking incomplete information. They also suggest that this explains the weak empirical support for backward induction found in laboratory experiments with the ultimatum game (the entry deterrence game can be viewed as a mini ultimatum game).

C. More cutting power in cheap-talk coordination games. Consider a 2×2 *Coordination game* with two strict Nash equilibria and one mixed-strategy Nash equilibrium, such as the game discussed above. While the mixed equilibrium was rejected by both static and dynamic evolutionary stability criteria, neither of the two strict Nash equilibria can be so rejected. Any strict Nash equilibrium is evolutionarily stable and dynamically stable in deterministic selection dynamics. However, Robson (1990), Wärneryd (1991), Schlag (1993, 1994) and Kim and Sobel (1995) have shown that if one extends a coordination game to include a pre-play communication stage, then evolutionary stability criteria, applied to the extended game, may reject Pareto inefficient *strict* equilibria in the underlying game, while the Nash equilibrium criterion has no cutting power at all. The possibility of rejecting strict equilibria in the underlying game is that these need not correspond to strict equilibria in the extended game.

In this literature one usually work with static criteria such as evolutionary stability and *neutral stability*, where the latter (due to Maynard Smith (1982)) is obtained from the first by weakening the strict inequality.⁹

The setting is the following: there is a symmetric and finite two-player base game to be played after a pre-play communication session. Communication takes the form of costlessly and simultaneously sent messages, one from each player. These messages are chosen from a finite set of available messages, the same set for both players. The sent messages are observed without error by both players before they choose a strategy to play in the base game. A pure strategy in this symmetric cheap-talk game thus consists of a message to send and a decision rule that prescribes a pure base-game strategy for every pair of sent messages.

Consider such a cheap-talk extension of the 2×2 *Coordination game* given above and suppose momentarily that individuals always play pure strategies. If a population plays the Pareto inefficient equilibrium (B,B) in this base game, and there exists some *unused* message in the message set, then the population is vulnerable to invasion by individuals who always send this unused message, play the Pareto efficient equilibrium (A,A) when they meet each other, and play a best reply to the (pure) strategy that the natives play when they receive this unused message. In this way the invaders earn a higher (expected) payoff than the natives, and hence play of (B,B) is not even neutrally stable.

More general results are available for the special case of 2×2 coordination games (not relying on unused messages or pure strategies). Schlag (1994) shows that those cheap-talk strategies that result in the Pareto efficient outcome together constitutes an evolutionarily stable set in the sense of Thomas (1985a). He also shows that there is only one cheap-talk strategy that is evolutionarily stable.¹⁰ When the message

⁹Each of these criteria implies a form of dynamic stability in the replicator dynamics. An evolutionarily stable strategy, viewed as a population state, is (locally) asymptotically stable (Taylor and Jonker (1978)), and a neutrally stable strategy is Lyapunov stable (Thomas (1985a) and Bomze and Weibull (1996)).

¹⁰Evolutionarily stable strategies, and evolutionarily stable sets of strategies, are dynamically

set is large, its payoff is close to the Pareto efficient outcome. Banerjee and Weibull (1997) show that, with finite message sets, the set of neutrally stable outcomes is finite and contains both strict Nash equilibrium payoffs. As the size of the message set increases, the set of neutrally stable outcomes converges to a certain countable limit set which has the Pareto efficient Nash equilibrium outcome as its unique cluster point.¹¹ Hence, evolutionary robustness, even in the weak form of neutral stability, has some cutting power when combined with (costless) preplay communication, and it slightly favors social efficiency.¹²

These results contrast sharply with those obtained with non-cooperative re-negotiations: any payoff value between the worst and best Nash equilibrium payoffs can be approximated by a (strictly) perfect Nash equilibrium in the cheap talk game when the message set is sufficiently large (Banerjee and Weibull, 1997).

5.3. Stochastic dynamic stability. It appears quite natural to model evolutionary processes as noisy: Especially the mutation process appears as inherently random. Random shocks are implicitly accounted for in deterministic selection dynamics by way of stability analysis: a stable population state is robust to isolated small perturbations of the population state. However, such stability analysis has little to say about robustness against sequences of small shocks or simultaneous small shocks that together make up a big perturbation. Such sequential or simultaneous cascades of shocks may take the population state out of one of the selection process basins of attraction, and may thus lead the population state far away. Although such cascades may be quite unlikely events in the case of statistically independent and rare mutations, this possibility changes the nature of the dynamic process in a fundamental way. Instead of being history dependent (depend on the initial population state), the process may become *ergodic*, i.e., have an asymptotic distribution that is history independent (the same for all initial population states).

This research route was pioneered by Foster and Young (1990), followed by Fudenberg and Harris (1992), Young (1993a,b), Kandori, Mailath and Rob (1993), Nöldeke and Samuelson (1993), Samuelson (1994), Kandori and Rob (1995), Binmore, Samuelson and Vaughan (1995), Bergin and Lipman (1996), Cabrales (1996), Robson and Vega-Redondo (1996) and Binmore and Samuelson (1997).

Below follows an account of Kandori, Mailath and Rob (1993) and Young (1993a), whereafter some of the other models are briefly commented.

A. Kandori-Mailath-Rob and Young. Consider the distribution of pure strategies in a finite population of individuals who are repeatedly matched in pairs to

(asymptotically) stable in the replicator dynamics, results due to Taylor and Jonker (1978) and Thomas (1985a), respectively.

¹¹We also show continuity at infinity in the sense that the limit set of neutrally stable outcomes is identical with the set of neutrally stable outcomes when the message set is countably infinite.

¹²A neutrally stable strategy is dynamically (Lyapunov) stable in the replicator dynamics, see Thomas (1985a) and Bomze and Weibull (1996).

play a symmetric two-player game. Suppose individuals change their strategy at discrete points in time, always changing to a best reply to the last period's distribution of strategies. This is a special case of the *selection process* in Kandori, Mailath and Rob (1993). In Young (1993a) individuals play a best reply to a finite *sample* from a finite record of past play. Hence, the selection process in Young's model is noisy, the more so the smaller is the sample in relation to the record of past play.¹³

Young shows that if the sample is sufficiently small in relation to the record in his selection process, i.e., the selection is sufficiently noisy, then the selection process converges almost surely in a certain class of games that he calls *weakly acyclic*. The defining property is that from every pure strategy profile in the game there exist a finite sequence of best replies, by one player at a time, that ends in a strict Nash equilibrium.¹⁴

Kandori, Mailath, Rob (1993) and Young (1993a) show that the addition of a small probability ε of mistakes, or mutations, has drastic consequences for the asymptotic outcome. Suppose every pure strategy has a positive probability of being mistakenly played, or mutated to. This is the *mutation process* in their evolutionary model. Combined with the selection process one obtains an evolutionary process that is an irreducible finite Markov chain. Since every such process has a unique invariant distribution, x_ε , to which the population state converges from *any* initial distribution, the addition of mutations makes the limit distribution of the so perturbed process unique, for any given mistake/mutation probability $\varepsilon > 0$. The authors establish the existence of the limit $x^* = \lim_{\varepsilon \rightarrow 0} x_\varepsilon$ of the sequence of such invariant distributions as ε is taken to zero, and they proceed to analyze its properties.

Consider again the 2×2 *Coordination game* discussed above. This game has two strict Nash equilibria, (A,A) and (B,B), where the first Pareto dominates the second. We now add the observation that the second *risk dominates* the first in the sense of Harsanyi and Selten (1988). The equilibrium (B,B) involves less strategic risk than (A,A) in the sense that strategy B is optimal even if the chance is only 40% that one's opponent will play B, while the chance need to be at least 60% that one's opponent will play A for A to be optimal. Hence (B,B) is preferable to (A,A) in terms of the strategic risks involved.

A best-reply selection process in discrete time and without mutations will settle in one step on any one of these two Nash equilibria.¹⁵ If initially more than 60% of the population play strategy A, then the whole population will move to strategy A in one step. If initially more than 40% play B, the population will move to B in one step. Hence, the long-run solution is history dependent. (In these simple games

¹³Note that these selection mechanisms presume more information than the earlier discussed selection dynamics: knowledge of one's payoff matrix and of past population state(s). Note also the built-in inertia: players believe that past play represents future play.

¹⁴In the absence of noise (i.e., if the sample is complete) there is a possibility that the selection process diverges - players may miscoordinate forever.

¹⁵The mixed-strategy Nash equilibrium, in which each player assigns probability .6 to pure strategy A, is unstable in such a process.

the long-run outcome is the same for the replicator dynamics as for this best-reply dynamics.)

Suppose now that each individual, having decided to play the optimal strategy, mistakenly adopts the other pure strategy with an exogenously specified probability ε (mistakes being i.i.d.). Since any population state in this way can be reached with positive probability from any other population state, the population process indeed constitutes an irreducible Markov chain. The asymptotic behavior of such a process is independent of the initial state. Instead, it is determined by the number of mistakes or mutations it takes to move the population from one of the two strict Nash equilibria to the other. The equilibrium which is most easily disrupted by mutations in this sense is given zero probability as the mutation probability ε goes to zero. The basin of attraction of strategy B (in the selection process) was above seen to be larger than that of strategy A. Therefore more mutations are needed to bring the population out of equilibrium (B,B) than out of equilibrium (A,A). As the probability ε of a mistake goes to zero, the first event becomes infinitely more likely than the second, and the population state will spend almost all time in (B,B) - despite the fact that the (B,B) is Pareto dominated by (A,A).

Note that what upsets a strict equilibrium here is not a sequence of independent small shocks but a large number of simultaneous shocks.

Bergin and Lipman (1996) point out that it is crucial for the results in these two models that the mutation probabilities go to zero at the same rate in all population states. Otherwise any invariant population distribution in the selection process can be turned into the limiting distribution when (suitably chosen state-dependent) mutation probabilities are taken to zero. The authors accordingly claim that what is missing from the models of Kandori, Mailath and Rob (1993) and Young (1993a) is a theory of mutations.

One such theory is that mutations are mistakes, where the probability of a mistake depends on the individual's *effort* to control his or her actions. Assume that there is a convex cost function of ε , with infinite marginal cost at $\varepsilon = 0$, and subtract this cost from the payoffs in the game. Assume that individuals choose their control effort optimally in each population state. Give a utility weight $\delta > 0$ to the cost term (when subtracted from payoffs), and let $\delta \rightarrow 0$, i.e., study the limit as control costs become arbitrarily small in comparison with the payoffs in the game. It turns out that in 2×2 *Coordination games* one then obtains in the limit the risk dominant Nash equilibrium, just as in the Kandori, Mailath and Rob (1993) and Young (1993a) models (van Damme and Weibull (1997)). One may argue that this theory of mutations is too rationalistic in spirit. However, the mentioned qualitative conclusion should hold even if only some individuals are (somewhat) rational in this sense while others make mistakes with the same probability in all population states.

B. Other stochastic models. Some of the other above-mentioned models in the literature on stochastic evolutionary processes favor the Pareto efficient outcome over the risk-dominant equilibrium in 2×2 coordination games. The result depends

on players' information, the matching technology, and on the selection and mutation processes. For example, Robson and Vega-Redondo (1996) show that if one changes the set-up in Kandori, Mailath and Rob (1993) so that there is random matching and imitation of successful players, then the Pareto efficient equilibrium is selected in the long run. Moreover, in their model population movements are much faster. Two mutating individuals that happen to be matched can initiate a cascade of imitators, and lead the population state to the other equilibrium. Nöldeke and Samuelson (1993) differ from the other studies in that they develop a stochastic evolutionary/learning process for games in extensive form.

6. SUMMARY AND REMAINING CHALLENGES

We have seen that evolution does not always lead to optimality, equilibrium and/or social efficiency. The economist's friend was thus at least partially justified in being sceptical about the general validity of the economist's *as if* approach.

Some desiderata concerning future research:

- Bring in institutions, and analyze market selection mechanisms. See Winter (1964, 1971), Nelson and Winter (1982), Blume and Easley (1992, 1995, 1996), Dutta (1992), Dutta and Radner (1993), Radner (1995), Nöldeke and Samuelson (1997), Vega-Redondo (1997).
- Study the evolutionary robustness of learning rules in games. Eventually, this may lead to models where individuals are endowed with intermediate cognitive capacity. See Harley (1981), Stahl (1993), Lu and Stahl (1993), Banerjee Weibull (1995), Vega-Redondo (1995), Binmore and Samuelson (1997).
- Analyze settings where instead of the same game being played over and over, *similar* games are being played. Although it is difficult to say what similarity is, robustness in this respect is needed for the validity of the *as if* approach.
- More analysis of evolutionary processes in extensive-form games. See Fudenberg and Levine (1993a,b, 1997), Nöldeke and Samuelson (1993).
- More approximation theory: We need to better understand the relations between deterministic and stochastic models. See Boylan (1992), Binmore, Samuelson and Vaughan (1995), Cabrales (1996), Börgers and Sarin (1996), Sarin (1996).
- More empirical analysis of human subjects in laboratory experiments, and in the field. See Roth and Erev (1995), Cooper and Feltovich (1996), Mookerjee and Sopher (1997), Erev and Roth (1997), Camerer (1997) and Camerer and Ho (1997).
- Search for structurally robust predictions. It appears that certain game substructures, such as sets of pure strategies closed under best replies and under weakly better replies, are robust attractors to evolutionary processes. See

Hurkens (1995), Ritzberger and Weibull (1995), Sanchirico (1996) and Young (1997a,b).

Two final comments:

Dependency on context: Some of the results in evolutionary game theory depend on the details of the selection and mutation processes involved. Is this a drawback? Perhaps not always. Maybe it is not so strange if predictions in some games depend on the context in which the game is played (a point made at several occasions by Ken Binmore). Of course, this is not a theoretically defensible position if we take the game to be a complete description of the situation at hand. However, in our search for models of economic agents with intermediate cognitive capacity - what the economist-researcher's friend asked for - why can't we view the game as only one part of the model, where the other is a description of the context, including boundedly rational processes of adaptation and learning? After all, apparently similar markets, institutions and whole economies sometimes produce very different outcomes. In some such cases it may be appropriate to consider these as different games, but in other cases we might view them as the same game played in different contexts.

Contact with behavior sciences: Recent developments in evolutionary game theory, of which some have been sketched here, and in models of learning in games, suggest new channels for communication with the other social sciences. In order to narrow down the relevant classes of selection, learning and mutation processes we need to know much more about the qualitative properties of how individuals, groups of individuals, organizations, firms, and whole societies adapt and learn over time. This in part a task for experimental game theory, but this is also an area where we could learn from the other social sciences.

REFERENCES

- [1] Akin E. (1980): Domination or equilibrium , *Mathematical Biosciences* 50, 239-250.
- [2] Aumann R. and A. Brandenburger (1995): Epistemic conditions for Nash equilibrium , *Econometrica* 63, 1161-1180.
- [3] Banerjee A. and J. Weibull (1995): Evolutionary selection and rational behavior , in Kirman A. and M. Salmon (eds.), *Learning and Rationality in Economics*, Blackwell (Oxford).
- [4] Banerjee A. and J. Weibull (1997): Neutrally stable outcomes in cheap talk coordination games , mimeo, MIT and Stockholm School of Economics.
- [5] Basu K. and J. Weibull (1991): Strategy subsets closed under rational behavior , *Economics Letters* 36, 141-146.
- [6] Benaim M. and J. Weibull (1997): Deterministic approximation of stochastic evolution , work in progress.
- [7] Bergin J. and B.L. Lipman (1996): Evolution with state-dependent mutations , *Econometrica* 64, 943-956.
- [8] Binmore K. (1987): Modeling rational players. Part I , *Economics and Philosophy* 3, 179-214.
- [9] Binmore K. (1988): Modeling rational players. Part II , *Economics and Philosophy* 4, 9-55.
- [10] Binmore K. (1990): *Essays on the Foundations of Game Theory*, Basil Blackwell (Cambridge MA).
- [11] Binmore K. and L. Samuelson (1992): Evolutionary stability in repeated games played by finite automata , *Journal of Economic Theory* 57, 278-305.
- [12] Binmore K. and L. Samuelson (1994): Drift , *European Economic Review* 38, 859-867.
- [13] Binmore K., L. Samuelson and R. Vaughan (1995): Musical chairs: Modelling noisy evolution , *Games and Economic Behavior* 11, 1-35.
- [14] Binmore K. and L. Samuelson (1997): Muddling through: Noisy equilibrium selection , *Journal of Economic Theory* 74, 235-265.
- [15] Björnerstedt J. and J. Weibull (1996): Nash equilibrium and evolution by imitation , in K. Arrow et al (eds.), *The Rational Foundations of Economic Behaviour*, Macmillan.

- [16] Blume A., Y.-G. Kim, and J. Sobel (1993): Evolutionary stability in games of communication , *Games and Economic Behavior* 5, 547-575.
- [17] Blume L. and D. Easley (1992): Evolution and market behavior , *Journal of Economic Theory* 58, 9-40.
- [18] Blume L. and D. Easley (1995): Evolution and learning in competitive markets , in Kirman A. and M. Salmon (eds.), *Learning and Rationality in Economics*, Blackwell (Oxford).
- [19] Blume L. and D. Easley (1996): Optimality and natural selection in markets , mimeo, Cornell University.
- [20] Bomze I. (1986): Non-cooperative two-person games in biology: a classification , *International Journal of Game Theory* 15, 31-57.
- [21] Bomze I. and B. Pötscher (1989): *Game Theoretical Foundations of Evolutionary Stability*, Springer Verlag (Berlin).
- [22] Bomze I. and J. Weibull (1996): Does neutral stability imply Lyapunov stability? , *Games and Economic Behavior* 11, 173-192.
- [23] Börgers T. and R. Sarin (1997): Learning through reinforcement and replicator dynamics , forthcoming, *Journal of Economic Theory*.
- [24] Boylan R. (1992): Laws of large numbers for dynamical systems with randomly matched individuals , *Journal of Economic Theory* 57, 473-504.
- [25] Bush R. and F. Mosteller (1951): A mathematical model for simple learning , *Psychological Review* 58, 313-323.
- [26] Cabrales A. and J. Sobel (1992): On the limit points of discrete selection dynamics , *Journal of Economic Theory* 57, 407-419.
- [27] Cabrales A. (1996): Stochastic replicator dynamics , Universitat Pompeu Fabra, mimeo.
- [28] Camerer C. (1997): Progress in behavioral game theory , forthcoming, *Journal of Economic Perspectives*.
- [29] Camerer C. and T.-H. Ho (1997): Experience-weighted attraction learning in games: a unifying approach , CALTECH Social Science Working Paper 1003.
- [30] Canning D. (1995): Learning and social equilibrium in large populations , in Kirman A. and M. Salmon (eds.), *Learning and Rationality in Economics*, Blackwell (Oxford).

- [31] Cooper D. and N. Feltovich (1996): Selection of learning rules: theory and experimental evidence , University of Pittsburgh WP 305.
- [32] Corradi V. and R. Sarin (1996): Continuous approximation of stochastic evolutionary game dynamics , mimeo, University of Pennsylvania and Texas A&M University.
- [33] Crawford V. (1990a): Nash equilibrium and evolutionary stability in large- and nite-population Playing the Field Models , *Journal of Theoretical Biology* 145, 83-94.
- [34] Crawford V. (1990b): An Evolutionary interpretation of Van Huyck, Battalio, and Beil s experimental results on coordination , *Games and Economic Behavior* 3, 25-59.
- [35] Cressman R. (1992): *The Stability Concept of Evolutionary Game Theory*, Springer Verlag (Berlin).
- [36] Cressman R. (1992): Evolutionarily stable sets in symmetric extensive two-person games , *Mathematical Biosciences* 108, 179-201.
- [37] Cross J. (1973): A stochastic learning model of economic behavior , *Quarterly Journal of Economics* 87, 239-266.
- [38] van Damme E. (1987): *Stability and Perfection of Nash Equilibria*, Springer Verlag (Berlin).
- [39] van Damme E. (1989): Stable equilibria and forward induction , *Journal of Economic Theory* 48, 476-509.
- [40] van Damme E. and J. Weibull (1997): Mutation rates determined by control costs , work in progress.
- [41] Dawkins R. (1976): *The Selfish Gene*, Oxford University Press (Oxford).
- [42] Dekel E. and S. Scotchmer (1992): On the evolution of optimizing behavior , *Journal of Economic Theory* 57, 392-406.
- [43] Dutta P. (1992): Bankruptcy and expected utility maximization , Rochester Center for Economic Research, WP 358.
- [44] Dutta P. and R. Radner (1993): Profit maximization and the market selection hypothesis , mimeo, AT&T Bell Laboratories.
- [45] Ellison G. (1993): Learning, local interaction and coordination , *Econometrica* 61, 1047-1072.

- [46] Ellison G. (1997): Learning from personal experience: one rational guy and the justification of myopia , *Games and Economic Behavior* 19, 180-210.
- [47] Easley D. and A. Rustichini (1996): Choice without beliefs , Center for Analytic Economics, WP #96-10, Cornell University.
- [48] Erev I. and A. Roth (1997): Modelling how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria , mimeo.
- [49] Fischer R.A. (1930): *The Genetical Theory of Natural Selection* (2nd ed. 1958), Clarendon Press (Oxford).
- [50] Foster D. and P. Young (1990), Stochastic evolutionary game dynamics , *Theoretical Population Biology* 38, 219-232.
- [51] Friedman D. (1991), Evolutionary games in economics , *Econometrica* 59, 637-666.
- [52] Friedman M. (1953), The methodology of positive economics , in Friedman M., *Essays in Positive Economics*, University of Chicago Press (Chicago).
- [53] Fudenberg D. and E. Maskin (1990): Evolution and cooperation in noisy repeated games , *AEA Papers and Proceedings* 80, No. 2, May, 274-279.
- [54] Fudenberg D. and C. Harris (1992): Evolutionary dynamics with aggregate shocks , *Journal of Economic Theory* 57, 420-441.
- [55] Fudenberg D. and D. Levine (1993a): Self-confirming equilibrium , *Econometrica* 61, 523-546.
- [56] Fudenberg D. and D. Levine (1993b): Steady-state learning and Nash equilibrium , *Econometrica* 61, 547-574.
- [57] Fudenberg D. and D. Kreps (1995): Learning in extensive-form games. I. Self-confirming equilibria , *Games and Economic Behavior* 8, 20-56.
- [58] Fudenberg D. and D. Levine (1997): *The Theory of Learning in Games*, forthcoming, MIT Press.
- [59] Gale J., K. Binmore and L. Samuelson (1995): Learning to be imperfect: the ultimatum game , *Games and Economic Behavior* 8, 56-90.
- [60] Gilboa I. and A. Matsui (1991), Social stability and equilibrium , *Econometrica* 59, No. 3, 859-867.
- [61] Hammerstein P. and R. Selten (1994): Game theory and evolutionary biology , in Aumann R. and S. Hart (eds.), *Handbook of Game Theory with Economic Applications*, Vol. 2, Elsevier Science Publ. (Amsterdam).

- [62] Harley C. (1981): Learning the evolutionarily stable strategy , *Journal of Theoretical Biology* 89, 611-633.
- [63] Harsanyi J. and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, MIT Press.
- [64] Hofbauer J. and K. Sigmund (1988): *The Theory of Evolution and Dynamical Systems*, Cambridge University Press (Cambridge UK).
- [65] Hofbauer J. and J. Weibull (1996): Evolutionary selection against dominated strategies , *Journal of Economic Theory* 71, 558-573.
- [66] Hurkens S. (1995): Learning by forgetful players , *Games and Economic Behavior* 11, 304-329.
- [67] Kalai E. and E. Lehrer (1993): Rational learning leads to Nash equilibrium , *Econometrica* 61, 1019-1046.
- [68] Kandori M., G. Mailath and R. Rob (1993): Learning, mutation, and long-run equilibria in games , *Econometrica* 61, 29-56.
- [69] Kandori M. and R. Rob (1995): Evolution of equilibria in the long run: a general theory and applications , *Journal of Economic Theory* 65, 383-414.
- [70] Kim Y.-G. and J. Sobel (1995): An evolutionary approach to pre-play communication , *Econometrica* 63, 1181-1193.
- [71] Kirman A. (1997): Evolutionary thinking in neo-classical economic theory , mimeo., G.R.E.Q.A.M..
- [72] Kohlberg E. and J.-F. Mertens (1986): On the strategic stability of equilibria , *Econometrica* 54, 1003-1037.
- [73] Leonard R. (1994): Reading Cournot, Reading Nash: The creation and stabilisation of the Nash equilibrium , *Economic Journal* 104, 492-511.
- [74] Losert V. and E. Akin (1983): Dynamics of games and genes: discrete versus continuous time , *Journal of Mathematical Biology* 17, 241-251.
- [75] Lu X. and D Stahl II (1993): An evolutionary process for a little more clever population and its implications , WP 9303, Department of Economics, University of Texas.
- [76] Mailath G. (1992): Introduction: Symposium on evolutionary game theory , *Journal of Economic Theory* 57, 259-277.
- [77] Matsui A. (1991a): Cheap-talk and cooperation in a society , *Journal of Economic Theory* 54, 245-258.

- [78] Matsui A. (1991b): Social stability and equilibrium , *Econometrica* 59, 859-867.
- [79] Matsui A. (1992): Best response dynamics and socially stable strategies , *Journal of Economic Theory* 57, 343-362.
- [80] Maynard Smith J. (1974): The theory of games and the evolution of animal conflicts , *Journal of Theoretical Biology* 47, 209-221.
- [81] Maynard Smith J. (1982): *Evolution and the Theory of Games*, Cambridge University Press (Cambridge UK).
- [82] Maynard Smith J. and G.R. Price (1973): The logic of animal conflict , *Nature* 246, 15-18.
- [83] Mookherjee D. and B. Sopher (1997): Learning and decision costs in experimental constant sum games , *Games and Economic Behavior* 19, 97-132.
- [84] Nachbar J. (1990): Evolutionary selection dynamics in games: convergence and limit properties , *International Journal of Game Theory* 19, 59-89.
- [85] Nachbar J. (1997): Prediction, optimization, and learning in repeated games , *Econometrica* 65, 275-309.
- [86] Nash J. (1950): Non-cooperative games , PhD thesis, Princeton University.
- [87] Nelson R. and S. Winter (1982): *An Evolutionary Theory of Economic Change*, Harvard University Press (Cambridge MA).
- [88] Nöldeke G. and L. Samuelson (1993): An evolutionary analysis of backward and forward induction , *Games and Economic Behavior* 5, 225-254.
- [89] Nöldeke G. and L. Samuelson (1997): A dynamic model of equilibrium selection in signaling markets , *Journal of Economic Theory* 73, 118-156.
- [90] Radner R. (1995): Economic Survival , *Nancy L. Schwartz Lecture*, Kellogg School of Management, Northwestern University.
- [91] Ritzberger K. and J. Weibull (1995): Evolutionary selection in normal-form games , *Econometrica* 63, 1371-1400.
- [92] Robson A.J. (1990): Efficiency in evolutionary games: Darwin, Nash and the secret handshake , *Journal of Theoretical Biology* 144, 379-396.
- [93] Robson A. and F. Vega-Redondo (1996): Efficient equilibrium selection in evolutionary games with random matching , *Journal of Economic Theory* 70, 65-92.

- [94] Roth A. and I. Erev (1995): Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term , *Games and Economic Behavior*, 164-212.
- [95] Samuelson L. (1988): Evolutionary foundations of solution concepts for finite, two-player normal-form games , in M. Vardi (ed.), *Theoretical Aspects of Reasoning about Knowledge*, Morgan Kaufmann Publ. (San Mateo, CA).
- [96] Samuelson L. (1991): Limit evolutionarily stable strategies in two-player, normal form games , *Games and Economic Behavior* 3, 110-128.
- [97] Samuelson L. and J. Zhang (1992): Evolutionary stability in asymmetric games , *Journal of Economic Theory* 57, 363-391.
- [98] Samuelson L. (1993): Does evolution eliminate dominated strategies? , in Binmore K., A. Kirman and P. Tani (eds.), *Frontiers of Game Theory*, MIT Press (Cambridge, MA).
- [99] Samuelson L. (1994): Stochastic stability in games with alternative best replies , *Journal of Economic Theory* 64, 35-65.
- [100] Samuelson L. (1997): *Evolutionary Games and Equilibrium Selection*, MIT Press (Cambridge MA).
- [101] Sanchirico C. (1996): A probabilistic model of learning in games , *Econometrica* 64, 1375-1393.
- [102] Schlag K. (1993a): Cheap talk and evolutionary dynamics , Discussion Paper B-242, University of Bonn.
- [103] Schlag K. (1993b): Dynamic stability in the repeated prisoners dilemma played by finite automata , Department of Economics, Discussion Paper B-243, University of Bonn.
- [104] Schlag K. (1996): Which one should I imitate? , Discussion Paper B-365, University of Bonn.
- [105] Schlag K. (1997): Why imitate, and if so, how? A bounded rational approach to multi-armed bandits , forthcoming, *Journal of Economic Theory*.
- [106] Schuster P. and K. Sigmund (1983): Replicator dynamics , *Journal of Theoretical Biology* 100, 533-538.
- [107] Selten R. (1975): Re-examination of the perfectness concept for equilibrium points in extensive games , *International Journal of Game Theory* 4, 25-55.
- [108] Selten R. (1978): The chain store paradox , *Theory and Decision* 9, 127-159.

- [109] Selten R. (1980): A note on evolutionarily stable strategies in asymmetric animal conflicts , *Journal of Theoretical Biology* 84, 93-101.
- [110] Selten R. (1983): Evolutionary stability in extensive-form two-person games , *Mathematical Social Sciences* 5, 269-363.
- [111] Selten R. (1988): Evolutionary stability in extensive-form two-person games - correction and further development , *Mathematical Social Sciences* 16, 223-266.
- [112] Selten R. (1991): Evolution, learning, and economic behavior , *Games and Economic Behavior* 3, 3-24.
- [113] Smith A.: *An Inquiry into the Nature and Causes of the Wealth of Nations*, reprinted in *The Glasgow Edition of the Works and Correspondence of Adam Smith*, Liberty Classics 1981 (Indianapolis).
- [114] Stahl D. (1993): Evolution of smart_n players , *Games and Economic Behavior* 5, 604-614.
- [115] Swinkels J. (1992a): Evolutionary stability with equilibrium entrants , *Journal of Economic Theory* 57, 306-332.
- [116] Swinkels J. (1992b): Evolution and Strategic Stability: from Maynard Smith to Kohlberg and Mertens , *Journal of Economic Theory* 57, 333-342.
- [117] Swinkels J. (1993): Adjustment dynamics and rational play in games , *Games and Economic Behavior* 5, 455-484.
- [118] Tan T. and S.R. Werlang (1988): The Bayesian foundations of solution concepts of games , *Journal of Economic Theory* 45, 370-391.
- [119] Taylor P. (1979), Evolutionarily stable strategies with two types of player , *Journal of Applied Probability* 16, 76-83.
- [120] Taylor P. and L. Jonker (1978), Evolutionary stable strategies and game dynamics , *Mathematical Biosciences* 40, 145-156.
- [121] Thomas B. (1985a): On evolutionarily stable sets , *Journal of Mathematical Biology* 22, 105-115.
- [122] Thomas B. (1985b): Evolutionarily stable sets in mixed-strategist models , *Theoretical Population Biology* 28, 332-341.
- [123] Thorndyke, E.L. (1898): Animal intelligence: An experimental study of the associative process in animals , *Psychological Monographs* 2.
- [124] Ullmann-Margalit E. (1978): Invisible-hand explanations , *Synthese* 39, 263-291.

- [125] V. Bhaskar (1992): Noisy communication and the evolution of cooperation , mimeo, Delhi School of Economics.
- [126] Van Huyck J., Cook J. and R. Battalio (1994): Selection dynamics, asymptotic stability, and adaptive behavior , *Journal of Political Economy* 102, 975-1005.
- [127] Vega-Redondo, F. (1995): Expectations, drift and volatility in evolutionary games , *Games and Economic Behavior* 11, 391-412.
- [128] Vega-Redondo, F. (1996): *Evolution, Games, and Economic Behavior*, Oxford University Press (Oxford).
- [129] Vega-Redondo F. (1997): The evolution of Walrasian behavior , *Econometrica* 65, 375-384.
- [130] Wärneryd K. (1991): Evolutionary stability in unanimity games with cheap talk , *Economics Letters* 36, 375-378.
- [131] Weibull J. (1995): *Evolutionary Game Theory*, MIT Press (Cambridge, MA).
- [132] Weissing F. (1991): Evolutionary stability and dynamic stability in a class of evolutionary normal form games , in Selten R. (ed.), *Game Equilibrium Models I*, Springer Verlag (Berlin).
- [133] Winter S. (1964): Economic natural selection and the theory of the firm , *Yale Economic Essays*, Vol. 4, No.1.
- [134] Winter S. (1971): Satisficing, selection, and the innovating remnant , *Quarterly Journal of Economics* 85, 237-261.
- [135] Young P. (1993a): Evolution of conventions , *Econometrica* 61, 57-84.
- [136] Young P. (1993b): An evolutionary model of bargaining , *Journal of Economic Theory* 59, 145-168.
- [137] Young P. (1995): Equilibrium selection through adaptation and experimentation , in Kirman A. and M. Salmon (eds.), *Learning and Rationality in Economics*, Blackwell (Oxford).
- [138] Young P. (1997a): Individual learning and social rationality , forthcoming, *European Economic Review*, Papers and Proceedings.
- [139] Young P. (1997b): *Individual Strategy and Social Structure*, forthcoming, Princeton University Press (Princeton).