



ELSEVIER

European Economic Review 38 (1994) 868-881

EUROPEAN
ECONOMIC
REVIEW

The 'as if' approach to game theory: Three positive results and four obstacles

Jörgen W. Weibull*

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

Abstract

The usual justifications of non-cooperative solution criteria are rationalistic in nature, building on individual rationality and expectation-coordination postulates. In contrast, the 'as if' approach to game-theoretic rationality is evolutionary, arguing that even if strategically interacting agents do not meet these epistemic conditions, their long-run aggregate behavior will nevertheless conform with them because of the workings of biological or social selection processes. The present essay discusses three implications of evolutionary selection dynamics in favour of the 'as if' paradigm, and four potential obstacles to the general validity of these implications.

Key words: Game theory; Evolution; Rationality

JEL classification: C72

1. Introduction

Many economic models rely on solution concepts from non-cooperative game theory such as Nash equilibrium and its refinements. Game theorists have investigated the decision-theoretic foundations of these constructs. It has been shown that strong epistemic conditions are needed to justify Nash equilibrium behavior, conditions which not only involve 'rationality' of the interacting agents but also 'consistency' between their beliefs about each

other.¹ An alternative to this rationalistic approach is to instead ask if evolutionary selection processes induce a tendency towards behavior which conforms with Nash equilibrium or other non-cooperative solution criteria. The processes in question can be some form of biological or social selection, or selection due to market competition. If the behaviors selected for by such evolutionary processes indeed do meet non-cooperative solution criteria, one may claim that, even if the interacting agents are not 'rational' and have 'coordinated' beliefs, they nevertheless behave, at least in the aggregate and in the long-run, *as if* they did meet these epistemic assumptions. In the context of perfectly competitive markets (hence non-strategic environments), the 'as if' approach was earlier advocated by Alchian (1950), Friedman (1953) and Winter (1971).

Recent developments in evolutionary game theory have established some positive results on the implications of evolutionary selection for 'game theoretic rationality'. One then imagines that individuals are randomly drawn from a large population, or from distinct large populations, to play the game in question, over and over again, and studies how the population distribution across pure strategies changes over time according to some evolutionary selection process. First, in a fairly wide class of such evolutionary selection dynamics, dynamic stability implies Nash equilibrium. Secondly, in the same class of models, dynamic convergence implies Nash equilibrium, granted all pure strategies of the game are present in the initial population. Thirdly, strictly dominated pure strategies are wiped out in the long run, even if the evolutionary dynamic path doesn't converge, again granted all pure strategies are initially present.

These positive results may sound comforting: not only does dynamic evolutionary selection justify the use of Nash equilibrium and the elimination of strictly dominated strategies; these results even suggest that the analyst should go beyond Nash equilibrium, viz. to those Nash equilibria which are robust to evolutionary selection pressures. However, this happy picture is somewhat clouded by certain obstacles to their general validity. The purpose of this essay is to discuss the above-mentioned three positive results in relation to four such obstacles.

The first obstacle, below labelled 'temporal lumpiness', refers to the volatility that may arise when large shares of the population simultaneously change behaviors. The second obstacle is related to the nature of the selection mechanism: if this is not biological but social or market-performance based, are the above-mentioned positive results then valid? The third obstacle is related to the 'richness' of the 'menu' of behaviors over which the evolutionary selection process takes place: what if this 'menu' contains some 'sophisticated' behaviors alongside more 'simplicistic' ones? The

* This paper summarizes results from joint research with Abhijit Banerjee, Jonas Björnerstedt, Martin Dufwenberg, Peter Norman and Klaus Ritzberger, see references in text. I thank Jonas Björnerstedt, Eric van Damme, Henrik Horn and Johan Stenck for helpful comments to an earlier version of this paper. This project was sponsored by the Industrial Institute for Economic and Social Research, Stockholm, Sweden.

¹ See e.g. Tan and Werlang (1988), Aumann and Brandenburger (1991).

fourth obstacle has to do with the multiplicity of populations taking part in the interaction: do the positive results still apply if the dynamics is applied to interactions between individuals from distinct populations?

2. Positive bench-mark results

The basic dynamic model of evolutionary selection is the so-called continuous-time single-population replicator dynamics. In this dynamics, one imagines that individuals in a large population interact pairwise, and the interaction takes the form of a finite and symmetric two-player game. Pairs of individuals are randomly matched to play the game, and payoffs represent fitness, i.e., the expected number of offspring. Each individual is genetically or otherwise 'programmed' to always play a certain pure strategy. Hence, at each instant, the population can be divided into sub-populations, one for each pure strategy i of the game. The non-negative population shares x_i sum up to one, so the vector x of population shares, the population state, is formally identical with a mixed strategy in the game in question.

It is assumed that each offspring inherits its (single) parent's strategy. Consequently, the growth-rate of each sub-population is its strategy's current expected payoff, implying that the growth rate of the population share x_i programmed to pure strategy i equals the difference between its current payoff $u_i(x)$ and the current population average payoff $\bar{u}(x)$. This is the so-called replicator dynamics (Taylor and Jonker, 1978):

$$\dot{x}_i(t) = [u_i(x(t)) - \bar{u}(x(t))]x_i(t) \quad \text{for } t \geq 0, \quad (1)$$

where the dot signifies the time derivative. Note that the first of these two payoffs, $u_i(x)$, equals the expected payoff to pure strategy i when played against mixed strategy x , and the second, $\bar{u}(x)$, equals the expected payoff to mixed strategy x when played against itself. A strategy pair, or profile, (x, y) constitutes a Nash equilibrium if x is a best reply to y and y is a best reply to x .

The above-mentioned three positive results on the connection between evolutionary selection dynamics and game-theoretic rationality apply to this particular evolutionary selection dynamics. More exactly:

- If a population state x^* is Lyapunov stable in (1), i.e., such that solution trajectories starting near x^* remain forever near x^* , then (x^*, x^*) constitutes a Nash equilibrium (Bomze, 1986).
- If (a) the initial population state $x(0)$ is interior, i.e., all population shares

$x_i(0)$ are positive, and (b) the state $x(t)$ converges to some state x^* as time t goes towards plus infinity, then (x^*, x^*) constitutes a Nash equilibrium (Nachbar, 1990).

- If (a) the initial population state $x(0)$ is interior, and (b) a pure strategy i is iteratively strictly dominated, then its population share $x_i(t)$ converges to zero as time t goes towards plus infinity (Samuelson and Zhang, 1992).

At a first glance, the third positive result, that strictly dominated strategies vanish, may appear obvious. However, it is not. For even a strictly dominated strategy can earn more than average in some population states, so the sub-population programmed to such a strategy may, at least temporarily, increase over time in the replicator dynamics.

3. Temporal lumpiness

In some biological models, evolutionary selection is modeled in discrete time where each time period represents a generation. If for each $t = 0, 1, 2, \dots$, $u_i[x(t)] \geq 0$ is the expected number of offspring to an individual in generation t programmed to pure strategy i , and $\bar{u}[x(t)] > 0$ is the average number of offspring in that generation, then the following discrete-time replicator dynamics results:

$$x_i(t+1) = \frac{u_i[x(t)]}{\bar{u}[x(t)]} x_i(t) \quad \text{for } t = 0, 1, 2, \dots \quad (2)$$

As shown by Nachbar (1990), the first two positive results above do hold also in this dynamics, i.e., dynamic stability, and likewise dynamic interior convergence, implies Nash equilibrium play. Note, however, that the sets of stable population states and converging interior dynamic paths, respectively, may differ from the corresponding sets in the continuous-time dynamics (1). Most likely, these sets are smaller in discrete time – due to the possibility of 'over-shooting' – so the predictive power of these two links between evolution and Nash equilibrium may be weaker in (2) than in (1).

Nachbar also established the following weaker version of the third positive result: if the iterated elimination of pure strategies which are strictly dominated by some pure strategy results in a single pure strategy, then the dynamics (2) converges from any interior initial population state towards the population state in which all individuals use that strategy. That the stronger result for the continuous-time replicator dynamics does not carry over to (2), was shown, by way of a counter-example, by Dekel and Scotchmer (1992). In

their example, a pure strategy which is strictly dominated by a mixed strategy, but by no pure strategy, survives in the replicator dynamics (2).

As pointed out by Cabrales and Sobel (1992), the reason why this counter-example works is not temporal discreteness per se, but its special form (2). Indeed, they show that if the time discretization is made sufficiently 'fine', so that only a small batch of individuals change strategy each time, then the third positive result is restored: all strictly dominated strategies are wiped out along all interior solution paths. Björnerstedt et al. (1993) establish this for a class of games containing the Dekel-Scotchmer example in an overlapping-generations (OLG) model of biological reproduction in which a randomly drawn population share δ reproduces simultaneously, at times $t=0, \delta, 2\delta, \dots$. We show that if the share δ is sufficiently large, then virtually all solution trajectories diverge and the population share programmed to the strictly dominated strategy does not converge to zero, while if δ is sufficiently small, all interior solution trajectories converge to the unique Nash equilibrium strategy of the game.

Note also that these problems for the 'as if' paradigm do not arise at all in continuous time, even if one assumes temporal 'lumpiness' in the sense that the over-all rate at which reproduction takes place varies strongly over time. To see this, suppose each right-hand side in (1) is multiplied by some positive factor $\lambda(t)$ which varies (Lipschitz) continuously over time t . For instance, the rate $\lambda(t)$ may be very large when t is a non-negative integer and close to zero most of the time in between, in this sense mimicking the discrete-time replication model (2). However, irrespective how $\lambda(t)$ varies with time t , the associated continuous-time solution curves in the mixed-strategy space are identical with those of (1), since the effect of λ is equivalent to a mere change of time scale.

Temporal lumpiness thus appears to be largely an artificial obstacle to the 'as if' paradigm.

4. Non-biological selection

For applications to economics, biological reproduction is not always a compelling parable for how behaviors (pure strategies) spread in a population, for such modeling purposes imitation of successful behaviors, boundedly rational search for optimal behaviors, market selection against unsuccessful behaviors etc., appear more relevant. A few simple models of such social transmission mechanisms will here be sketched.

These models build on the notion that individuals live forever but now and then review their strategy choice. Let $q_i(x)$ be the time rate at which an individual who is currently using strategy i reviews her strategy choice when

the population state is x , and let $p_{ij}(x)$ be the probability that a reviewing individual will adopt strategy j . In particular, $p_{ii}(x)$ is the probability that she will choose not to change strategy. In terms of expected values, the resulting outflow per time unit from sub-population i will then be $x_i q_i(x) \sum_{j \neq i} p_{ij}(x)$ and the inflow will be $\sum_{j \neq i} x_j q_j(x) p_{ji}(x)$. Rearranging terms, one obtains the following dynamics in an infinitely large population. Such a system of differential equations induces a well-defined dynamics if the involved functions, q and p , are Lipschitz continuous.

$$\dot{x}_i = \sum_j x_j q_j(x) p_{ji}(x) - q_i(x) x_i \tag{3}$$

If each choice probability $p_{ij}(x)$ is an increasing function of x_j , the population share of the potential 'target' strategy j , then the transmission mechanism could be said to be *imitative*: more prevalent strategies, ceteris paribus, are more likely to be adopted. The notion that individuals with less successful strategies are more inclined to review their strategy choice can be formalized by having all review rate $q_i(x)$ non-increasing in the player's current expected payoff $u_i(x)$. Likewise, the notion that more successful behaviors are more prone to being initiated can be formalized by letting each choice probability $p_{ij}(x)$ be non-decreasing in the expected payoff $u_j(x)$ of the potential target strategy j .

A simple example of such payoff-oriented imitation is obtained if one sets all review rates equal and constant, say, $q_i(x) = 1$ for all i and x , and lets the switching probabilities $p_{ij}(x)$ be increasing in the population share x_j and in the corresponding payoff $u_j(x)$, as follows: $p_{ij}(x) = x_j \exp[\sigma u_j(x)] / [\sum_k x_k \exp[\sigma u_k(x)]]$ for some positive scalar σ . The higher the value of the parameter σ is, the stronger is the tendency of reviewing individuals to adopt currently optimal strategies (granted some such strategy is present). Indeed, as σ tends to plus infinity, the probability $p_{ij}(x)$ associated with any pure strategy j which is *not* a best reply to the current state x tends to zero. For any fixed $\sigma > 0$, the induced selection dynamics (4) is

$$\dot{x}_i = \left[\frac{e^{\sigma u_i(x)}}{\sum_k x_k e^{\sigma u_k(x)}} - 1 \right] x_i \tag{4}$$

Such a dynamics is *monotonic* with respect to payoffs in the sense that a pure strategy with a higher current payoff has a higher current growth rate.

It is known that the first two positive 'as if' results carry over to all such dynamics, while the third result, on strictly dominated strategies, has to be weakened in the same way as in the case of the discrete-time replicator dynamics (2) (Nachbar, 1990; Samuelson and Zhang, 1992).

An alternative payoff-monotonic dynamics, with a flavor of market selection, is suggested in Björnerstedt (1993). Here one can imagine some outside observer, a 'bank', who now and then samples a small fraction of the population of firms', observes their payoffs, or 'profits', with a small random error, and abandons all firms who use a strategy which has the lowest observed payoff. An interesting feature of this (continuous time) dynamics, shown in Björnerstedt (1993), is that a strictly dominated strategy may survive in the long run if it is not strictly dominated by any *pure* strategy. The reason is that the strategy may yield a low but fairly steady payoff, and hence may frequently avoid being the worst performing strategy.

While some positive 'as if' results apply to some of the mentioned selection dynamics, more research on specific social and market-performance based selection mechanisms is needed before we know which non-biological selection dynamics lend support to the 'as if' paradigm.

5. Rich menu of behaviors

Robson (1990), Dekel and Scotchmer (1992), Banerjee and Weibull (1991, 1992, 1993) and Stahl (1992) have pointed out that the 'as if' approach to game-theoretic rationality may be in serious trouble if the population is heterogeneous with respect to 'sophistication'. In particular, *none* of the mentioned three positive results on game-theoretic rationality then need hold even under the standard replicator dynamics (1).

For illustration, consider the following thought-experiment: a biologist studies evolutionary selection in a large population of 'programmed' individuals who are randomly matched to play some finite and symmetric two-player game. Without the biologist's knowledge, an economist comes by and injects a few individuals of the species *homo oeconomicus* in the population. These 'new' agents are endowed with the capacity to recognize the type of agent they meet. In particular, they correctly predict at each encounter the strategy to be used by each of their biological 'opponents', and play a best reply against it. They also recognize each other, and then play some iteratively strictly undominated strategy. In contrast, the original 'biological' agents go on using their pre-programmed pure strategies as before. What will happen?

As shown in Banerjee and Weibull (1991, 1992), biological agents programmed to a strictly dominated strategy can then survive in the long run,

they can even be part of an asymptotically stable population state.² For if programmed to a strictly dominated pure strategy which earns a payoff above the payoff that its best reply earns in such matchings, these biologically programmed agents can do quite well. In terms of non-cooperative game theory: they benefit, in encounters with *homo oeconomicus*, from a commitment to an aggressive strategy. Suitably re-formulated as a game between programmed players, where *homo oeconomicus* is identified with an artificial, additional pure strategy, the argument is a standard application of the replicator dynamics (1).

More generally, the above example can be embedded in a richer model framework, and then it turns out that the three initial positive benchmark results on evolution and rationality are restored, albeit in a somewhat weaker and more complex form. In a sense, one then embeds the example in a model in which biological evolutionary selection operates at the level of 'decision procedures' rather than at the level of 'decisions' or 'actions'.

First, suppose we have, as usual, a population of individuals who are randomly matched to play a symmetric two-player game. However, now each individual has one of finitely many observable physical traits, or, equivalently, at each matching sends one and the same of finitely many possible 'signals' or 'messages' before playing the game. In such a richer world, individuals can be genetically or otherwise 'programmed' to rules prescribing which strategy to play depending on the signal received from their 'opponent'. Mathematically, such a behavior rule is a function f from the finite set C of possible signals (traits, messages) to the pure-strategy set of the game in question: if my rule is f and my opponent's signal (trait, message) is c , then I use pure strategy $i = f(c)$, etc.

Applied to the above thought-experiment, the set C could consist of one distinct signal for each type of biological agent (one for each pure strategy in the game), plus one signal for *homo oeconomicus*. A biological agent then is programmed to the *constant* rule which prescribes that his pure strategy i be used in all encounters, and each *homo oeconomicus* is programmed to the 'best reply' rule which prescribes a best reply to each signal sent by a biologically programmed agent and an iteratively undominated strategy to the *homo oeconomicus* signal.

Secondly, given such a set C of signals, messages or observable traits, a model of evolutionary selection should allow for *all* possible choice rules, rather than, as in the above example, only a few special rules. But once this generalization is made, one ends up in an equivalent 'meta game', in which a pure strategy is a pair (c, f) . Formally, such a meta-game is identical with a cheap-talk game with (pre-play) message space C . Moreover, the three

² A population state x is asymptotically stable in a given dynamics if solution trajectories starting near x remain forever near x and converge over time to x .

above-mentioned positive results on evolutionary selection and rational play apply to the *meta-game*. Evolutionary selection thus has the 'usual' positive implications when applied at the level of 'decision procedures' or 'cognitive designs'. The implications of this *cognitive evolution* for the resulting 'decisions' or 'actions', i.e., the pure strategies used in the base game, are a bit less direct.

Expressed in terms of the underlying 'base game', the main results are the following (for details, see Baneijee and Weibull (1993)). First, each Lyapunov stable population state in the replicator dynamics (1), applied to the meta game, corresponds to some *convex combination* of symmetric and/or asymmetric Nash equilibria of the base game. Secondly, if an interior solution trajectory to the replicator dynamics converges, then its limit state is again some convex combination of Nash equilibria. In both of these results, there thus is a possibility that aggregate behavior ('actions') does *not* conform with play of a single symmetric Nash equilibrium. However, in such cases it is as if the sub-aggregate of all individuals with the same trait (signal, message) play some symmetric or asymmetric base-game Nash equilibrium with every other such sub-aggregate of individuals, and each such sub-aggregate plays some symmetric Nash equilibrium with itself. Thirdly, if one observes the relative *frequency* with which some strictly dominated base-game strategy is used in the matchings, one will find that this frequency converges to zero over time, along *any* interior dynamic solution trajectory to the replicator dynamics for the meta-game. In this sense, evolution does select 'rational play' even in this more complex setting of strategy choice rules. (See also the hierarchical model in Stahl (1992))

Why did this not happen in the above laboratory experiment?³ The answer is simply that the initial population state was not interior. With k pure strategies and $k+1$ types, there are k^{k+1} possible strategy-choice rules, while we only allowed for $k+1$.

6. Multiple populations

A fourth potential difficulty for the 'as if' approach to non-cooperative game theory arises if the interacting individuals do not all come from one and the same population. For instance, suppose that in a symmetric or asymmetric two-player game the row-player is drawn from one population (of, e.g., 'buyers') and the column-player from another (e.g., 'sellers'). Indeed, this identification of each player 'role' or 'position' with a distinct population was suggested already by Nash when he first developed his equilibrium concept. I am grateful to Harold Kuhn for providing Leonard (1993) in which the following passage from Nash's Ph.D. dissertation is quoted.

"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 participants for each position is 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" (Nash, 1950, p. 21).

In an evolutionary setting for an n -player game, a population state x is formally identical with a mixed-strategy profile, each component x^k of which is a mixed strategy for player (position) k , where $k=1,\dots,n$. Just as in the single-population replicator dynamics, its n -population counterpart assumes that each individual in the population associated with player position k is 'programmed' to one of the pure strategies available to that player position. The n -population continuous-time replicator dynamics becomes

$$\dot{x}^k(t) = [u_k(x(t)) - \bar{u}_k(x(t))]x^k(t) \quad \text{for } t \geq 0, \quad (5)$$

where $u_k(x)$ is the payoff to player k when using pure strategy i against mixed-strategy profile x , and $\bar{u}_k(x)$ is the same player's payoff when x is played. In terms of player-populations, $u_k(x)$ is the payoff to all those individuals in the k th player population who are programmed to strategy i , and $\bar{u}_k(x)$ is the average payoff in that player population.³

It is easily shown that our three positive results for the single-population continuous-time replicator dynamics (1) all carry over to their n -population versions (5). On the surface, thus, the mere multiplicity of populations seems to be no obstacle, per se, to the 'as if' paradigm. However, although dynamic stability implies Nash equilibrium, few population states are stable in the multi-population setting. More exactly: a population state x is asymptotically stable in the multi-population replicator dynamics (4) of an n -player game if and only if x is a strict Nash equilibrium, i.e., if and only if each strategy x_i is the *unique* best reply to x (Ritzberger and Vogelsberger (1990); see also Hofbauer and Sigmund (1988), and Samuelson and Zhang (1992)). Since many games of interest for economics lack strict Nash equilibria, this is a serious problem for the 'as if' justification of Nash equilibrium play. In contrast, the predictive power of the result on the elimination of strictly

³ Just as in the single-population setting, one can of course elaborate alternative dynamics based on social or market-oriented selection. However, we here try to isolate the role of population multiplicity per se, and hence stick to the replicator dynamics.

dominated strategies is unaffected by this observation, since it presupposes neither stability nor convergence.

One way to restore asymptotic stability is to consider sets X of population states, i.e., mixed strategy profiles, rather than individual population states x .⁴ And indeed, Swinkels (1993) establishes a result which implies that if a closed and convex set X is asymptotically stable in the replicator dynamics (5), then X contains a set Y of Nash equilibria. He also shows that there always exists such a subset Y which is *strategically stable* in the sense of Kohlberg and Mertens (1986).⁵ This is one of the most stringent refinements in non-cooperative game theory, so there is indeed a link, in terms of set inclusion, between set-wise dynamic evolutionary stability non-cooperative set-wise robustness. The predictive usefulness of this link in part depends on whether one can operationally identify asymptotically stable sets X , and whether these sets X are 'large' or 'small'.

It turns out that the above negative result on point-wise asymptotic stability has a set-wise counterpart which reduces the class of 'candidate' sets X . For by the same technique as is used to establish the above result, one can show that no *interior* closed set X is asymptotically stable in the replicator dynamics (see Ritzberger and Weibull, 1993).

Moreover, for a certain class of non-interior sets X , there is a simple necessary and sufficient condition for their asymptotic stability in the replicator dynamics (5). This condition can be expressed in terms of the so-called *better-reply correspondence* γ (Ritzberger and Weibull, 1993), the correspondence which for each player k maps any mixed-strategy profile x to those pure strategies i for the player which give him at least the payoff he gets under x . Formally:⁶

$$\gamma^k(x) = \{i: u_{ik}(x) \geq \bar{u}_k(x)\}. \quad (6)$$

The class of subsets X in question are the (finitely many) 'subspaces' of the space of mixed-strategy profiles which one obtains by restricting each player's mixed strategies to a fixed subset of his pure strategy set. Call a subspace X *closed* under the better reply correspondence if every strategy

⁴ A closed set X is called *asymptotically stable* if solution trajectories starting near X remain forever near to X and converge over time to X . In particular, the set of all population states is always asymptotically stable.

⁵ A strategically stable set is a minimal closed and nonempty set Y of Nash equilibria with the property that all slightly perturbed versions of the game, in the sense of 'trembling hand' disturbances of strategies, have some Nash equilibrium near to Y .

⁶ Cf. the usual (pure strategy) *better-reply* correspondence β^k which assigns to any strategy profile x those pure strategies i which are optimal for player k against x : $\beta^k(x) = \{i: u_{ik}(x) \geq u_{jk}(x) \forall j\}$.

profile x in X is mapped by γ to pure strategies in the subsets defining (spanning) X .

The characterization result is that a subspace X is asymptotically stable in the multi-population replicator dynamics (5) if and only if it is closed under γ (Ritzberger and Weibull, 1993). Hence, even if not all agents in the population use pure strategies in the subset, eventually they will, granted sufficiently few of them initially used other strategies. Since every game possesses at least one subspace which is closed under γ , and there are finitely many such subspaces, there moreover exists at least one *minimal* subspace X with this property.

In some games, the only minimal subspace closed under γ is the full mixed-strategy space itself, and hence the present set-valued characterization has no 'cutting power'. In other games, some such minimal subspaces contain few pure strategies for each player, and the result's cutting power is accordingly stronger. However, irrespective of this, Swinkels' (1993) conditions are met so any subspace X which is closed under γ contains a subset Y which is *strategically stable* in the sense of Kohlberg and Mertens (1986). The set-valued approach thus restores a positive but only set-inclusive implication for game-theoretic rationality.

7. Conclusions and directions for further research

The first studied obstacle to the 'as if' justification of game-theoretic rationality, temporal 'lumpiness', was argued to be essentially artificial. The second obstacle, that the selection dynamics may be social or market-performance based rather than biological, was seen to require more research. The third obstacle, that selection may operate over 'decision rules' or 'cognitive designs', was shown to weaken the implications for Nash equilibrium play towards convex combinations thereof. The fourth obstacle, that the interaction may involve individuals from distinct populations, was seen to weaken the evolutionary foundation of Nash equilibrium behavior towards set-valued inclusions.

Evidently, these are not the only obstacles or complications which have to be faced before one can claim that there is a robust evolutionary foundation for game-theoretic rationality. Among other relevant obstacles, one can mention the random-matching, large population format of the evolutionary set-up. For many applications to economics, this is not so natural a framework. Moreover, one may ask what are the implications from asymptotic, i.e., very 'long-run' properties such as those discussed above, for predictions in the 'medium' term. Also, such elements as 'mutations' and 'mistakes' are treated only implicitly in the discussed models, viz. in terms of

dynamic stability criteria. Explicit stochastic models of 'evolutionary noise', and for matchings and payoffs, are needed. Finally, the strategic interaction in question was here represented by a game in normal form, and one may ask if evolutionarily relevant aspects of the associated extensive forms are well captured in the normal form.

Several of these additional complications are currently being researched, and some results have already been obtained, see e.g. Binnmore and Samuelson (1992), Kandori et al. (1993), Nöldeke and Samuelson (1993), and Young (1993), to mention a few examples of fine research along some of these lines.

References

- Alchian, A., 1950. Uncertainty, evolution and economic theory. *Journal of Political Economy* 58, 211-222.
- Aumann, R. and A. Brandenburger, 1991. Epistemic conditions for Nash equilibrium. *Mimeo.* (Hebrew University, Jerusalem).
- Banerjee, A. and J. Weibull, 1991. Evolutionary selection and rational behavior. in: A. Kirman and M. Salmon, eds., *Rationality and learning in economics* (Blackwell, Oxford) Forthcoming.
- Banerjee, A. and J. Weibull, 1992. Evolution and rationality: Some recent game-theoretic results. *Proceedings from the Tenth World Congress of the International Economic Association* (Blackwell, Oxford) Forthcoming.
- Banerjee, A. and J. Weibull, 1993. Evolutionary selection with discriminating players. WP 1637 (Department of Economics, Harvard University, Cambridge, MA).
- Binnmore, K. and L. Samuelson, 1992. Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory* 57, 278-305.
- Björnerstedt, J., 1993. Experimentation, imitation and evolutionary dynamics. *Mimeo.* (Department of Economics, Stockholm University, Stockholm).
- Björnerstedt, J., M. Dutwenberg, P. Norman and J. Weibull, 1993. Evolutionary selection dynamics and irrational survivors. *Mimeo.* (Department of Economics, Stockholm University, Stockholm).
- Bomze, I., 1986. Non-cooperative two-person games in biology: A classification. *International Journal of Game Theory* 15, 31-57.
- Cabrates, A. and J. Sobel, 1992. On the limit points of discrete selection dynamics. *Journal of Economic Theory* 57, 407-419.
- 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: M. Friedman, *Essays in positive economics* (University of Chicago Press, Chicago, IL).
- Hofbauer, J. and K. Sigmund, 1988. *The theory of evolution and dynamical systems* (Cambridge University Press, Cambridge).
- Kandori, M., G. Mailath and R. Rob, 1993. Learning, mutation, and long-run equilibria in games. *Econometrica* 61, 29-56.
- Kohlberg, E. and J.-F. Mertens, 1986. On the strategic stability of equilibria. *Econometrica* 54, 1003-1037.
- Leonard, R., 1993. *Reading Cournot, reading Nash*. *Mimeo.* (University of Québec at Montréal, Montréal).
- Nachbar, J., 1990. 'Evolutionary' selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory* 19, 59-89.
- Nash, J., 1950. Non-cooperative games. Ph.D. dissertation (Princeton University, Princeton, NJ).
- Nöldeke, G. and L. Samuelson, 1993. An evolutionary analysis of forward and backward induction. *Games and Economic Behavior* 5, 425-454.
- Ritzberger, K. and K. Vogelsberger, 1990. The Nash field. RR 263 (Institute for Advanced Studies, Vienna).
- Ritzberger, K. and J.W. Weibull, 1993. Evolutionary selection in normal-form games. WP 383 (The Industrial Institute for Economic and Social Research, Stockholm, Stockholm).
- Robson, A.J., 1990. Efficiency in evolutionary games: Darwin, Nash and the secret handshake. *Journal of Theoretical Biology* 144, 379-396.
- Samuelson, L. and J. Zhang, 1992. Evolutionary stability in asymmetric games. *Journal of Economic Theory* 57, 363-391.
- Stahl, D.O., 1993. Evolution of smart_n players. *Games and Economic Behavior* 5, 604-617.
- Swinkels, J., 1993. Adjustment dynamics and rational play in games. *Games and Economic Behavior* 5, 455-484.
- Tan, T. and S.R. Weibull, 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.
- Winter, S., 1971. Satisficing, selection, and the innovating remnant. *Quarterly Journal of Economics* 85, 237-261.
- Young, P., 1993. Evolution of conventions. *Econometrica* 61, 57-84.