

Weibull, Jörgen W.

Working Paper

The 'As if' Approach to Game Theory: 3 Positive Results and 4 Obstacles

IUI Working Paper, No. 392

Provided in Cooperation with:

Research Institute of Industrial Economics (IFN), Stockholm

Suggested Citation: Weibull, Jörgen W. (1993) : The 'As if' Approach to Game Theory: 3 Positive Results and 4 Obstacles, IUI Working Paper, No. 392, The Research Institute of Industrial Economics (IUI), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/94929>

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.



Industriens Utredningsinstitut

THE INDUSTRIAL INSTITUTE FOR ECONOMIC AND SOCIAL RESEARCH

A list of Working Papers on the last pages

No. 392, 1993

THE 'AS IF' APPROACH TO GAME THEORY: 3 POSITIVE RESULTS AND 4 OBSTACLES

by

Jörgen W. Weibull

September 1993

| | | | | |
|------------------|---------------|--------------|----------|------------|
| Postadress | Gatuadress | Telefon | Bankgiro | Postgiro |
| Box 5501 | Industrihuset | 08-783 80 00 | 446-9995 | 19 15 92-5 |
| 114 85 Stockholm | Storgatan 19 | Telefax | | |
| | | 08-661 79 69 | | |

THE 'AS IF' APPROACH TO GAME THEORY: 3 POSITIVE RESULTS AND 4 OBSTACLES*

Jörgen W. Weibull[†]

27 September 1993

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.

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

*Paper presented at the European Economic Association Meeting in Helsinki, August 27-29, 1993. The paper summarizes results from joint research with Abhijit Banerjee, Jonas Björnerstedt, Martin Dufwenberg, Peter Norman and Klaus Ritzberger, see references in text.

[†]Department of Economics and Institute for International Economic Studies, Stockholm University, S-106 91 Stockholm, Sweden. I thank Jonas Björnerstedt, Henrik Horn and Johan Stennek 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.

¹See e.g. Tan and Werlang [24], Aumann and Brandenburger [2].

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 [1], Friedman [13] and Winter [27].

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 "simplistic" ones? The 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 $\bar{u}(x)$.² This is the so-called *replicator dynamics* (Taylor and Jonker [25]):

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

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 [9]).
- 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 [17]).

²Note that the first of these two payoffs equals the (expected) payoff to pure strategy i when played against mixed strategy x , and the second equals the (expected) payoff to mixed strategy x when played against itself.

- 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 [22]).

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 modelled 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 [17], 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 convergent 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 [11].

This example is based on the Rock-Scissors-Paper (RSP) game; a constant-sum, symmetric 3×3 game with a unique Nash equilibrium. In equilibrium, both players randomize uniformly over the three pure strategies. Dekel and Scotchmer [11] modify somewhat the payoffs in this game and add a fourth pure strategy. The payoff matrix of the expanded RSP-game is

$$A = \begin{pmatrix} 1 & 2.35 & 0 & 0.1 \\ 0 & 1 & 2.35 & 0.1 \\ 2.35 & 0 & 1 & 0.1 \\ 1.1 & 1.1 & 1.1 & 0 \end{pmatrix} \quad (3)$$

The fourth pure strategy is strictly dominated by $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$, and the symmetric strategy-profile (x^*, x^*) is the unique Nash equilibrium of this game. Since strategy 4 is strictly dominated, its population share shrinks towards zero along all interior solution paths to the *continuous*-time replicator dynamics (1). In contrast, Dekel and Scotchmer [11] prove that this population share does *not* tend to zero in the discrete-time version (2), except when the initial state is perfectly symmetric with respect to the three first strategies.

As pointed out by Cabrales and Sobel [10], 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 [8] establish this for a class of games containing the Dekel-Scotchmer example in an alternative, overlapping-generations (OLG) model of biological reproduction in which a randomly drawn population share δ reproduces simultaneously (see Appendix). First, we show that, on the sub-simplex where x_4 is zero, the OLG-dynamics spirals inwards, towards the Nash equilibrium strategy x^* , when δ is small enough, and outwards, towards the boundary of the sub-simplex, when δ is large, see Figure 1 below.

Figure 1: The OLG replicator dynamics for game A when $x_4 = 0$.

Secondly, in the full mixed-strategy simplex, strategy 4 earns less than average inside an egg-shaped region containing x^* , and more than average outside this region. If the share δ of simultaneously revising individuals is sufficiently large, then the dynamics leads the population state out of this egg-shaped region and x_4 does *not* converge to zero, while if δ is sufficiently small, then all interior solution trajectories sooner or later enter the egg-shaped region and converge to the Nash equilibrium strategy x^* .

In sum: temporal lumpiness, as modelled in discrete time, is a threat to the "as if" paradigm. First, the explanatory power of evolutionary selection for Nash equilibrium behavior may be reduced. Secondly, the elimination of strictly dominated strategies is not guaranteed. However, if the population share of simultaneously switching individuals is sufficiently small, then the positive continuous-time result is restored.

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




Figure 1.

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 inbetween, 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 modelling purposes, imitation of succesful behaviors, experimentation with alternative behaviors, search for optimal behaviors etc, seem more appropriate. A few examples of such boundedly rational dynamic selection processes will here be sketched. In these models of "social evolution", the basic data are the game, the random matching interaction scheme, and the "switching technology", i.e., the mechanism whereby individuals change strategy.

A basic ingredient in the selection dynamics to follow is the assumption that a revising individual makes a *noisy* observation of the current *expected* payoff of one or more pure strategies j . Let the observed payoff value be $v_j(x) = u_j(x) + \varepsilon_j$, where the first term is the expected payoff value and the second a random error term. For any pair of pure strategies i and j , let $p_{ij}(x)$ denote the probability that the observed payoff value for i exceeds that of j , i.e., $p_{ij}(x) = \Pr[v_i(x) > v_j(x)]$.

One can model a process of imitation as follows (Weibull [26]): each revising individual samples at random another individual from the population, with equal probability for all other individuals. Hence, the probability that the sampled individual uses pure strategy j is x_j . The conditional probability that the sampled individual's payoff will be observed to be higher than the sampling individual's own strategy is $p_{ji}(x)$, in which case the sampling individual switches to the sampled individual's strategy. Hence, the probability that the sampling individual will switch from his strategy i to another strategy j is $x_j p_{ji}(x)$. This results in the following continuous-time *imitation* dynamics:

$$\dot{x}_i = \sum_j x_j [p_{ij}(x) - p_{ji}(x)] x_i. \quad (4)$$

Under mild regularity conditions on the probability functions $p_{ij}(x)$, such a dynamics is *payoff monotonic* in the sense that a pure strategy i with higher current payoff $u_i(x)$ than another pure strategy has a higher current growth rate

\dot{x}_i/x_i than the latter (see [8]). Because of this monotonicity, the above positive results on the connections between continuous-time evolutionary selection and Nash equilibrium are easily established. However, for other monotonic selection dynamics than the replicator dynamics (1), such as (4), the implication for strict dominance is weaker: pure strategies which are strictly dominated by other *pure* strategies vanish along such solution paths (Samuelson and Zhang [22]). Whether a strictly dominated strategy which is not strictly dominated by any pure strategy may survive the particular dynamics (4) has not yet been investigated.³ Hence, more research is needed before we know the implications of evolutionary selection by imitation for "game-theoretic rationality".

An individual who imitates successful individuals in the way described above may do very well. If she samples often and with little noise, she can much of the time play a best reply to the current population mixture, and rarely be using poorly performing strategies. Another class of selection dynamics arises if we instead assume that each revising individual makes noisy observations of *all* pure strategies' expected payoffs, and shifts to the strategy with the highest observed payoff. Unlike in the case of imitation, the conditional switching probabilities for an individual now is functionally independent of her own current strategy (unless the sampling scheme is biased towards or against this). Despite its kinship in spirit with game-theoretic rationality, the induced *noisy best-reply* dynamics turns out to be non-monotonic with respect to payoffs. Again, more research is needed before we know whether such selection mechanisms are "threats" or "promises" to the "as if" paradigm (see Björnerstedt et al [8] for a preliminary investigation).

A dynamics with a flavor of market selection, and which in a sense is a dual to the noisy best-reply dynamics, is suggested in Björnerstedt [7]. Here one can imagine some outside observer, a "bank", who regularly 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 the pure strategy which is observed to have the lowest observed payoff. Again invoking reasonable regularity properties of the underlying probability functions $p_{ij}(x)$, the induced continuous-time *cut-off dynamics* can be shown to be payoff monotonic and hence the two first positive results apply, while the third, on strictly dominated strategies, need not hold. In fact, Björnerstedt [7] shows that the strictly dominated fourth pure strategy in Dekel's and Scotchmer's [11] example can survive in the long run. Hence, the harsh policy of abandoning the weakest need not weed out strictly dominated strategies - if these give a fairly stable payoff while other strategies' payoffs are volatile.

Conclusion: while some positive "as if" results are known to be valid for all payoff monotonic selection dynamics, not all "plausible" selection dynamics are such. More research on specific social and market-performance based selec-

³See Björnerstedt [7] for a study of a class of generalized imitation-experimentation dynamics.

tion mechanisms is needed before we know whether these constitute threats or promises to the "as if" paradigm.

5 Rich menu of behaviors

Robson [21], Dekel and Scotchmer [11], Banerjee and Weibull [3], [4], [5] and Stahl [23] 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 not 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". They also recognize each other, and then play some iteratively strictly undominated (pure or mixed) strategy. In contrast, the original "biological" agents go on using their pre-programmed pure strategies as before. What will happen?

Borrowing an example from Banerjee and Weibull [3],[4], suppose the payoff matrix of the game is

$$A = \begin{pmatrix} 3 & 1 & 6 \\ 0 & 0 & 4 \\ 1 & 2 & 5 \end{pmatrix}. \quad (5)$$

This game is strictly dominance solvable: strategy 2 is strictly dominated, both by strategy 1 and 3, and, once strategy 2 has been eliminated, strategy 3 is strictly dominated by strategy 1. Hence, the biologist expects, by the third positive result above, the population to converge from any interior initial state to the monomorphic population state in which *all* agents use strategy 1 in every encounter. Note also that had instead all agents in the population belonged to the precious species *homo oeconomicus*, then these would all instantly have chosen strategy 1.

What if the initial population is mixed, and, for instance, contains many biological agents programmed to the strictly dominated strategy 2, few biological agents programmed to strategies 1 and 3, and some *homo oeconomicus*? The latter would earn payoff 3 when meeting each other and payoff 2 when meeting biological agents of "type 2", i.e., programmed to strategy 2. In contrast, type-2 agents would earn zero when meeting each other but 4 when meeting *homo*

oeconomicus. The type-2 agent so to speak benefit from their "commitment" to the strictly dominated but "aggressive" strategy 2. Hence, for a sufficiently large population share of *homo oeconomicus*, and low shares of biological agents programmed to strategies 1 and 3, biological agents programmed to strategy 2 do well; they may even earn *more* than *homo oeconomicus*.

In fact, one can show that, starting from any initial population mixture containing positive shares of all three types of biological agents, and of *homo oeconomicus*, the population share of biological agents of type 3 vanishes asymptotically in the continuous-time replicator dynamics. Once this population share is small, two things can happen, depending on the initial state. Either the population state moves towards a continuum of states in which *all* agents use strategy 1, as predicted by evolutionary and non-cooperative game theory alike. In this case, the survivors are some *homo oeconomicus* mixed with some biological agents programmed to strategy 1. As put by Stahl [23]: "being right is just as good as being smart". Alternatively, the population state moves towards the state at which $\frac{2}{3}$ of the population belong to *homo oeconomicus* and $\frac{1}{3}$ are programmed to the strictly dominated strategy 2. In the latter case, the three pure strategies of the game are used in proportions $\frac{4}{9}$, $\frac{1}{3}$ and $\frac{2}{9}$, respectively; definitely at variance with the biologist's expectation. Moreover, even if the distressed biologist perturbs this population state by injecting, say, a few "good" agents, programmed to strategy 1, the replicator dynamics leads the population back towards its "bad" habit of, in aggregate, playing the strictly dominated mixed strategy $(\frac{4}{9}, \frac{1}{3}, \frac{2}{9})$. In fact, this outcome is asymptotically stable in the replicator dynamics 1, as applied to this situation.⁴

To see how this can be established, first note that the situation in the biologist's lab is equivalent to letting a population of *programmed* agents play the following 4×4 game:

$$A^* = \begin{pmatrix} 3 & 1 & 6 & 3 \\ 0 & 0 & 4 & 4 \\ 1 & 2 & 5 & 1 \\ 3 & 2 & 6 & 3 \end{pmatrix}. \quad (6)$$

For, to be a *homo oeconomicus* in the original game is equivalent to being programmed to strategy 4 in this expanded game, since this strategy earns the best-reply payoff against all strategies, and all other strategies earn that payoff against strategy 4 which they get when meeting their best replies. The rest of the argument is a standard application of the replicator dynamics (1). As shown in Banerjee and Weibull [4], the population share $x_3(t)$ decreases to zero along

⁴A state x is *asymptotically stable* (in any given dynamics) if it is Lyapunov stable, i.e., such that solution trajectories starting near x remain forever near x and, moreover, asymptotically approach x as time goes towards plus infinity.




Figure 2

any interior solution path, and the dynamics on the sub-simplex where strategy 3 is extinct looks as in Figure 2 below.

Figure 2: The replicator dynamics (1) for game A^* when $x_3 = 0$.

The above example can be embedded in a richer model framework, and then it turns out that the three initial positive bench-mark results on evolution and rationality are restored, albeit in a somewhat weaker and more complex form. In a sense, we embed the example in a model in which biological evolutionary selection operates at the level of "decision procedures" rather than on 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 four distinct signals, one for each type of biologically programmed agent (strategies 1,2,3) and one for *homo oeconomicus*. In that example, each agent of type $c = i \leq 3$ is programmed to the *constant* rule which prescribes strategy i for all encounters, and each agent of type $c = 4$ (*homo oeconomicus*) is programmed to the "best reply" rule which prescribes the unique best reply to each signal $c \leq 3$ and the unique iteratively undominated strategy to signal $c = 4$.

Given 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) (much in the same way as the matrix A^* was construed from the payoff matrix A above). Formally, such a meta-game is identical with a cheap-talk game with (pre-play) message space C . Moreover, the three 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 Banerjee and Weibull [5]). First, each Lyapunov sta-

ble 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.

Why did this not happen in the above laboratory experiment? The answer is simply that the initial population state was not interior. With 3 pure strategies and 4 types, there are $3^4 = 81$ possible strategy-choice rules, and we only allowed for 4. For instance, a small injection of agents "looking" exactly like those programmed to strategy 2, but using the rule which prescribes strategy 1 against all opponents, would do better than the "original" type 2 agents etc. If such mutations were allowed for, or the initial state was interior, the process would eventually move towards a population state in which only strategy 1 was used in every matching.

Conclusion: the "as if" paradigm does not apply to all situations of "rich menus of behavior". However, if applied at the associated level of "decision procedures" or "cognitive designs", the paradigm does apply, albeit with somewhat weaker and more complex implications at the "base level" of actual decisions. (See also the hierarchical model in Stahl [23].)

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"). This modification of the set-up is evidently needed for evolutionary analyses of arbitrary finite n -player games, but is also relevant for certain symmetric games. 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:

"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." (Nash [18], p. 21.)

In this 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

$$\dot{x}_i^k(t) = [u_{ki}(x(t)) - \bar{u}_k(x(t))] x_i^k(t) \quad \text{for } t \geq 0, \quad (7)$$

where $u_{ki}(x)$ 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, $u_{ki}(x)$ 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.⁵

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 (7). The results now read as follows. First, if a population state x is Lyapunov stable in (7), then x is a Nash equilibrium (Nachbar [17], Friedman [12]). Secondly, if (a) the initial population state $x(0)$ is interior, i.e., all population shares $x_i^k(0)$ are positive, and (b) the state $x(t)$ converges to some state x^* as time t goes towards plus infinity, then x^* is a Nash equilibrium (Nachbar [17]). Thirdly, if (a) the initial population state $x(0)$ is interior, and (b) a pure strategy i for player k is iteratively strictly dominated, then its population share $x_i^k(t)$ converges to zero as time t goes towards plus infinity (Samuelson and Zhang [22]).

On the surface, thus, the mere multiplicity of populations seems to be no obstacle, per se, to the "as if" paradigm. However, a closer examination reveals that the above positive results may become virtually vacuous when the

⁵Just 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.

player populations are distinct. For instance, in certain symmetric two-player games (such as the so-called Hawk-Dove game) there are symmetric Nash equilibria which are asymptotically stable in the single-population dynamics (1) but unstable in the corresponding two-population dynamics (7). The reason is the possibility that the two populations "polarize", i.e., abandon asymmetric states near a symmetric equilibrium. Consequently, although dynamic stability implies Nash equilibrium, fewer population states are stable in the multi-population setting. More generally and exactly: a population state x is asymptotically stable in the multi-population replicator dynamics (7) 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 [19]; see also Hofbauer and Sigmund [14], Samuelson and Zhang [22], and Ritzberger and Weibull [20]). 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. (The predictive power of the result on the elimination of strictly dominated strategies is unaffected by this observation, since it presupposes neither stability nor convergence.)

One way to restore asymptotic stability, and thereby hopefully re-establish some non-vacuous positive link between evolutionary selection and Nash equilibrium play, is to consider *sets* X of population states (mixed strategy profiles), rather than individual population states x .⁶ It turns out that, for a certain class of such sets X , there is a simple necessary and sufficient condition for their asymptotic stability in the replicator dynamics (7). This condition can be expressed in terms of the so-called *better-reply correspondence* γ ([20]), 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:⁷

$$\gamma^k(x) = \{i : u_{ki}(x) \geq \bar{u}_k(x)\}. \quad (8)$$

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. Such a "subspace" is thus the set of those mixed strategy profiles which assign *no* probability mass to pure strategies outside the subsets in question. In terms of population states: it is the set of population states in which *all* individuals in each player population use *only* those pure strategies which are in the corresponding subset. One special such subspace X is evidently the full mixed-strategy space itself, other special spaces are the singleton-sets which contain a pure-strategy profile.

⁶Such a closed and non-empty set X (a special case being $X = \{x\}$ for some state x) is called *asymptotically stable* in the replicator dynamics (7) if solution trajectories to (7) starting near X remain forever near to X and converge over time to X .

⁷Cf. the usual (pure strategy) *best-reply* correspondence β^k which assigns to any strategy profile x those pure strategies i which are optimal for player k against x : $\beta^k(x) = \{i : u_{ki}(x) \geq u_{kj}(x) \forall j\}$.

A subspace X is said to be *closed* under the better reply correspondence if every strategy profile x in X is mapped by γ to pure strategies in the subsets defining ("spanning") X . Clearly the full mixed-strategy space is closed in this sense, and a singleton set X containing a strict (and hence pure) Nash equilibrium, if such exists, is also closed under γ . The result alluded to above is two-fold. First, a subspace X is asymptotically stable in the multi-population replicator dynamics (7) if and only if it is closed under γ (Ritzberger and Weibull [20]). Hence, closure under γ , which is easy to verify in simple games, implies that the associated subsets of pure strategies together constitute an *attractor* in the replicator dynamics, i.e., even if not all agents in the population use strategies in the subsets, eventually they will, granted the not too many deviated initially. Secondly, each such closed subspace *contains* a subset Y which is *strategically stable* in the sense of Kohlberg and Mertens [16], i.e. a minimal closed and nonempty set Y of Nash equilibria with the property that all slightly perturbed games (in the sense of "trembling hands") have some Nash equilibrium near to Y . This is one of the most stringent refinements in non-cooperative game theory. Hence, there is a link, here expressed in terms of set inclusion, between set-wise dynamic evolutionary stability and stringent non-cooperative set-wise robustness requirements.

Since every game possesses at least one subspace which is closed under γ , and there are finitely many such subspaces, there 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 evolutionary approach has no "cutting power". In other games, some such minimal subspaces contain few pure strategies for each player, and the model's cutting power is accordingly stronger.

The conclusion concerning the effect of population multiplicity for the "as if" approach is that all three positive results remain formally correct, but the predictive power of the first two, on Nash equilibrium, is much reduced, due to lost stability. The power of the discussed set-valued approach is game dependent, and requires more research.

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 not to constitute a real problem in its simplest formulation, once the implications

were somewhat relaxed. The fourth obstacle, that the interaction may involve individuals from distinct populations, was seen to weaken the predictive power of the evolutionary approach towards set-valued inclusions.

Evidently, these are not the only obstacles or complications which have to be faced before one can claim that we have 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" are here 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, some results have already been obtained, but more research is needed. The interested reader is advised to consult e.g. Binmore and Samuelson [6], Kandori, Mailath and Rob [15], Young [28], to mention a few examples of fine research along some of these lines.

Nash equilibrium and other non-cooperative solution concepts have become almost indispensable tools for economics. Further research on the rationalistic and evolutionary foundations of these are therefore highly relevant. Evolutionary game theory has supplied us with some powerful analytical tools for this, tool which we need to sharpen and apply, but, perhaps more importantly, we need to create new tools designed for the institutional framework and cognitive machinery of human beings who interact strategically in markets and within organizations.

8 Appendix: an OLG-model of biological reproduction

Suppose that reproduction/death takes place r times per time unit, each time involving a randomly drawn population fraction of size $\delta = \frac{1}{r}$, all individuals having equal chance of being drawn. Let the times of reproduction/death be $t = 0, \delta, 2\delta, \dots$. If $u_i[x(t)]$ is the number of offspring of an individual reproducing at time t , then the corresponding *overlapping-generations replicator dynamics* becomes

$$x_i(t + \delta) - x_i(t) = \frac{u_i[x(t)] - \bar{u}[x(t)]}{r - 1 + \bar{u}[x(t)]} x_i(t) \quad \text{for } t = 0, \delta, 2\delta, \dots \quad (9)$$

The discrete-time replicator dynamics (2) evidently corresponds to the special case of one reproduction time per time unit, $r = 1$, and the continuous-time replicator dynamics (1) is obtained in the limit as the number r of reproduction times per time units tends towards plus infinity, or, equivalently, as the population fraction $\delta = \frac{1}{r}$ of simultaneously reproducing individuals tends to zero.

References

- [1] Alchian A. (1950), "Uncertainty, evolution and economic theory", **Journal of Political Economy** 58, 211-222.
- [2] Aumann R. and A. Brandenburger (1991), "Epistemic conditions for Nash equilibrium", mimeo., Hebrew University.
- [3] Banerjee A. and J.W. Weibull (1991), "Evolutionary selection and rational behavior", Kirman A. and M. Salmon (eds.), **Rationality and Learning in Economics**, Blackwell (forthcoming).
- [4] Banerjee A. and J.W. Weibull (1992), "Evolution and rationality: some recent game-theoretic results", **Proceedings from the Tenth World Congress of the International Economic Association**, Blackwell (forthcoming).
- [5] Banerjee A. and J.W. Weibull (1993), "Evolutionary selection with discriminating players", Department of Economics, Harvard University, WP 1637.
- [6] Binmore K. and L. Samuelson (1992), "Evolutionary stability in repeated games played by finite automata", **Journal of Economic Theory** 57, 278-305.
- [7] Björnerstedt J. (1993), "Experimentation, imitation and evolutionary dynamics", mimeo., Department of Economics, Stockholm University.
- [8] Björnerstedt J., Dufwenberg M., Norman P. and J.W. Weibull (1993), "Evolutionary selection dynamics and irrational survivors", mimeo, Department of Economics, Stockholm University.
- [9] Bomze I. (1986), "Non-cooperative two-person games in biology: a classification", **International Journal of Game Theory** 15, 31-57.
- [10] Cabrales A. and J. Sobel (1992), "On the limit points of discrete selection dynamics", **Journal of Economic Theory** 57, 407-419.
- [11] Dekel E. and S. Scotchmer (1992), "On the evolution of optimizing behavior", **Journal of Economic Theory** 57, 392-406.

- [12] Friedman D. (1991), "Evolutionary games in economics", **Econometrica** 59, 637-666.
- [13] Friedman M. (1953), "The methodology of positive economics", in Friedman M., **Essays in Positive Economics**, University of Chicago Press, Chicago.
- [14] Hofbauer and Sigmund (1988), **The Theory of Evolution and Dynamical Systems**, Cambridge University Press.
- [15] Kandori M., G. Mailath and R. Rob (1993), "Learning, mutation, and long-run equilibria in games", **Econometrica** 61, 29-56.
- [16] Kohlberg E. and J.-F. Mertens (1986), "On the strategic stability of equilibria", **Econometrica** 54, 1003-1037.
- [17] Nachbar J. (1990), "'Evolutionary' selection dynamics in games: convergence and limit properties", **International Journal of Game Theory** 19, 59-89.
- [18] Nash J. (1950), "Non-cooperative games", PhD thesis, Princeton University.
- [19] Ritzberger K. and K. Vogelsberger (1990), "The Nash field", Institute for Advanced Studies, Vienna, RR 263.
- [20] Ritzberger K. and J.W. Weibull (1993), "Evolutionary selection in normal-form games", The Industrial Institute for Economic and Social Research, Stockholm, WP 383.
- [21] Robson A.J. (1990), "Efficiency in evolutionary games: Darwin, Nash and the secret handshake", **Journal of Theoretical Biology** 144, 379-396.
- [22] Samuelson L. and J. Zhang (1992), "Evolutionary stability in asymmetric games", **Journal of Economic Theory** 57, 363-391.
- [23] Stahl D.O. (1992), "Evolution of smart_n players", Department of Economics, University of Texas, mimeo.
- [24] Tan T. and S.R. Werlang (1988), "The Bayesian foundations of solution concepts of games", **Journal of Economic Theory** 45, 370-391.
- [25] Taylor P. and L. Jonker (1978), "Evolutionary stable strategies and game dynamics", **Mathematical Biosciences** 40, 145-156.
- [26] Weibull J.W. (1992), "An introduction to evolutionary game theory", The Industrial Institute for Economic and Social Research, Stockholm, WP 347.
- [27] Winter S. (1971), "Satisficing, selection, and the innovating remnant", **Quarterly Journal of Economics** 85, 237-261.
- [28] Young P. (1993), "Evolution of convention", **Econometrica** 61, 57-84.