

Game Theory: The Next Stage*

Eric van Damme[†]

November 1993

Revised, March 1995

Abstract

This paper surveys some recent developments in (non-cooperative) game theory and provides an outlook on the near future of that theory. In particular, attention is focused on the limitations inherent in normative game theory and on attempts to construct a behavioral version of the theory that incorporates aspects of procedural and bounded rationality. It is argued that a redirection towards more empirical work may be called for.

*Based on a talk held at “Economics, the Next Ten Years”, a conference to celebrate the 10th anniversary of GREQE, Marseille, September 7-9, 1992. The author thanks Hans Carlsson, Stef Tijs, Oscar Volij and an anonymous reviewer for comments on an earlier version. The usual disclaimer applies.

[†]Mailing address: CentER for Economics Research, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, the Netherlands.

The next stage develops when the theory is applied to somewhat more complicated situations in which it may already lead to a certain extent beyond the obvious and the familiar. Here theory and application corroborate each other mutually. Beyond this lies the field of real success: genuine prediction by theory. (Von Neumann and Morgenstern (1947, p. 7, 8.)

1 Introduction

Game theory provides a framework, a language, for modelling and analyzing interactive decision situations, that is, situations in which multiple decision makers with (partially) conflicting objectives interact. It aims at understanding human behavior in such conflict situations and at grasping how the resulting outcome depends on the “rules of the game”. Such understanding then enables advise on which changes in the rules might allow more desirable outcomes to be reached. Three different types of game theory might be distinguished:

- (i) *normative game theory*, in which one analyses the consequences of strategic behavior by superrational players,
- (ii) *descriptive game theory*, which is concerned with documenting how people actually make decisions in game situations, and
- (iii) *prescriptive game theory*, which aims at giving relevant and constructive advise that enables players to reach better decisions in game situations.

Of course, it is not always straightforward to categorize a game theoretic contribution as descriptive, normative, or prescriptive. To enhance understanding of an actual conflict situation, elements from all branches may be needed.

In the past, mainly the normative branch of game theory has been developed. This theory is based on strong rationality principles. The players are assumed to know what they want, they are exclusively motivated to obtain what they want and they have an unbounded, costless, calculating ability for analyzing any situation. In addition, the theory assumes that these players analyze the situation by means of a common model. The rules of the game are assumed common knowledge, i.e. everybody knows that everybody knows that such and such are the rules. As Aumann (1987) puts it: “The common knowledge assumption underlies all of game theory and much of economic theory. Whatever the model under discussion, (...), the model itself must be assumed common knowledge; otherwise the model is insufficiently specified, and the analysis incoherent”.

Normative game theoretic analysis is deductive, the theory analyses which outcomes will result when (it is common knowledge that) the game is played by rational individuals. The main aim is “to find the mathematically complete principles which define ‘rational behavior’ for the participants in a social economy, and to derive from them the general characteristics of that behavior” (Von Neumann and Morgenstern (1947, p. 31)). In other words “the basic task of game theory is to tell us what strategies rational players will follow and what expectations they can rationally entertain about other rational players’ strategies”. (Harsanyi and Selten (1988, p. 342).) Of course, the theory of rationality should not be self-destroying, hence, in a society of people behaving according to the theory, there should be no incentive to deviate from it. In consequence, normative theory has to prescribe the play of a Nash equilibrium.

In the last two decades, game theoretic methods have become more and more important in economics and the other social sciences. Many scientific papers in these areas have the following basic structure: A problem is modeled as a game, the game is analyzed by computing its equilibria, and the properties of the latter are translated back into insights relevant to the original problem. The close interaction between theory and applications has, inevitably, led to an increased awareness of the limitations of the theory. It has been found that the tools may not be powerful enough or that they may yield results which do not provide a useful benchmark for the analysis of actual behavior. For example, many models admit a vast multiplicity of equilibrium outcomes so that the predictive power of game theoretic analysis is limited. To increase understanding, it may, hence, be necessary to perfect the tools. In other models, such as in Selten’s (1978) chain store paradox, the theory yields a unique recommendation, but it is one that sensible people refuse to take seriously as a guide for actual behavior. Hence, new tools need to be developed as well.

Of course, one should not be surprised to find discrepancies between predictions based on normative game theory and outcomes realized in practice. The rationality of human beings is limited. Human players are not mental giants with nerves of steel, they do not always know what they want, they may not be motivated to obtain what they want and their cognitive abilities are severely bounded. Furthermore, in many real life games it is

not clear what the rules of the game are. Even if they are clear, it is not sure that people are aware of them, let alone that they are common knowledge. Consequently, one may raise the important question of the empirical relevance of normative game theory. How can it be that a theory based on such idealizing assumptions can say anything sensible about the real world? Can it actually say something sensible? In which contexts does a game theoretic solution concept, or a prescription on the basis of it, make sense?

Harsanyi (1977) expresses an optimistic attitude. According to Harsanyi, a normative solution concept is not only useful to clarify the conceptual issues involved in the definition of rationality. It is prescriptively relevant since it can serve as a benchmark for actual behavior:

- (i) It can help with explaining and predicting the behavior of players in those cases where they can be expected to behave *as if* they are rational.
- (ii) It can lead to a better understanding of actual behavior in situations different from those covered by (i), i.e. the behavior might be explained as an understandable deviation from rationality.

Of course, this leaves open the question of when people can be expected to behave as if they are rational, hence, in which contexts a solution concept is a useful benchmark.

Other game theorists are much more pessimistic than Harsanyi. For example, Raiffa expresses the frustration that he experienced after having accepted an appointment at the Harvard Business School, just after *Games and Decisions* was published in 1957:

“I began by studying loads of case studies of real-world problems. Practically every case I looked at included an interactive, competitive decision component, but I was at a loss to know how to use my expertise as a game theorist.”

(Raiffa (1982, p. 2)

Raiffa continues by pointing out the limitations of the normative theory, the gap that exists with descriptive theory, and the difficulties of erecting a meaningful prescriptive theory on the basis of foundations provided by existing theory:

“The theory of games focuses its attention on problems where the protagonists in a dispute are superrational, where the “rules of the game” are so well understood by the “players” that each can think about what the others are thinking about what he is thinking, ad infinitum. The real business cases I was introduced to were of another variety: Mr. X, the vice-president for operations of Firm A, knows he has a problem, but he’s not quite sure of the decision alternatives he has and he’s not sure that his adversaries (Firms B and C) even recognize that a problem exists. If Firm A, B, and C behave in thus-and-such a way, he cannot predict what the payoffs will be to each and he doesn’t know how he should evaluate his own payoffs, to say nothing about his adversaries’ payoffs. There are uncertainties all around besides those that relate to the choices of Firms B and C; no objective probability distributions for those ancillary uncertainties are available. Mr. X has a hard time sorting out what he thinks about the uncertainties and about the value tradeoffs he confronts, and he is in no frame of mind to assess what Mr. Y of Firm B and Mr. Z of Firm C are thinking about what he’s thinking. Indeed, Mr. X is mainly thinking about idiosyncratic issues that would be viewed by Y and Z as completely extraneous to their problems. Game theory, however, deals only with the way in which ultrasmart, all-knowing people *should* behave in competitive situations, and has little to say to Mr. X as he confronts the morass of his problem.” (Raiffa (1982, p. 2).)

The challenge of game theory is to bridge the gap from the ideal world of mathematics to the study of actual behavior in the real, complex world. A main lesson from the past seems to be that exclusive development of a normative theory does not bring success: the hope of obtaining precise, reasonable predictions on the basis of general rationality principles was idle. A really useful theory with high predictive power that can be used for prescriptive purposes has to stand on two legs, a deductive one and descriptive one. At the present stage, the marginal return to developing the latter leg is higher.

To realize the potential of game theory it is first of all necessary to improve its empirical base. In order to obtain a broad set of facts on which to theorize on, we

need to do fieldwork and careful laboratory experimentation. We have to study actual human behavior in order to find regularities in that behavior so as to be able to construct meaningful theories of procedural and bounded rationality. We indeed need to construct such theories. In order to successfully develop applied game theory, it is also useful to document what has been learned in the past, to make an overview of the cases in which existing game theory has been successfully applied as well as to document why the theory does not work in the cases where it does not work. In which situations can we expect existing theory to improve our understanding of the real world? What is the range of successful applications of existing theory? Which models and which solution concepts are most appropriate in which contexts?

In this paper I comment on some recent developments in game theory that are building blocks towards a better theory. Since the bulk of the work is in normative (non-cooperative) game theory, I mainly restrict myself to this branch. To provide a perspective, I start in Section 2 by giving a broad overview of the developments in game theory in the last forty years. In Section 3, I discuss game theoretic models and some applications of game theory: What can experiments and fieldwork tell us about the domain of applicability of the models and the solution concepts? In Section 4, I discuss three rationales underlying the notion of Nash equilibrium, one involving perfect rationality, another involving limited rationality and the final one being based on perfect absence of rationality. Section 5 is devoted to issues of bounded rationality. It describes some recent research that investigates the consequences of the players being bounded information processing devices and discusses the difficulties associated with modelling human reasoning processes. Section 6 offers a brief conclusion.

2 History: Problems Solved and Unsolved

In the preface of the *Contributions to the Theory of Games Volume 1*, published in 1950, the editors Harold Kuhn and Albert Tucker, list 14 outstanding problems in the theory of games. Half of these deal with two-person zero-sum games in which the most important one listed is “to find a computational technique of general applicability for finite zero-

sum two-person games with a large number of pure strategies.” Other problems in this class concern the existence of a value for games with an infinite number of strategies, to characterize the structure of the set of solutions and to construct efficient algorithms to find solutions. As evidenced by the little activity in this area at present, these problems have been solved satisfactorily. In particular, optimal strategies can be found efficiently by linear programming techniques.

In cooperative game theory, the problems that Kuhn and Tucker list are (i) to ascribe a formal value to an arbitrary n -person game, (ii) to establish significant asymptotic properties of n -person games for large n , (iii) to establish existence of Von Neumann-Morgenstern stable sets for arbitrary n -person games and to derive structural characteristics of such stable sets and (iv) to extend the theory to the case where utility is not transferable. To some extent these problems have also been satisfactorily solved. Shapley (1953) defined and axiomatized a value for n -person games. Extensions to NTU-games were given in Harsanyi (1963) and Shapley (1968), with axiomatizations being provided in Aumann (1985) and Hart (1985). Asymptotic properties were studied in Aumann and Shapley (1974). Of particular importance were the equivalence theorems for large games and markets, see Aumann (1964), Debreu and Scarf (1963) and Mas-Colell (1989). As far as vNM-stable sets are concerned, the picture is somewhat less satisfactory: stable sets do not always exist (Lucas (1968)) and the concept is so difficult to work with that few general structural properties are known.

One problem on the Kuhn/Tucker list has cooperative as well as noncooperative aspects. It is “to study n -person games with restrictions imposed on the forming of coalitions (e.g. embargo on side payments or on communication, or on both), thus recognizing that the cost of communication among the players during the pregame coalition-forming period is not negligible but rather, in the typical economic model with large n , is likely to be the dominating consideration. One approach to this question might be to formalize the coalition-forming period as a noncooperative game in the sense of Nash.” In this area we have seen less progress than in the ones mentioned before. The process of coalition-formation (and coalition dissolution) is still not well understood and it is only very recent that we have seen noncooperative formalizations of the coalition forma-

tion process along the lines suggested by Kuhn and Tucker (see Harsanyi (1974), Selten (1981), Selten (1991a), and, for an overview of the special case where only 2-person buyer/seller coalitions are relevant, Osborne and Rubinstein (1990)). Much of the recent work, however, is plagued by the fact that the extensive forms studied admit infinitely many equilibria. It seems safe to conjecture that we will see more work in this area in the near future.

Within noncooperative game theory, Kuhn and Tucker mention two problems which “need further classification and restatement above all”. Here “we find the zone of twilight and problems which await clear delineation”. The problems are

- (a) “To develop a comprehensive theory of games in extensive form with which to analyze the role of information, i.e. the effect of changes in the pattern of information”.
- (b) “To develop a dynamic theory of games:
 - (i) In a single play of a multimove game, predict the continuation of the opponent’s strategy from his early moves.
 - (ii) In sequence of plays of the same game, predict the opponent’s mixed strategy from his early choices of pure strategies”.

The seminal work of Kuhn (1953) and Selten (1965, 1975) on extensive form games, and of Harsanyi (1968) and Aumann (1976) on information and knowledge, has enabled more formal restatements of these problems as well as theory development. In the past two decades much effort has been devoted to try to solve problem (a), i.e. to define solution concepts that capture rational behavior in extensive form games. Various concepts have been proposed and their properties have been vigorously investigated. In the process severe difficulties in the foundations of game theory have been uncovered. Some authors (e.g. Basu (1990), Reny (1986)) have even argued that game theory is inconsistent since the theory’s basic assumption, the common knowledge of players’ rationality, cannot hold in nontrivial extensive form games.

By using the newly developed solution concepts, a start has been made to study the influence of the distribution of information on the outcome of play. It was found that

there is a disturbing lack of robustness of the rational outcome: small changes in the information structure may have drastic consequences. In particular, inserting a tiny bit of irrational behavior may have a large impact on the rational solution of a game. Furthermore, when rationality is not common knowledge, a rational player may benefit by pretending to be somewhat irrational, so that the superiority of “rational behavior” over other behavior is not clear. The latter poses a challenging problem for game theory since it goes to the heart of it. Namely, as the founding fathers already wrote “The question remains as to what will happen if some of the participants do not conform. If that should turn out to be advantageous to them (...) then the above “solution” would seem very questionable.” (Von Neumann and Morgenstern (1947, p. 32). We will return to this problem area in Section 4.

Within normative theory, the Kuhn/Tucker problem (bi) about learning and prediction reduces to a routine computation using Bayes’ rule, since this theory imposes the assumptions of perfect rationality and equilibrium. Important results were obtained, for example, concerning information revelation in repeated games of incomplete information. As a consequence of the growing awareness of the limitations of the rationality and equilibrium assumptions, however, there has recently been a renewed interest in actual learning of boundedly rational players in dynamic games, as well as in processes of evolutionary selection in such situations. We return to this topic in Sections 4.2 and 4.3.

With respect to the Kuhn/Tucker problem (bii) about repeated games, we note that the most important result obtained is the Folk Theorem. This result establishes a link between cooperative and noncooperative game theory: repetition (with ‘enough’ information being transmitted to players during the game) may allow players to reach more efficient outcomes in equilibrium (Aumann and Shapley (1976), Rubinstein (1978), Fudenberg and Maskin (1986)). Since the result shows that under weak assumptions almost anything can happen, it can also be interpreted as a negative result. Cooperation might result but it need not, also less efficient outcomes may be sustained by equilibria of repeated games. There is an embarrassing multiplicity of equilibria and this raises the important question of why one should expect any equilibrium at all. We return to this

question in the Section 4.1.

To summarize, although the theory has been developed substantially, progress uncovered many new conceptual and technical problems. These include problems with the foundations of the theory and with the justification of the solution concepts, problems of multiplicity (making the analysis inconclusive) and problems of lack of robustness of the outcomes with respect to the assumptions underlying the model. At the same time, the better developed theory enabled more extensive application. Inevitably, the increased number of applications, led to an increased awareness of the limitations and weaknesses in the theory. In particular, the applications threw doubt on the relevance of strong rationality assumptions, i.e. it was found that game theoretic solutions may be hard to accept as a guide to successful practical behavior. In the next sections we describe these drawbacks in more detail and discuss how game theorists try to overcome them.

3 The Rules of the Game

Two types of game models can be distinguished, cooperative and noncooperative. Cooperative models assume that binding contracts can be concluded outside the formal rules of the game. In a noncooperative game, no external enforcement mechanism is available. To put it differently, these models differ in the amount of detail that they require to be specified. At the macro level, we have the cooperative model in which the actual process of play is abstracted away from. In this model, attention is focused on the options and alternatives that the various coalitions have. At the other extreme, we have the noncooperative, extensive form model, which allows faithful modelling of actual real life institutions. In this model, emphasis is on the process of play and the analyst is forced to specify the fine structure of the game, i.e. the strategic variables that are relevant, the timing of the moves and the distribution of information among the players at the relevant points in time. At the intermediate level we find the (noncooperative) strategic form which focuses on the strategies that individual players have available and the resulting outcomes. Each of these models has its attractive features as well as its drawbacks. Broadly speaking, incorporating more detail allows a more refined prediction

but that prediction might depend very strongly on those details, hence, the prediction might not be robust.

Already in the modelling stage game theory can provide important insights as it forces the analyst to go through a checklist of questions that afford a classification of the situation at hand. (How many players are there? Who are they? What do they want? What can they do? When can they do it? What do they know? Can they sign binding contracts? etc. etc.) This classification in turn allows one to see similarities in different situations and allows the transfer of insights from one context to another.

The theoretical development of the extensive form model and of games with incomplete information that followed the seminal work of Aumann, Harsanyi, Kuhn and Selten made noncooperative game theory more suited for application and in the past two decades game theoretic methods have pervaded economics as well as the other social sciences. Game theoretic methods have come to dominate the area of industrial organization, and game theoretic tools have been essential for the understanding of economies in which information is asymmetrically distributed. The rapid growth of the use of game theory in economics in the last two decades can be attributed in part to the fact that the extensive form model allows a tremendous degree of flexibility in modelling. Any real life institution can be faithfully modelled and the explicitness allows scrutinizing the model's realism and limitations. The richness of the model also has its drawbacks, however. First of all, richer models allow more flexibility in classifying conflict situations, hence, it is more difficult to obtain general insights. Secondly, it is more difficult to do sensitivity analysis. Thirdly, and most importantly, it will only rarely be the case that the situation at hand dictates what the model should be. Frequently, there is considerable scope for designing the model in various ways, each of them having something going for it. Judgement on the appropriateness of the model is essential.

If there is freedom in filling in the details of the game one would ideally want the outcome to be robust with respect to inessential changes. Unfortunately, this happy state of affairs only seems to materialize for solution concepts that are based on weak rationality assumptions, i.e. concepts that do not make sharp predictions. Outcomes generated by stronger solution concepts, such as subgame perfect equilibrium or stable

equilibrium (Kohlberg and Mertens (1986)) frequently depend critically on these details. The application of these concepts then requires the modelling of all details. Indeed, Kohlberg and Mertens do make the assumption that the model is isomorphic to reality rather than an abstraction of it: “we assume that the game under consideration fully describes the real situation — that any (pre-)commitment possibilities, any repetitive aspect, any probabilities of error, or any possibility of jointly observing some random event, have already been modelled in the game tree” (Kohlberg and Mertens (1986, fn. 3)). The appropriateness of such an assumption for applied work may be questioned, especially since human beings are known not to perceive all details of a situation.

In the recent past we have seen a tendency for models to get more detailed and more complex. For example, models with incomplete information rely on Harsanyi’s (1968) trick of adding a fictitious chance move to ensure the common knowledge of the model. This construction quickly yields a complicated game so that, even in the case where the analyst can solve the game, the question remains as to how relevant that solution is for real life players with bounded computational complexities. Furthermore, at the intuitive level one might argue that the more detail one adds to the model, the less likely it is that this detail is commonly perceived by all parties involved in the actual conflict, hence, the less credible the common knowledge assumption. The extensive quotation from Raiffa in Section 1 suggests that by making the assumption that the model is common knowledge, game theory abstracts away from the most basic problem of all: “What is the problem to be solved?”. We return to this issue in Section 5.

3.1 Laboratory Experiments

The structural characteristics of the game can be controlled in laboratory experiments, but even in experiments it is hard to control the preferences and beliefs of the players and how players perceive the situation. Experiments with ultimatum bargaining games have shown that proposers do not ask for (almost) the entire cake, but we do not know why they do not do so. We do not know what motivates people and what they believe about the motivations of others. Do people make “fair” proposals because they are intrinsically motivated by “fairness” or because they believe the opponent cares about

the distribution of payoffs and not just about his own share? The experimental data contradict the joint hypotheses that is common knowledge that (a) players are only interested in their own monetary payoffs and (b) want to maximize these payoffs, but this conclusion is not very informative. We want to dig deeper and get to know why the results are as they are and why the results depend on the context in the way they do. (See Güth and Van Damme (1994) for a systematic investigation how the results depend on the amount of information that is transmitted from the proposer to the responder.)

Experiments may give us a better idea of the settings in which the use of a game theoretic solution concept, like Nash equilibrium is justified. An important and intriguing puzzle is offered by the experimental research on double auctions (Plott (1987), Smith (1990)). The experiments show that Nash equilibria may be reached without players consciously aiming to reach it. However, the Nash equilibrium that is obtained is not an equilibrium of the complex, incomplete information game that the players are playing, rather it is the (Walrasian) equilibrium of the associated complete information game. In addition, giving players information that allows them to compute the equilibrium, may make it less likely that this equilibrium is reached. In these experiments there is a number of traders, each trader i assigning a value $v_i(n)$ to n units of the good that is traded. All the information is private. Agents act both as price makers (by announcing bids and asks) and as price takers (by accepting bids or asks from other traders). The prices and allocations converge quickly to the Walrasian equilibria of the complete information economy, even though none of the traders has the information to compute this outcome and none is aware of the fact that traders are maximizing profits in this way. Subjects participating in these experiments describe them as “nontransparent”, hence, the question to be answered is “Which institutional aspects help players in reaching the equilibrium?” Which mechanisms provide sufficient feedback for players to learn the equilibrium? We will return to this question at a more abstract level in the next section.

The interaction between individual behavior and the institutional setting is not very well understood yet and, no doubt, much research will be done in this area in the next decade. It is interesting to compare the ‘good news’ from the previous paragraph with the ‘bad news’ obtained by Kagel and Levin (1986). These authors study common value

auctions in which participants can fall prey to the winners' curse. It turns out that, given enough experience, players learn to avoid making losses. However, they certainly do not learn to understand the situation and play the equilibrium. In fact, their learning does not allow them to cope with a change in the circumstances: If the number of bidders is increased, the players increase their bids while equilibrium behavior would force them to shade their bids even more. As a consequence, players make losses for some time until they have learned to cope with the new situation. It would be interesting to know what would happen if players gained experience in many circumstances. An example, from a completely different context, suggests that learning from different environments may allow people to learn more: Selten and Kuon (1992) study 3-person bargaining games in which only 2-person coalitions can form. They find that players who gain experience with many diverse situations learn to understand the logic of the quota solution and come to behave in accordance with it. In contrast, behavior of bargainers who draw from a more limited set of experiences does not seem to settle down.

3.2 Applied Game Theory

One of the areas in which game theoretic tools have been extensively applied in the last two decades is industrial organization. Several commentators (Fisher (1989), Pelzman (1991)) have argued that these applications have not been successful as the theory has been unable to reach general insights. Part of the reason, no doubt, lies in the three drawbacks of the extensive form model that were mentioned above. The flexibility of that model has enabled the theoretical explanation of any type of behavior. In addition, it seems that, when writing the details in the extensive form, theorists have not been guided by detailed empirical work. Of course, the problem might not lie in the tool but in the subject of study: it might be that no general insights are possible. Furthermore, the rapid spread of game theoretic ideas through this literature can probably be attributed more to the absence of good competitors to game theory than to the fact that the field is ideally suited to apply game theoretic techniques: The basic assumptions underlying game theory (common knowledge of rationality and of the rules of the game) are not particularly realistic in this area. Since the formal rules governing competition are not

very detailed, industrial organization is not the first area one thinks of for successful application of game theory based on the extensive form.

Game theory may be more successful in situations that are closer to its base, situations in which the rules are clear and where one can have more faith in the players' rationality. Financial markets immediately come to mind: The rules are clear, the game offers opportunities for learning and the stakes are high, so that one could at least hope that irrational behavior is driven out. However, as evidenced by the large number of anomalies in this area, one should also not expect too much here (De Bondt and Thaler (1992)). Nevertheless, it seems that game theory could contribute something to the analysis of financial markets. For example, following the Big Bang in London, European stock exchanges have gone through a series of restructurings in order to try to increase their competitiveness. Such restructurings involve changes in the rules of the trading game, hence, problems of mechanism design (Pagano and Roell (1990)).

Auctions are another case in which the context pins down the structural characteristics of the game, i.e. the actions and their timing, not, however, the distribution of information. Standard auction models can make definite predictions about outcomes and indeed in some cases the predictions match the data reasonably well (Hendriks and Porter (1988)).

Another case in which the institutional structure of interaction is well-defined is the one of centralized matching markets studied by Al Roth (Roth (1984), Roth and Sotomayor (1990), Roth and Xing (1994)). Although the rules of the noncooperative game are well-defined, Roth's point of departure for the analysis lies in cooperative game theory, viz. in the Gale and Shapley (1962) investigation of whether the core of such a game, the set of stable matchings, is nonempty. Roth investigates whether rules that produce stable outcomes will remain in place and whether rules that produce unstable outcomes will induce market failure and, hence, be abolished after some time. The hypothesis indeed is powerful to explain the success of various market rules, hence, we have a clear example where not only the rules determine the play but, vice versa, the play influences the rules. It would be nice to see more examples of this phenomenon and to know whether there is a tendency for "efficient" rules to evolve. What is particularly

interesting here is the combination of ‘cooperative’ and ‘noncooperative’ elements in the analysis, providing an example that might be successful also in other contexts.

I expect that in the next decade the pendulum will swing back again from noncooperative theory in the direction of cooperative game theory. In some cases it will not pay the analyst to model the game to the greatest possible detail. Rather it might be more attractive to consider the situation at a more aggregate level and to make broad qualitative predictions that hold true for a large range of detailed specifications of the actual process. To some extent, this redirection is already occurring in 2-person bargaining theory. The noncooperative underpinnings of Nash’s solution (that were enabled by the seminal paper Rubinstein (1982)) are useful in that they increase our confidence in that solution and since they show us how to include “outside options” in Nash’s original cooperative model. (See Binmore et al. (1992) for an overview.) Once this has been established, rational advice to bargainers concerning what policies to pursue, as well as comparative statics properties (for example, concerning risk aversion) can be derived from (an appropriate modification) of Nash’s original cooperative model.

Of course, cooperative game theory is still underdeveloped according to several dimensions. We know little about the dynamic processes of coalition formation and coalition dissolution and very little about cooperation under incomplete information. I expect to see some work on these problems in the near future.

4 Equilibrium and Rationality

Noncooperative game theoretic analysis centers around the notion of Nash equilibrium, hence, it is essential to address the relevance of this solution concept. Why do we focus on Nash equilibria? When, or in which contexts is Nash equilibrium analysis appropriate? Where do equilibria come from? How can one choose among the equilibria?

There are at least three interpretations (justifications) of the notion of Nash equilibrium:

- (i) it is a necessary requirement for a self-enforcing theory of rational behavior;
- (ii) it results as the outcome of a learning process;

(iii) it results as the outcome of an evolutionary process.

In the three subsections that follow, we discuss these justifications in turn as well as some recent literature dealing with each topic.

4.1 Perfect Rationality

The first justification of Nash equilibrium is a normative one. Nash equilibrium arises in addressing the question “What constitutes rational behavior in a game?” A theory of rationality that prescribes a definite (probabilistic) choice (or belief) for each player has to prescribe a Nash equilibrium, since otherwise it is self-contradictory. In Nash’s own words: “By using the principles that a rational prediction should be unique, that the players should be able to deduce and make use of it, and that such knowledge on the part of each player of what to expect the others to do should not lead him to act out of conformity with the prediction, one is led to the concept” (Nash (1950)). Nash also comments on the limited scope of this justification: “In this interpretation we need to assume the players to know the full structure of the game in order to be able to deduce the prediction for themselves. It is quite strongly a rationalistic and idealizing interpretation” (Nash (1950)).

This rationalistic interpretation relies essentially on the assumptions that each game has a unique rational solution and that each player knows this solution. To address the question of how players get to know the solution, one needs a formal model that incorporates players’ knowledge. Such a model has been developed by Aumann and the reader is referred to Aumann and Brandenburger (1991) for a discussion of the epistemic conditions underlying Nash’s concept. We just remark here that in the 2-player case less stringent conditions suffice than in general n -player games. (Roughly speaking, in the two player case, mutual knowledge of beliefs, rationality and payoffs suffices, while in the n -player case, one needs common knowledge assumptions, as well as a common prior on the beliefs.)

While the rationalistic interpretation of the Nash equilibrium concept relies on the assumption that each game has a unique strategy vector as its rational solution, it is the case that many games admit multiple Nash equilibria. Hence, the rationalistic

interpretation requires to address two questions: (i) Can a theory of rational behavior prescribe any Nash equilibrium? (ii) What constitutes rational behavior in case there are multiple equilibria? These questions are addressed, respectively, in the literatures on equilibrium refinement and equilibrium selection.

The research that has been performed on extensive form games has made it clear that the answer to the first question must be in the negative: Certain Nash equilibria are not compatible with perfect rationality as they rely on incredible threats. To rule out these equilibria, Selten (1965) started a program of refining the equilibrium concept. Many different variations were proposed, each imposing somewhat stronger rationality requirements than the Nash equilibrium does. (See Van Damme (1987) for an overview). Game theorists have not yet agreed upon the ultimate refinement: We certainly do not yet have a convincing answer to the question: “What constitutes rational behavior in an extensive form game?”

Recently, the relevance of the refinements program has been questioned. Namely, most refinements (in particular, subgame perfect and sequential equilibrium) insist on “persistent rationality”, i.e. it is assumed that no matter what has happened in the past, it is believed that a rational player will play rationally in the future. This assumption might well be a sensible one to make for perfectly rational players, but it is a problematic one in the applications of the theory, especially, if the application involves a simplified model. Human players are not perfectly rational, they make mistakes and they might deviate from perfect rationality in a systematic way. Once, in a real game, one sees that a player deviates from the rational solution of the game, one should not exclude the possibility that one’s model of the situation or one’s model of that player is wrong. Of course what one should then believe cannot be determined by that original model. The model has to be revised. If the model of the exogenous environment is appropriate, one is forced to enrich the model by incorporating actual human behavior. Hence, in extensive form games, the perfectly rational solution might not be a good benchmark against which to compare actual behavior.

With regard to changes in the model, it is worrisome that the perfectly rational solution might change drastically with minor changes in the data of the game. For

example, if there is a small probability of there being irrational players around, rational players might play very differently than in the case where this possibility does not exist (Kreps et al. (1982)). Human players have free will; if a player can profit from behaving differently than the theory of perfect rationality prescribes, there is nothing that can prevent the player from doing so.

Nash already stressed that the normative interpretation of equilibrium requires one to solve the problem of equilibrium selection. A solution to this problem has been provided in Harsanyi and Selten (1988), in which a coherent single-valued theory of rationality for interactive decision situations has been constructed. However, the Harsanyi/Selten book also shows that such a theory necessarily has to violate certain intuitively desirable properties. For example, a theory of rationality that only depends on the best reply structure of the game necessarily has to pick a Pareto inferior Nash equilibrium in some games. In the stag hunt game $g(x)$ of Figure 1, (c, c) is the Pareto dominant equilibrium if $x < 2$. This game, however, is best-reply-equivalent to a common payoff coordination game with diagonal payoffs equal to $2 - x$ and x and off-diagonal payoffs equal to zero. In the latter game, the Pareto dominant equilibrium is (d, d) if $x > 1$.

	c	d
c	$2, 2$	$0, x$
d	$x, 0$	x, x

Figure 1: Stag Hunt Game $g(x)$ ($0 < x < 2$)

Since there is room for alternative theories of equilibrium selection — the Harsanyi/Selten theory violates certain desirable properties — as well as improvements, some work on rational equilibrium selection will continue. Carlsson and Van Damme (1993b) compare several (partial) theories that all derive their inspiration from Nash’s (1953) bargaining paper. Nash (1953) suggested to select equilibria on the basis of their relative stability properties. Specifically, he suggested to investigate robustness with respect to perturbations in the knowledge structure of the game. Hence, an equilibrium should survive in a “more realistic” version of the model. In a specific implementation of this

idea, Carlsson and Van Damme (1993a) show that "absence of common knowledge" may serve as an equilibrium selection device. Only some equilibria may be viable when there is just "almost common knowledge" (see also Rubinstein (1989)). Quite interestingly, in the Carlsson/Van Damme model, there is a form of "spontaneous coordination", one does not need to assume equilibrium behavior in the perturbed game to obtain equilibrium selection in the original game, iterated dominance arguments suffice. Hence, the model illustrates a possibility of how rational players might derive the solution.

In the Carlsson/Van Damme model it is common knowledge that a game from a certain class has to be played; players make observations on which game is played, but observations are noisy, with the errors of different players being independent. As a concrete example, suppose the game is as in Figure 1 but each player i makes a noisy observation $x_i = x + \varepsilon_i$ on the parameter characterizing the game. As a result of the noise, rational players have to analyze all games $g(x)$ at the same time: What is optimal for a player i at x_i depends on what his opponent does at points in the interval $[x_i - 2\varepsilon, x_i + 2\varepsilon]$ which in turn depends on what the opponent believes i will do on $[x_i - 4\varepsilon, x_i + 4\varepsilon]$. Clearly, player i will choose $c(d)$ if x_i is close to zero (two) since then chances are good that this action is dominant for the actual value of x . Having determined the behavior at the 'end points', a recursive argument allows determination of the optimal behavior at the other observations. Carlsson/Van Damme show that with vanishing noise players will coordinate on the risk dominant equilibrium, i.e. they will play c if and only if $x < 1$.

Formally, in a 2×2 game, one equilibrium is said to risk-dominate another if its associated (Nash) product of deviation losses is larger. In Fig. 1, a player's deviation loss from (c, c) is $2 - x$, while that from (d, d) is x , hence (c, c) is risk dominant if and only if $(2 - x)^2 > x^2$. This concept of risk dominance was introduced in Harsanyi and Selten (1988) and it has proved important also in other contexts. We will return to it below. It should be noted, however, that for games larger than 2×2 , the definition of risk-dominance is by means of the tracing procedure and that in these cases, the equilibrium selected by the Carlsson/Van Damme technique may be different from the risk dominant one. (See Carlsson and Van Damme (1993b) for further details. In particular, that paper derives the intuitive result that cooperation is more difficult when there are more

players: In an n -player version of Fig. 1, the observation where players switch from c to d is strictly decreasing in n .)

One interpretation of the Carlsson/Van Damme model is that the observations correspond to the players' models of the actual situation. Models of different players are highly similar, but they are not identical. The conclusion then is that this more realistic modelling implies that certain Nash equilibria are not viable. (Note the link with the discussion on perception in Section 5.)

Incorporating more realistic knowledge assumptions need not always reduce the number of equilibria. For example, Neyman (1989) shows that small changes in the knowledge structure may allow new equilibria to arise. He demonstrates that in the finitely repeated prisoner's dilemma, if players do not have a common knowledge upper bound on the length of the game, cooperation until near the end of the game is an equilibrium outcome. In the simplest of Neyman's models, the actual length n of the game is a draw from a geometric distribution, and each player i gets a signal n_i on the length of the game with $|n_i - n| = 1$ and $|n_1 - n_2| = 2$. Hence, player i knows that the actual length is either $n_i - 1$ (and his opponent has signal $n_i - 2$) or $n_i + 1$ (with the opponent having the signal $n_i + 2$). It is now easily seen that if each player follows the strategy of defecting only after a previous defection, or in case he is sure that the current round is the last one, an equilibrium results. In this equilibrium players cooperate until the next to last round.

From the above two examples, it is clear that much work remains to be done before we have a clear picture of how the equilibrium outcomes depend on the distribution of knowledge in the game. Hence, there is a need for further development of rationalistic theories, even though it may be questioned whether such theories will provide a useful benchmark for actual decision making.

4.2 Limited Rationality: Learning

The second interpretation of equilibrium, Nash calls 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 or 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” (Nash (1950)). If we assume that there is a stable frequency with which each pure strategy is used, players will learn these frequencies and they will play best responses against them. Consequently, a necessary condition for stability is that the frequencies constitute a Nash equilibrium. Nash remarks that “Actually, of course, we can only expect some sort of approximate equilibrium, since the information, its utilization, and the stability of the average frequencies will be imperfect.”

It is clear that Nash’s remarks raise many intriguing questions which cry for an answer. Under which conditions will there exist stable population frequencies? What will happen if the frequencies do not settle down? When will there be a limit cycle? Is it possible to have strange attractors or chaos? How does an approximate equilibrium look like? How long does it take before the process settles down? In what contexts is the long run relevant? How does the outcome depend on the information that players utilize? Does limited information speed up the process? Or might more limited information lead to completely different outcomes? What if the game is one in the extensive form and players only get to see the actual path of play and not the full strategies leading to the path? Do we get to more refined equilibrium notions? Can the concept of subgame perfect equilibrium be justified by some learning concept? Might non-Nash equilibria be asymptotically stable fixed points of learning processes? How does the outcome depend on the complexity of the reasoning processes that players utilize? Although some of these questions were already addressed in the fifties and sixties, in particular in relation to the Brown/Robinson process of fictitious play, interest dwindled after Shapley (1964) had given an example of a non-zero sum game for which this process does not converge, but rather approaches a limit cycle. Recently, interest has shifted again in this direction. At present, the above questions are being vigorously researched, research that will continue in the near future.

The interest in these questions, of course, arises from the observation that “There are situations in economics or international politics in which, effectively, a group of interests are involved in a noncooperative game without being aware of it, the non-awareness

helping to make the situation truly noncooperative” (Nash (1950)). Hence, players have only local information, they do not see through the system, but nevertheless they get such feedback from the system so as to be able to reach an equilibrium (cf. the discussion on the double auction in Section 3). Consequently, this second interpretation suggests a domain of relevance for Nash equilibrium that is completely different from the one suggested by the first interpretation. Whereas the first make sense in simple situations (“with an obvious way to play” (Kreps (1990))), the second relies on the situation not being obvious at all.

Formal models that address the questions raised above postulate that players behave according to certain rules, players are seen as information processing machines. The machine has a certain memory and an output rule that associates a decision to each possible state of memory. After each stage of play, the machine processes the information about this period’s play and incorporates it in its memory, thereby possibly changing its state. A collection of rules, one for each player, then determines the evolution of play, hence, the payoffs for each player (cf. the discussion in Section 5). The questions mentioned above then correspond to asking what consequences various rules will have. In the future we should expect to see both purely theoretical research (analyzing the properties of mathematically tractable processes), simulation studies of more complicated processes as well as empirical research: What kinds of learning processes do people adopt and what types of outcomes do these processes imply?

Most of the research done till now has been theoretical, but there have also been some simulation studies (for example, see Marimon et al (1990)). The work is so diverse and vast that it is impossible to summarize it here. I will confine myself to a simple illustration which is based on Young (1993). (Also see Kandori et al. (1993), Ellison (1993).) Assume that the game $g(x)$ from Figure 1 is played by members of two finite populations of size N . Each time period one member of each population is picked at random to play the game. In deciding what to do, a player (randomly) asks $k(k \ll N)$ members of his population what their most recent experience was in playing this game (i.e. what their opponents did) and he then plays a best response against this frequency distribution. Obviously, the system has two stationary states; all play c or all play d .

Hence, in the long run, an outside observer will only see a Nash equilibrium being played. Now let us add some small amount of noise. Assume that each player's memory may be imperfect: With small probability ε , a player remembers $c(d)$ when the actual experience was $d(c)$. The imperfection implies that the system may move from one equilibrium to the other. However, such movements are unlikely. To move away from "all c " one needs simultaneous mutation (i.e. imperfect recall) of a fraction $1 - x/2$ of the sample. To move away from "all d ", a fraction $x/2$ needs to mutate simultaneously. If ε is very small, then the first possibility is much more likely if $x > 1$, while the second is much more likely if $x < 1$. Hence, if $x < 1$ ($x > 1$) the system will remain much longer in "all c " ("all d ") than in "all d " ("all c "). In the ultra long run we get equilibrium selection according to the risk-dominance criterion. Again, the introduction of random variation leads to equilibrium selection.

4.3 Zero Rationality: Evolution

The third justification of Nash equilibrium has its origins in biology and was proposed first in Maynard Smith and Price (1973). In this interpretation there is no conscious choice at all: Individuals are programmed to play certain strategies, more successful strategies reproduce faster than others so that eventually only the most successful strategies survive. If the population reaches a stable state, all existing strategies must be equally successful, and strategies that are not present cannot be more successful. Hence, a stable state must be a Nash equilibrium. This interpretation, then, involves perfect absence of rationality.

In the most basic model of this type there is an infinite population of individuals who are randomly matched in pairs. Individuals are programmed to play strategies from a certain set S and if an s -individual meets a t -individual then the expected number of offspring to s is $u(s, t)$ where u is some symmetric bimatrix game. A monomorphic population in which only s^* -individuals are present is stable if any mutant $s \neq s^*$ who enters in the population with a small frequency is selected against. The formal condition for such stability is that (s^*, s^*) is a symmetric Nash equilibrium of u with $u(s^*, s) > u(s, s)$ for all alternative best replies s against s^* . A strategy s^* satisfying these conditions is

said to be an evolutionarily stable strategy or ESS. Hence, ESS is a refinement of Nash equilibrium. In the game of Figure 1, for example, both c and d are ESS, but the mixed strategy equilibrium does not correspond to an ESS. Within this framework one can also investigate the evolution of a polymorphic population. If the set of all possible strategies is finite, then, if the time between successive generations is small, the population proportions evolve according to the replicator dynamics $\dot{x}_s = x_s(u(s, x) - u(x, x))$. (In this expression, x_s denotes the fraction of s -individuals in the population, $u(s, x)$ is the expected number of offspring of an s -individual and $u(x, x)$ is the average fitness of the population). Broadly speaking, s^* is an ESS if and only if it is an asymptotically stable fixed point of the replicator dynamics.

Within this area several questions are presently being investigated. The answers at present are far from complete so that the research will continue. A typical question is whether evolutionary forces will wipe out irrational behavior, i.e. if $x(t)$ is a trajectory of the replicator equation and s is an (iteratively) dominated strategy, will $x_s(t)$ tend to zero as t gets large? Another question is whether evolutionary forces will produce equilibria, i.e. in which contexts does $\lim_{t \rightarrow \infty} x(t)$ exist and is such a limit an equilibrium of the game? Also we want to know the properties of ESS in specific classes of games. For example, do evolutionary pressures lead to efficient equilibria? In repeated games, does evolution force cooperation? (Axelrod (1984).) Furthermore, the basic model of a symmetric strategic form game is very limited, hence, how should evolutionary stability be defined in extensive form games or in asymmetric games? What are the properties of ESS's in these games? Furthermore, how should the definition be modified if mutants appear more frequently, or if there is local interaction, i.e. viscosity? What happens if there can be drastic innovations that can change the character of the game? (Holland (1992).)

Preliminary research on the above questions has shown that many games fail to have ESS. Theorists have been reluctant to give up the idea of equilibrium and they have proposed weaker concepts with better existence properties. Also set-valued concepts have been proposed. These are attractive for extensive form games, especially if the strategies in the set differ only off the equilibrium path. Finally, it has been argued that

in economic contexts it might be more appropriate to assume a bit more rationality on the part of mutants: New strategies will be introduced only if they have some chance of survival. Corresponding solution concepts have been defined and the properties of these are currently being investigated. As the literature dealing with this topic is vast I only give one example, and refrain from further comments. I refer to Van Damme (1994) for further details and references.

The example concerns the evolution of language. There is the common wisdom that, if players could communicate before playing the game $g(x)$ of Figure 1, they would talk themselves into the efficient equilibrium (c, c) . However, Aumann (1990) has argued that the conventional story may not be fully convincing. Furthermore, the intuition has been hard to formalize using equilibrium concepts that are based on perfect rationality. Recently, some progress has been made by using evolutionary concepts. The basic idea is very simple: In a population playing the inefficient equilibrium (d, d) , a mutant who sends a special signal and who reacts to the signal by playing c could possibly invade. Things are not that simple, however, success is not guaranteed. If the existing population punishes the use of the new signal (for example by playing the mixed strategy in response to it), then the mutant does worse than the existing population. Hence, if the mutant enters at the wrong point in time it will die out. However, the existing population cannot guarantee such punishment. Strategies that do not punish and behave on the equilibrium path just as other members of the population do equally well as the population and they can spread through it. If the mutant arises at a point in time when there are only few punishers around, it will thrive and eventually take over the entire population. Hence, with communication, the outcome (d, d) is not evolutionarily stable, the population will drift to (c, c) . (See Kim and Sobel (1991).)

5 Bounded Rationality

Standard game theoretic analysis assumes that it is common knowledge that players are rational in the Bayesian sense, hence, it is assumed to be common knowledge that each player (i) knows the set of all feasible alternatives, (ii) knows the exact consequences of

each action combination and (iii) has a globally consistent preference relation on the set of all possible consequences. The behavior of each player is assumed to be substantively rational, i.e. “it is appropriate to the achievement of given goals within the limits imposed by given conditions and constraints” (Simon (1976)). Hence, each player has a skill in computation that enables him to calculate infinitely fast, and without incurring any costs, the action that is optimal for him in the situation at hand.

Experiments and field work have shown that already in relatively simple situations human subjects may not behave *as if* they are substantively rational, at least, it may take a very long time before they behave this way. Hence, the empirical relevance of the Bayesian theory is limited. One of the virtues of game theory is that, by taking the Bayesian model to its logical extremes it has clearly revealed the limitations of that model. As Simon already wrote in 1955 “Recent developments (...) have raised great doubts as to whether this schematized model of economic man provides a suitable foundation on which to erect a theory – whether it be a theory of how firms do behave or how they rationally “should” behave” (Simon (1955)). He also wrote that the task we face is “to replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist” (Simon (1955, p. 99)).

Recent game theory papers have taken up this task. They deal with all kinds of cognitive limits — with respect to perception, memory, information processing capacity, knowledge and computational abilities, etc — by including them explicitly in the model and they deduce the consequences of the presence of these limits on outcomes. In these models, agents are information processing devices and the bounded rationality arises from the bounded capacity of the system. The models retain the optimization assumption: How should the machine be optimally designed within the constraints that are specified? Hence, it is assumed that the “system designer” is substantively rational. One interpretation of this meta-player is as an unconscious evolutionary process. Alternatively, it is a superrational player who correctly foresees the constraints that he will face and who takes these optimally into account.

As there are infinitely many possibilities for adding constraints, work of this type will certainly continue for quite a while. We will describe some of it in subsection 5.1. It is noteworthy that these models are not based on the empirical knowledge of actual thinking processes. In subsection 5.2 we discuss why not more input from psychology is used.

5.1 An Optimization Approach

The first models of bounded rationality in the game theory literature deal with repeated games and they depart from perfect rationality by taking complexity costs of implementing strategies into account. It is assumed either that strategies that are too complicated cannot be used (Neyman (1985)) or that more complex strategies have higher costs (Rubinstein (1986), Abreu and Rubinstein (1988)). Hence, Neyman's approach amounts to eliminating strategies from the original game, while Rubinstein's approach changes the payoffs. Both models view a strategy as an information processing rule, as a machine. The machine has a number of states and each state induces an action. In addition there is a transition function: Depending on the information that the machine receives (i.e. which action combination is played by the opponents) the machine moves to another state. The complexity of a strategy is measured by the number of states in the machine. Neyman assumes that players only have a certain number of states available, Rubinstein assumes that states are costly and that players care, lexicographically, about repeated game payoffs and complexity costs. Each player has to choose a machine at the beginning of the game, the chosen machines then play the repeated game against each other and each player receives the resulting payoff. Hence, we have a game in strategic form where the strategy set of each player is the set of all possible machines and we can investigate the Nash equilibria of this "machine game".

One can easily see what causes the action in Neyman's setup: By eliminating strategies one might create new, possibly more attractive equilibria. Indeed Neyman shows that, if players cannot use too complicated machines, cooperation might be an equilibrium outcome in the finitely repeated prisoner's dilemma. In Rubinstein's setup the results are driven by the observation that there are many ties in the payoff matrix of a repeated

game (many strategy profiles induce the same path), hence, small changes in these payoffs may have large effects. Introducing explicit cost for implementing a strategy indeed has drastic consequences: In the repeated prisoner's dilemma, for example, only the "diagonals" of the set of feasible payoffs can be obtained as Nash equilibrium payoffs of the machine game (Abreu and Rubinstein (1988)). By introducing costs also for the number of transitions, the set of equilibrium payoffs shrinks even further: Only the repetition of the one-shot equilibrium survives (Banks and Sundaram (1990)). The contrast with the Folk Theorem is remarkable.

Note that in these models the cost of calculating an equilibrium strategy are not taken into account, and that this calculation might be more complicated than calculating an equilibrium in the unrestricted game. In addition, there is the question of why Nash equilibria of the machine game are relevant. Binmore and Samuelson (1992) argue in favor of an evolutionary interpretation of the machine game in which equilibrium results from an evolutionary adaption process. Hence, nature might endow the players with the equilibrium and there is no issue of finding or computing it.

A next generation of models builds on the above ideas by incorporating limits on the information processing abilities of players. A player is viewed as an information processor: information flows in, is processed in some way and a decision results as an output. The processor has limited capacity, he can only carry out a certain number of operations per time period. Perhaps the capacity can be extended, but extensions are costly. A seminal paper is Rubinstein (1993) in which the consequences of heterogeneity in information processing ability are investigated. Some players can only distinguish high prices from low ones, they cannot make fine distinctions. The *ex ante* decision such a player has to make is which prices to classify as low ones and which as high ones, knowing that his final decision (whether to buy or not) can only depend on the classification of the price and not on the price itself. The question addressed is how one can optimally exploit such "naive" players. Formally, the model is a 2-stage game in which a kind of sequential equilibrium is computed: Players optimize taking their constraints and those of other players into account.

In Rubinstein's model, there is a monopolistic shop owner and two types of customers

A and B . The shop owner is privately informed about which state of nature prevails and to maximize his profit he would like to reveal this information to the type B individuals, but not to those of type A . Specifically, in a certain state of nature the monopolist would prefer to sell only to type B . It is assumed that the only signal that the monopolist has available is the price that he sets. Since the optimal price reveals the state, the monopolist's most desired outcome cannot be realized if all consumers can perfectly perceive the price: the type A consumers would correctly infer the state from the price. However, if perceptions of the type A consumers are imperfect, then the monopolist can do better. He can add some noise to his price signal and force consumers to pay attention to this noise by, possibly, hiding some relevant information behind it. By distracting consumers attention, they might not notice information that is really essential and the monopolist might be better off.

Fershtman and Kalai (1993) consider a similar model of a multimarket oligopolist with a limited capacity to handle information. The oligopolist can only pay attention to a limited number of markets and he has to decide how to allocate his attention: Should he stay out of markets where there is competition and where in order to play well he is forced to monitor the competitors' behavior closely, or should he rather devote much effort to those markets and go on "automatic pilot" in the monopolistic markets?

Models of limited attention like the above (but see also Radner and Rothschild (1975) and Winter (1981)) seem to me to be extremely relevant for actual decision making and to evaluate the role game theory can play in such situations. For example, should a business manager with limited time and attention, best focus his attention on the strategic interaction with the competitors or is he better off by trying to improve the organization of production within the firm? It is obvious that in real life we are involved in many games at the same time and that we do not devote equal time to analyzing each of them. *Ceteris paribus*, more important games deserve more attention, but certainly also the complexity of a game plays a role. I think it is important to find out how much time to devote to each game that one plays and expect to see some research in this area in the future.

Summarizing the above, we might say that, although the initial papers in this area

dealing with complexity of executing strategies where perhaps not directly practically relevant, they were tremendously important and improved our tools for the analysis of other aspects of bounded rationality. Nevertheless, a drawback is that this work does not take into account the cost of computing an optimal strategy. Probably, most actual situations are so complex that it is simply impossible to find an optimal strategy within the time span that is allowed. In such cases, one has to settle for a “good” solution. Such a solution may either be obtained from solving a drastically simplified problem exactly or it may be the result from a heuristic procedure applied directly to the complex situation. Game theory at present does not offer much advice on what to do when one has to rely on heuristics. The literature focuses exclusively on the question “What is optimal given the constraints?” It does not address the question “What is an efficient procedure for coming up with a reasonable solution?” The theory does not deal with “satisficing behavior”, it has not yet made the transition from studying behavior that is substantively rational to behavior that is procedurally rational, i.e. “behavior that is the outcome of appropriate deliberation” (Simon (1976)).

It seems reasonable to expect that, if computations are complicated and costly and if the computation process is not deterministic, one will not be able to determine exactly the point at which the other players stop computing, hence, one will not be able to figure out what the others will do. Each player will face uncertainty and there will be private information. Each player will stop computing only if he has a strategy that is a reasonably good response against the average expected strategy of the others. We do not necessarily end up at an equilibrium. It will be extremely interesting to see what such “robust choices” look like and whether or not they bear any relationship to existing game theoretic solution concepts.

5.2 A Behavioral Approach

The research discussed in the previous subsection is firmly entrenched in the deductive branch of game theory: A well-specified game is set up and the equilibria of the game are analyzed. Elements of bounded rationality are introduced within a given well-specified model: The problem to be solved is assumed given. Most actual decision problems,

however, are unstructured and complex. In reaching a decision one first has to construct a model and then one has to evaluate the decisions within that model. The papers discussed above do not deal with the question of how to generate an appropriate model. In most actual decision taking situations, however, most time is spent on trying to visualize and understand the situation, hence, on the formulation of a model that is appropriate for the situation. Hence, it is probably at the modelling stage that aspects of bounded rationality are most important. It is remarkable that none of the papers discussed above contains an explicit model of the reasoning process of a player, let alone that the papers take detailed empirical knowledge of actual human thinking processes into account. In this subsection we discuss some of these behavioral aspects of bounded rationality.

Most actual decision taking situations are complex. It is better to speak of the emergence of decisions than of decision taking. Broadly speaking, in reaching a decision, the actor has to perform the following steps of perceiving, thinking and acting:

1. Perception of the situation and generation of a model for analyzing it.
2. Problem solving. Searching for patterns, for similarities with other models and situations, and for alternative plans of action.
3. Investigating the consequences of (a subset of the) actions and evaluating them.
4. Implementing an action.
5. Learning: Store relevant information in memory so as to facilitate solving a similar problem later.

Selten (1978) develops an informal model of the human reasoning process that takes these steps into account. He emphasizes the importance of perception: How does a player view the problem, what elements does he consider to be relevant? What patterns does a player see? What types of similarities with other decision situations does a player notice? Selten suggests a model of the human reasoning process that is based on the idea that a decision may be reached on three levels, the levels of routine, imagination and reasoning. It is assumed that each higher level uses more information and needs more

effort than the lower one. Hence, because of the costs involved the player may decide not to even activate all levels. Furthermore, it is not necessarily true that the decision reached by the highest activated level will be taken. As Selten writes: “The reason is quite simple. It is not true that the higher level always yields a better decision. The reasoning process is not infallible. It is subject to logical and computational mistakes.” (Selten (1978, p. 150).)

Actually Selten makes an argument for decisions arising from the level of imagination. In game situations it is important to put oneself in the shoes of the other players in order to form expectations about their behavior. Since a player who makes decisions at the routine level “is likely to make some mistakes which can be easily avoided by imagining oneself to be in the other player’s position”, this level is unattractive in game situations. On the other hand “If a player tries to analyze the game situation in a rigorous way, then he will often find that the process of reasoning does not lead to any clear conclusion. This will weaken his tendency to activate the level of reasoning in later occasions of the same kind. ” Furthermore, rigorous reasoning has to be applied to a model of the situation, and to construct such a model, one has to rely on the level of imagination. Since “the imagination process is not unlikely to be more reliable as a generator of scenarios than as a generator of assumptions for a model of the situation,” this level will yield good solutions in many cases, so that Selten concludes that “one must expect that the final decision shows a strong tendency in favor of the level of imagination even in such cases where the situation is well structured and the application of rigorous thinking is not too difficult.”

The concrete problem Selten (1978) addresses is why a game theorist familiar with the backward induction argument in the finite horizon chain store game and who accepts the logical validity of that argument, nevertheless decides to fight entry if entry would occur in the first period. Hence, Selten discusses why a human player might deliberately neglect the advice offered by a rationality based theory. Hence, he deals with “motivational bounds” of human rationality, rather than with “cognitive bounds”. The solution that Selten offers is that, at the level of imagination, a player might not perceive the situation in the same way as the game theorist views the extensive game. A player’s model of

the situation need not contain all detail that the extensive form provides. Psychologists tell us that an observer exercises control over the amount of detail he wishes to take in and that people sometimes see things which are not there (also see Schelling (1960, fn. 18 on p. 108)). It matters: If the entrants believe that the monopolist classifies the situation just according to whether the horizon is far away or near and that he views a game with a horizon that is far away as one with an infinite horizon, then the deterrence equilibrium becomes possible.

The classical interpretation of a game is as a full description of the physical rules of play. Following Selten's lead implies taking seriously the idea that a player's model of a situation depends on how the player perceives the situation. In a game context, the fact that a player's perception of the situation need not coincide with the actual situation forces us to discuss a player's perception of the other players' perceptions. Rubinstein (1991) advocates viewing the extensive form as the players' common perception of the situation rather than as an exhaustive description of the situation. Hence, the model should include only those elements which are perceived by the players to be relevant. It is unknown what the consequences are of this reinterpretation of the game model. However, it should be noted that Schelling already stressed that the locus where strategic skill is important is in the modelling stage: The trick is to represent the situation in such a way that the outcome of the resulting model is most favorable to one's side (Schelling (1960, p. 69)).

It is a challenging task to formalize the above ideas, i.e. to develop models of procedural rationality. However, procedural rationality cannot be fruitfully studied without taking empirical knowledge of the actual thinking processes into account. Without having a broad set of facts on which to theorize, there is a certain danger of spending too much time on models that are mathematically elegant, yet have little connection with actual behavior. At present our empirical knowledge is inadequate and it is an interesting question why game theorists have not turned more frequently to psychologists for information about the learning and information processing processes used by humans. One reason might be that this knowledge is not available in sufficiently precise form in order to be meaningfully incorporated into a formal mathematical model. Another

reason might be that other social sciences might not have much knowledge available. Wärneryd (1993) explains that economists might have little to learn from psychologists since psychologists have shown remarkable little interest in economic issues. Also Selten is of the opinion that little of value can be imported. His 1989 Nancy L. Schwarz memorial lecture is entirely devoted to the question: “What do we know about the structure of human economic behavior?” After having discussed this question for 18 pages he concludes:

“I must admit that the answer is disappointing. We know very little (...). We know that Bayesian decision theory is not a realistic description of human economic behavior (...) but we cannot be satisfied with negative knowledge — knowledge about what human behavior fails to be (...). We must do empirical research if we want to gain knowledge on the structure of human economic behavior.”

To improve our understanding of human behavior, laboratory experimentation is essential. Unlike many current experiments which just inform us that the rationalistic benchmark is not very relevant, we need experiments that inform us why the deviations occur and how players reason in these situations.

6 Conclusion

The starting point of this paper was the observation that, while standard game theory is based on the assumption of perfectly rational behavior in a well-defined model that is common knowledge, real world problems, in contrast, are often unstructured and human players far from rational. As a consequence, it follows that much of standard game theory is largely irrelevant for prescriptive (or explanatory) purposes. In order to increase the applicability of the theory it is necessary to develop models that incorporate aspects of bounded rationality. At present such models are indeed developed, yet they retain the assumption of optimizing behavior. Hence, they limit themselves to situations that are simple enough to enable the optimization to be carried out. Complex situations offer

less scope for optimizing behavior and they force to address the problem solving aspects associated with procedural rationality: How is the situation perceived, how is it modelled and how do humans go about solving them?

That aspects of mutual perception and joint problem solving might be more important than individual optimization was already stressed by Schelling, who formulated the essential game problem as “Players must together find ‘rules of the game’ or together suffer the consequences” (Schelling (1960, p. 107)). Up to now, the road that Schelling pointed to has not been frequently traveled. Game theorists have instead followed the road paved by Nash. I conjecture that there will be a reorientation in the near future, i.e. that game theory will focus more on the aspects of imagination stressed by Schelling than on those of logic stressed by Nash. Of course, I might be wrong: The game of which route to take is one of coordination with multiple equilibria: Being on one road is attractive only if sufficiently many (but not too many) others travel that road as well. As the discussion of Figure 1 has shown there is no reason to expect the Pareto efficient equilibrium to result.

Schelling also already stressed that in order to increase the relevancy of game theory, it is necessary to develop its descriptive branch. Prescriptive theory has to stand on two strong legs:

“A third conclusion (...) is that some *essential* part of the study of mixed-motive games is necessarily empirical. This is not to say just that it is an empirical question how people do actually perform in mixed-motive games, especially games too complicated for intellectual mastery. It is a stronger statement: that the principles relevant to *successful* play, the *strategic* principles, the propositions of a *normative* theory, cannot be derived by purely analytical means from a priori considerations.” (Schelling (1960, p. 162, 163).)

Hence we may conclude with a message that is somewhat depressing for theorists. Just as at the inception of the theory, it might still be true that

“the most fruitful work may be that of careful patient description; indeed

this may be by far the largest domain for the present and for some time to come.” (Von Neumann and Morgenstern (1947, p. 2).)

References

- Abreu, D. and A. Rubinstein (1988). “The Structure of Nash Equilibrium in Repeated Games With Finite Automata”, *Econometrica* **56**, 1259-1282.
- Aumann, R.J. (1964). “Markets with a Continuum of Traders”, *Econometrica* **32**, 39-50.
- Aumann, R.J. (1976). “Agreeing to Disagree”, *The Annals of Statistics* **4**, 1236-1239.
- Aumann, R.J. (1985). “An Axiomatization of the Non-Transferable Utility Value”, *Econometrica* **53**, 599-612.
- Aumann, R.J. (1987). “Game Theory”, in J. Eatwell, M. Milgate and P. Newman (eds.), *The New Palgrave Dictionary of Economics*, 460-482.
- Aumann, R.J. (1990). “Nash Equilibria are not Self-Enforcing”, in J.J. Gabszewicz, J.-F. Richard and L.A. Wolsey (eds.), *Economic Decision-Making: Games, Econometrics and Optimisation* 201-206.
- Aumann, R.J. and A. Brandenburger (1991). “Epistemic Conditions for Nash Equilibrium”, Working paper 91-042. Harvard Business School.
- Aumann, R.J. and L.S. Shapley (1974). *Values of Non-Atomic Games*. Princeton University Press, Princeton NJ.
- Aumann, R.J. and L.S. Shapley (1976). “Long-Term Competition - A Game-Theoretic Analysis”. Mimeo Hebrew University. (Also published as WP 676, September 1992, Dept. of Econ. UCLA).
- Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York.

- Banks, J. and R. Sundaram (1990). "Repeated Games, Finite Automata and Complexity", *Games and Economic Behavior* **2**, 97-117.
- Basu, K. (1990). "On the Non-Existence of a Rationality Definition for Extensive Games", *International Journal of Game Theory* **19**, 33-44.
- Binmore, K., M.J. Osborne and A. Rubinstein (1992). "Noncooperative Models of Bargaining". chapter 7, 179-225, in: R.J. Aumann and S. Hart, (eds.), *Handbook of Game Theory, Vol. 1*, North-Holland, Amsterdam.
- Binmore, K. and L. Samuelson (1992). "Evolutionary Stability in Repeated Games Played by Finite Automata", *Journal of Economic Theory* **57**, 278-305.
- Carlsson, H. and E. van Damme (1993a). "Global Games and Equilibrium Selection", *Econometrica* **61**, 989-1018.
- Carlsson, H. and E. van Damme (1993b). "Equilibrium Selection in Stag Hunt Games". In K.G. Binmore and A. Kirman (eds.) *Frontiers of Game Theory*, 237-254. MIT Press, Cambridge MA.
- Damme, E. van (1987). *Stability and Perfection of Nash Equilibria*. Springer Verlag, Berlin. Second edition 1991.
- Damme, E. van (1994). "Evolutionary Game Theory". *European Economic Review* **38**, 847-858.
- De Bondt, W.F.M. and R. Thaler (1992). "Financial Decision Making in Markets and Firms: A Behavioral Perspective", Mimeo, School of Business, Univ. of Wisconsin, Madison.
- Debreu, G. and H. Scarf (1963). "A Limit Theorem on the Core of an Economy", *International Economic Review* **4**, 236-246.
- Ellison, G. (1993). "Learning, Local Interaction, and Coordination", *Econometrica* **61**, 1047-1072.

- Fershtman, C. and E. Kalai (1993). "Complexity Considerations and Market Behavior", *Rand Journal of Economics* **24**, 224-235.
- Fisher, F.M. (1989). "Games Economists Play: A Noncooperative View", *Rand Journal of Economics* **20**, 113-124.
- Fudenberg, D. and E. Maskin (1986). "The Folk Theorem in Repeated Games with Discounting and with Incomplete Information", *Econometrica* **54**, 533-554.
- Gale, D. and L.S. Shapley (1962). "College Admission and the Stability of Marriage", *American Mathematics Monthly* **69**, 9-15.
- Güth, W. and E. Van Damme (1994). "Information, Strategic Behavior and Fairness in Ultimatum Bargaining", Discussion Paper 9465 CentER for Economic Research, Tilburg University.
- Hart, S. (1985). "An Axiomatization of Harsanyi's Nontransferable Utility Solution", *Econometrica* **53**, 1295-1314.
- Harsanyi, J.C. (1963). "A Simplified Bargaining Model for the n -person Cooperative Game", *International Economic Review* **4**, 194-220.
- Harsanyi, J.C. (1968). "Games with Incomplete Information Played by "Bayesian" Players. Parts I, II, and III". *Management Science* **14**, 159-182, 320-332, 468-502.
- Harsanyi, J.C. (1974). "An Equilibrium Point Interpretation of Stable Sets and a Proposed Alternative Definition", *Management Science* **20**, 1472-1495.
- Harsanyi, J.C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press, Cambridge.
- Harsanyi, J.C. and R. Selten (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge MA.
- Hendriks, K. and R.H. Porter (1988). "An Empirical Study of an Auction with Asymmetric Information", *American Economic Review* **78**, 301-314.

- Holland, J.H. (1992). *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge MA.
- Kagel, J.H. and D. Levin. (1986). "The Winner's Curse and Public Information in Common Value Auctions", *American Economic Review* **76**, 894-920.
- Kandori, M., G.J. Mailath and R. Rob (1993). "Learning, Mutation and Long Run Equilibria in Games", *Econometrica* **61**, 29-56.
- Kim, Y.-G. and J. Sobel (1991). "An Evolutionary Approach to Pre-Play Communication", Mimeo, University of California at San Diego.
- Kohlberg, E. and J.-F. Mertens (1986). "On the Strategic Stability of Equilibria", *Econometrica* **54**, 1003-1039.
- Kreps, D.M. (1990). *Game Theory and Economic Modeling*. Oxford University Press, Oxford.
- Kreps, D.M., P. Milgrom, J. Roberts and R. Wilson (1982). "Rational Cooperation in the Finitely-Repeated Prisoner's Dilemma", *Journal of Economic Theory* **27**, 245-252.
- Kuhn, H.W. (1953). "Extensive Games and the Problem of Information", *Annals of Mathematics Studies* **28**, 193-216.
- Kuhn, H.W. and A.W. Tucker (1950). "Contributions to the Theory of Games", Vol. 1. *Annals of Mathematics Studies* **24**. Princeton University Press, Princeton.
- Lucas, W.F. (1968). "A Game with no Solution", *Bulletin of the American Mathematics Society* **74**, 237-239.
- Luce, R.D. and H. Raiffa (1957). *Games and Decisions: Introduction and Critical Survey*. Wiley, New York.
- Marimon, R., E. McGrattan and T. Sargent (1990). "Money as a Medium of Exchange in an Economy with Artificial Intelligent Agents". *Journal for Economic Dynamics and Control* **14**, 329-373.

- Mas-Colell, A. (1989). “An Equivalence Theorem for a Bargaining Set”, *Journal of Mathematical Economics* **18**, 129-139.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- Maynard Smith, J. and G.R. Price (1973). “The Logic of Animal Conflict”, *Nature, Lond.* **246**, 15-18.
- Nash, J.F. (1950). *Non-Cooperative Games*. PhD Dissertation, Princeton University.
- Nash, J.F. (1953). “Two-Person Cooperative Games”, *Econometrica* **21**, 128-140.
- Neumann, J. von, and O. Morgenstern (1947). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton.
- Neyman, A. (1985). “Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner’s Dilemma”, *Economics Letters* **19**, 227-229.
- Neyman, A. (1989). “Games Without Common Knowledge”, Notes, Hebrew University.
- Osborne, M.J. and A. Rubinstein (1990). *Bargaining and Markets*. Academic Press, San Diego.
- Pagano, M. and A. Roell (1990). “Trading Systems in European Stock Exchanges: Current Performance and Policy Options”, *Economic Policy*, **10**, 63-115.
- Pelzman, S. (1991). “The Handbook of Industrial Organization: A Review Article”, *Journal of Political Economy* **99**, 201-217.
- Plott, C.R. (1987). “Dimensions of Parallelism: Some Policy Applications of Experimental Methods”. In: A.E. Roth (ed.) *Laboratory Experimentation in Economics: Six Points of View*. Cambridge University Press, Cambridge MA.
- Radner, R., and M. Rothschild (1975). “On the Allocation of Effort”. *Journal of Economic Theory* **10**, 358-376.

- Raiffa, H. (1982). *The Art and Science of Negotiation*. Harvard University Press, Cambridge MA.
- Reny, P. (1986). “Rationality, Common Knowledge and the Theory of Games”, PhD-dissertation, Princeton University.
- Roth, A.E. (1984). “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory”, *Journal of Political Economy* **92**, 991-1016.
- Roth, A.E. and M. Sotomayor (1990). *Two-sided matching: A Study in Game Theoretic Modelling and Analysis*. Cambridge University Press, 1990.
- Roth, A.E. and X. Xing (1994). “Jumping the Gun: Patterns and Pathologies in the Timing of Market Transactions”, *American Economic Review* **84**, 992-1044.
- Rubinstein, A. (1978). “Equilibrium in Supergames with the Overtaking Criterion”, *Journal of Economic Theory* **21**, 1-9.
- Rubinstein, A. (1982). “Perfect Equilibrium in a Bargaining Model”. *Econometrica* **50**, 97-109.
- Rubinstein, A. (1986). “Finite Autonomy Play the Repeated Prisoner’s Dilemma”, *Journal of Economic Theory* **39**, 83-96.
- Rubinstein, A. (1989). “The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge”. *American Economic Review* **79**, 385-391.
- Rubinstein, A. (1991). “Comments on the Interpretation of Game Theory”. *Econometrica* **59**, 909-924.
- Rubinstein, A. (1993). “On Price Recognition and Computational Complexity in a Monopolistic Model”, *Journal of Political Economy* **101**, 473-484.
- Schelling, T.C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge MA.

- Selten, R. (1965). "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträglichkeit", *Zeitschrift für die gesamte Staatswissenschaft* **12**, 301-324.
- Selten, R. (1975). "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games", *International Journal of Game Theory* **4**, 25-55.
- Selten, R. (1978). "The Chain Store Paradox", *Theory and Decision* **9**, 127-159.
- Selten, R. (1981). "A Noncooperative Model of Characteristic Function Bargaining", in V. Böhm and H.H. Nachtkamp (eds.), *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*. 131-151.
- Selten, R. (1991a) (ed.). *Game Equilibrium Models, Vol 3. Strategic Bargaining*. Springer Verlag, Berlin.
- Selten, R. (1991b). "Evolution, Learning and Economic Behavior", *Games and Economic Behavior* **3**, 3-24.
- Selten, R. and B. Kuon (1992). "Demand Commitment Bargaining in Three-Person Quota Game Experiments", Discussion Paper B219 University Bonn.
- Selten, R., M. Mitzkewitz and H.R. Uhlich (1988). "Duopoly Strategies Programmed by Experienced Players", Discussion Paper B106, Bonn University.
- Shapley, L.S. (1953). "A Value for n -person Games", in H.W. Kuhn and A.W. Tucker (eds.), *Contributions to the Theory of Games. Vol. 2. Annals of Mathematics Studies*. **28**, 305-317. Princeton University Press, Princeton.
- Shapley, L.S. (1964). "Some Topics in Two-Person Games", *Annals of Mathematics Studies* **52**, 1-28.
- Shapley, L.S. (1968). "Utility Comparison and the Theory of Games", in G.T. Guilbaud *La Décision: Aggregation et Dynamique des Ordres de Preference*. Paris, Editions du CNRS.
- Simon, H. (1955). "A Behavioral Model of Rational Choice", *Quarterly Journal of Economics* **69**, 99-118.

- Simon, H. (1976). "From Substantive to Procedural Rationality", in S.J. Latsis (ed), *Methods and Appraisal in Economics*. 129-148, Cambridge University Press.
- Smith, V.L. (1990). "Experimental Economics: Behavioral Lessons for Microeconomic Theory and Policy", Nancy L. Schwarz Memorial Lecture. J.L. Kellogg Graduate School of Management.
- Young, P. (1993). "The Evolution of Conventions", *Econometrica* **61**, 57-84.
- Wärneryd, K.E. (1993). "A Closer Look at Economic Psychology", Discussion Paper 9362 CentER for Economic Research, Tilburg University.
- Winter, S. (1981). "Attention Allocation and Input Proportions", *Journal of Economic Behavior and Organization* **2**, 31-46.