Analyse & Kritik 22/2000 (© Lucius & Lucius, Stuttgart) p. 114–129

*The EdK-Group**

Exit, Anonymity and the Chances of Egoistical Cooperation

Abstract: This paper presents the results of computer simulations with a community of actors playing a large number of voluntarily iterated two-person-PD. The simulations are designed to enable uncooperative actors to exploit partners, leave them and find a new partner who knows nothing about their previous behavioral history. Hit-and-run exploitation should thrive under these conditions. However, as Schuessler (1989; 1990) has shown, the setting is highly unfavorable to uncooperative players. The present study extends this result to a wider set of strategies which can alternatively stay with defectors (and try to improve them) or leave them quickly. In addition, a class of seemingly clever strategies is introduced which try to exploit the expected dynamics of looking for a partner. Still, a high amount of egoistical cooperation can persist in the present scenario.

0. Introduction

In the iterated Prisoner's Dilemma (= PD) the strategy of 'an eye for an eye' makes all see better. Robert Axelrod has not only established this result but brought it out of the retreats of game-theory and into the public realm. The basic problem to be solved, however, was old. Hobbes had already pointed out that rational egoists can become trapped in a war of all against all because they exclusively strive for their own benefits. Other political theorists have followed his lead. The age-old fear that human societies might cause their own breakdown by propagating rational egoism was spurred by the spread of free market economies in the 19th century. This development provided the background for general speculations of some classical sociologists (cf. Durkheim 1977; Toennies 1935) which culminated in the assumption that social order cannot emerge among pure rational egoists in any conceivable society. Without an effective normative socialization of the individual actor,

* This paper was written by Rudolf Schuessler and Ulrich Sandten. Contributions come from Carsten Bredanger, Oliver Doeschner, Eike Lange, Uwe Matzat, Sonja Nerad, Tilman Oeckinghaus, Thomas Schneider, Michael Steller, Attila Videk, Michael Zeman. These people formed the EdK-group, a research group at the University of Duisburg.

the social bond is bound to tear—say these theorists. It is a major aim of the rational-choice-program in modern sociology to prove this contention wrong, and the game-theoretical analysis of cooperation in iterated Prisoner's Dilemmas (PD) is one of the major roads towards this end, since the PD presents the Hobbesian problem of order in a structural nutshell.

A first big step on the way to a rational-choice explanation of social order was taken, when game-theorists confirmed the possibility of rational, cooperative solutions for the PD-supergame. These solutions emerge on the basis of threats of retaliation against exploitation which become possible in the iterated PD. But it is one thing to know that such threats are possible and another to demonstrate that they are credible even in the presence of attempts at evasion. It was Axelrod in *The Evolution of Cooperation* who stated this last point with verve and clarity. His computer tournaments were open to simulated players with exploitative strategies as complex as their author wished to contrive. Nevertheless, cooperation succeeded and TIT FOR TAT, a paradigmatically simple, friendly and self-defending contender, won the contests.

The cooperators' success in Axelrods tournaments, however, does not provide complete salvation from the fear of social disorder, not even if we put aside all the theoretical hostilities against the rational-choice paradigm and game-theoretical methods commonly found among sociologists. A typical supergame assumes that a PD is iterated among a fixed set of players. No player can leave or enter the game while a supergame is underway. No player can usually terminate the supergame intentionally. Therefore, rational actors know that they will face the same partners with some probability in the future, and that these partners hold the power to punish uncooperative behavior by refusing to cooperate with uncooperative players (=defectors). Under suitable conditions this knowledge creates a 'shadow of the future' which renders the cooperative behavior of all players rational. This was the good message, now comes the bad. The sketched boundary conditions of a typical supergame, which were also assumed by Axelrod, preclude a high mobility of actors. Modern societies, on the other hand, tend to be more and more mobile, in fact as well as potentially. The mobility and anonymity of the society to come were important ingredients of the gloomy scenarios of sociology's last fin de siècle. An answer to its worries, therefore, should probably include some inquiry into the fate of egoistical cooperation under these conditions.

This means that it is worthwhile to analyze iterated PD-scenarios which include an exit-option. One of our authors has investigated the effects of this additional strategic feature in earlier work (Schuessler 1989; 1990). Furthermore, research on the iterated PD with exit-option has become a small but steadily growing academic industry during the last decade. Most of the work in this area is not very closely related to our approach. Usually, different authors investigate different scenarios. The differences betweeen them is large enough to make it difficult to immediately compare results. Still, we will try to relate our findings to the results of others in the final section of this paper. Prior to that we will mainly present our own model, which is a sequel to the model of Schuessler (1989; 1990). This older model concentrated on the effects of anonymity on egoistical cooperation in an iterated PD with exit-option. Presently we mainly aim at a broader set of more complex strategies for very much the same scenario. Some strategies embody a kind of seemingly far-sighted opportunism which tries to anticipate and exploit the dynamics of the whole model. These strategies should help us to find out whether betting on presumed evolutionary tendencies in our scenario might impede egoistical cooperation. Therefore, it seems appropriate to sketch the original scenario first, before we proceed and add the pecularities of the present model.

1. The Basic Scenario and Earlier Results

Idealtypical exchange is dyadic and symmetric in regard to the strategic options of the participating actors. These options comprise non-violent fraud and taking advantage of informational asymmetries. The whole situation, therefore, assumes the structure of a two-person 2×2-Prisoner's Dilemma (cf. Rapoport/Chammah 1965) grounding in the question: deceive or not deceive? In consequence, one of the basic problems of exchange cooperation can be reduced to the problem of egoistical cooperation in an iterated PD.

In modern societies, voluntary interaction plays a central role both in norm and practice. Therefore, our model allows for a voluntary iteration of interaction between partners in a two-person PD, but also for a change of partners. All actors have an exit choice out of an ongoing interaction after each single game. If one of the players decides to terminate an interaction, both interacting partners have to look for a new 'spouse'. The model is designed to guarantee that all players who look for a new partner will get one in the next round of PD-games. Thus, all actors play a PD in every round. This assumption may seem arbitrary at first sight but note that the choice of partners is afflicted by anonymity in our model. Anonymity is implemented on the basis of certain restrictions on the information which players have. No player can obtain or remember information about an agent he is not currently playing with. No player can appraise the cooperativeness or non-cooperativeness of a new partner before their first mutual game. With these assumptions our approach differs sharply from others which base egoistical cooperation on some kind of 'translucency' of the players' characters and intentions (cf. Orbell/Dawes 1991). But note that our restrictions still do not entail complete anonymity. Players can remember how their partner behaved in a longer lasting partnership. Still, the kind of anonymity chosen justifies our assumptions about partner search. Since no one has any significant information about a new partner, there are no criteria for the selection of partners. It seems acceptable then, that all actors play a PD-game in every round, because nobody has an incentive to wait for better partners.

On the same grounds the assumption seems permissible that pairs of 'singles' are formed randomly with equal probability before a new round starts. Since the actors lack criteria for a purposive choice of partners, random matching is as good as any other procedure for the assignment of partners. As this does not seem to conform to the conditions of real life, it is important to recognize which analytic reasons guided our modeling decision. The model is designed for the investigation of just one cooperation generating mechanism in isolation. This explains, why reputation effects are ignored, although actors' reputations, of course, have an impact on partner choice in the real world. Furthermore, our model asks whether egoistical cooperation can survive even under presumably bad conditions. The model does not aim at representing actual strategic features of social reality but tries to formulate a bad-or worst-case-scenario. Therefore, we assume that leaving a partner and finding a new one involves no transaction costs or costs of any other sort. Significant separation costs are usually thought to lead to a supergame setting, but they are excluded here to get a clearer picture of the chances of purely voluntary cooperation.

Based on the mentioned assumptions the sequence of games for an actor can consist of an iterated interaction with the same partner as well as of a sequel of one-shot games with varying partners. This possibility devaluates threats of retaliation by defection which stabilize cooperation in a standard supergame. Cooperators have to face the risk of hit-and-run exploitation since the anonymity of the actors is postulated in the model. This anomymity is designed in exactly such a way that it allows defectors to take the profits of fraudulent action, move to a new victim and remain unidentifiable as defectors. New partners know absolutely nothing about the behavioral history of their companions. In this largely anonymous world the chances of cooperation seem bad. The common wisdom of gametheoretical cooperation analyses says that cooperative equilibria will not emerge if defectors can submerge in a "sea of anonymous others" (cf. Axelrod 1984, 100). And normativistic sociologists from the classics to modern communitarians also deny the possibility of social cooperation among anonymous, rational egoists (cf. Durkheim 1977; Parsons 1937; Tönnies 1935). But even well-entrenched dogmas need not be true. There is an effect that might support cooperation even in a world that knows no diffussion of actor's reputations.

This effect, which was discussed in Schuessler (1989; 1990), is a kind of adverse selection against uncooperative behavior in the voluntarily iterated PD. It emerged on grounds of the following behavioral assumptions. Let us assume that the first (one-sided or simultaneous) defective move in a sequence of PD-games with the same partner terminates the partnership. This seems to some degree plausible, since defectors may want to take the profit and run immediately, avoiding possible retaliative measures of an exploited partner. The exploited partner will also probably quit interaction with an uncooperative partner immediately, in order to punish her and find a new and better match. A pair of fully cooperative partners, on the other hand, will not separate, because nothing would be gained by it. Given these premises, the fraction of mutually cooperative, stable partnerships will increase with time, because cooperative matches do not break apart while new ones are formed round after round. Parallel to the cooperative actors' becoming absorbed in stable partnerships, the fraction of more or less uncooperative actors among the 'singles' increases. A growing number of new partnerships will therefore end with simultaneous defection. This also means that uncooperative actors will gradually hurt other uncooperative actors more likely than cooperative ones and thereby create an adverse selection effect against uncooperative behavior.

It is possible to speak of adverse selection in our model mainly because we assume an evolutionary dynamics for the diffusion of strategies. This dynamics may serve as a yard-stick for the success of behavioral strategies and its basic idea is simple. We assume with so many evolutionary approaches that players adopt, retain and give up behavioral strategies as a result of their success in gaining PD-payoff. It is often assumed that more than average success of the players of a certain strategy leads to a diffusion of the strategy among players whereas lower than average success leads to a vanishing of the strategy. There are several ways to implement this basic assumption and we chose to let top-ranking strategies gradually displace low-scoring ones. Thus, we arrive at a simulated process which selects egoistically successful strategies. Adverse selection against defectors implies in this context that the chances for defectors to score high deteriorate with time.

But can this effect make good for the losses to cooperators caused by hit-and-run exploitation? And can robust egoistical cooperation emerge merely on the grounds of the sketched adverse selection effect? The computer simulations of Schuessler (1989; 1990) have shown that strong and robust egoistical cooperation can emerge indeed, even when exit options are present and the chances for cooperators prima facie look dim. Similar studies have corroborated these results (Vanberg/Congleton 1992). In Schuessler's model, egoistical cooperation flourishes even if a moderate exogenous probability for the breaking apart of mutually beneficial partnerships is added (a catch-all-factor for external shocks, internal misunderstanding and trembling hands). Thus, egoistical cooperation seems to be less fragile than many observers would like to believe. Schuessler's results, however, provide only a first step to a fuller investigation of cooperation in the described scenario. They were based on the assumption that uncooperative actions immediately lead to a separation of partners. But not all actors will agree in ending a partnership immediately if a partner misbehaves. There is a very common conviction that one should stand by a misbehaving partner and try to correct him or her. In consequence, staying and trying to improve a partner should be taken seriously as a behavioral strategy. It seems natural then to expand the sketched model in the direction of allowing actors to continue partnerships and react as they see fit, when a partner defects. The present inquiry pursues this idea and asks, whether egoistical cooperation will also thrive under thus changed conditions.

2. The New Model

2.1 Basics

Like its predecessor, our new model is based on a voluntarily iterated two-person PD. The complete process of strategic interaction is a sequence of N rounds of PD-games among K players. It seems obvious that the number of players should be large enough to exclude effects of population size on the modeling results. Therefore, our computer simulations with the model were run with 4500 players, which is a sufficient number to ensure size-independency. Players can again stay with a partner and iterate a PD-game or change to a new partner at no cost after each round in the aforementioned way. The main difference between the old and the new model is to be found in the strategic behavior of players. Whereas the assumption of immediate exit on defection allows for a radical simplification of the strategy space and the linear ordering of strategies, the possibility of staying with a defector creates opacity. Therefore, successful strategies can only be discovered by trial and error starting from more or less informed guesses. The changed exit conditions also created an increased demand on a player's memory for storing strategically relevant data.

In our model, a vector of dimension ten is regarded as sufficient to satisfy this demand. The entries in this vector are mainly parameters derived from the past PD-choices of a partner in an on-going interaction. Due to anonymity, all entries which relate to past behavior of a partner are deleted after a separation has occured. In addition, parameters can also reflect an important feature of the social environment of the actors. The number of players who look for a new partner between two rounds is quite high in our scenario. These players thus form what we call a 'search pool', which varies in size during the game-process. Although we assume that the players are ignorant of the unengaged players' strategies and behavioral records, the size of the search pool after each round is known to all actors. The players can use this knowledge and their recollection of variations in the size of the search pool for strategic calculations. How they do this, can be gleaned from the description of strategies.

2.2 Strategic Options

2.2.1 General Options

The present model offers a multitude of behavioral options to the players. In a very general way they can opt:

1. to be friendly and not try to exploit a partner (or the opposite);
2. to stay with a partner or to leave;
3. to use or not to use certain kinds of information about partners or their environment.

Option 1 is a standard alternative in the repertoire of PD-supergame analyses. Like most other studies we have decided to compare friendly strategies (never first in playing D), primitive unfriendly strategies (always playing D), and more sophisticated unfriendly strategies (play D if some condition X is fulfilled).

Option 2 establishes the difference between analyses of voluntarily iterated PD-games and studies of inevitably longer lasting supergames such as Axelrod's. In order to inquire into the implications of the present model we looked at pure stayer strategies (never opt for leaving your partner), pure mover strategies (always leave your partner after a PD-round), and more moderate strategies whose staying or leaving depends on certain conditions. Note that a unilateral decision to leave always suffices to terminate a partnership. Players therefore cannot enforce their wish to continue a partnership.

Option 3 adds further complexity to the domain of possible behavioral strategies. Players have information about the moves (i.e. choices of C or D in a PD-game) their present opponents have made in the past rounds of jointly played PD-games. Apart from that, they are only able to screen how many players look for a new partner in a round. This means that the 'intentions' (so to say) of other players are not open to inspection and that the strategies others play are not known but can only be guessed on the described informational basis.

Concerning option 3 we assumed that it would be most interesting to analyze two categories of strategies. The strategies of the first category react to patterns in their opponents moves. TIT FOR TAT e.g. just mirrors the last move of its opponent (and in our model opts for prolonging their partnership). Other strategies look at the history of a partnership more thoroughly. They screen the history of an iterated game and condense it to reference numbers on which they base their decision to cooperate or defect, to move or to stay. The second category of strategies, on the other hand, contains members which use the number of players for looking for a partner as a parameter in their behavioral decisions. Strategies of this kind, which we call "milieu oriented", are designed to embody the following considerations: If the number of players looking for a new partner is unusually large, it may be assumed that more cooperative agents than usual will be found among them. Since the probability for successful hit-and-run defection increases, it seems prudent to follow a hit-and-run tactic in this situation. Obviously, such and similar considerations allow to represent sophisticated forms of opportunism, which are commonly regarded as very detrimental to cooperation among egoists. It will pay, therefore, to have a special eye on the performance of the milieu-strategies.

It is obvious, that a strategy could combine features out of every mentioned category of strategies in its decision procedure for its next move. This, indeed, happens in our model. The fifteen strategies, which were selected for our computer tournament, show how these features can be used to locate a strategy in a hypothetical space of strategic options.

2.2.2 The Strategies

Some of our fifteen strategies show rather intricate behavior and it would take some time to explain our rationale for adopting these strategies. In order not to impede the reception of our inquiry, we therefore have moved the description of the strategies to an appendix. Presently, it may suffice to introduce some starring actors of our contest. Among these are TFT and CONCO, representing two species of cooperative but easily provocable strategies. TFT is our version of TIT FOR TAT, which begins cooperatively, reciprocates cooperation and defection, and always opts for staying with a partner. CONCO always plays C and tries to stay with cooperative partners but leaves as soon as the partner begins to defect. On the unfriendly side of the strategic spectrum we have SALLD and MALLD. Both strategies always play D, but SALLD (= Stay-ALLD) never opts for leaving a partner itself, whereas MALLD (= Move-ALLD) changes partners after every round. MIL3, finally, stands in marked contrast to such primitive defectors, being a rather mean and opportunistic milieu-strategy. MIL3 has an alert eye on the number of players looking for new partners. If it regards this number as low, it tries to stay with its present partner. Nevertheless, MIL3 tries to exploit this partner by starting defective, but chooses cooperation after the first round of an iterated game, if the partner proves provocable. This pattern of behavior changes, however, if the number of players who look for new partners is regarded as high. In this case, MIL3 sees a chance to exploit naive singles. It defects and quits its present partnership, in order to play MALLD (i.e. hit and run) with new partners until the overall number of singles decreases again. Then, MIL3 once more will seek shelter in a stable partnership. We have looked at several other opportunistic strategies of the MIL-variety (see appendix), but it may suffice to introduce MIL3 because of its success. Whether it was successful enough to challenge more friendly cooperators like TFT and CONCO will soon be discussed.

2.2.3 The Evolution

In our computer tournament players adopt one of the above listed strategies and compete for PD-payoffs. The measure of their success, however, is not just the sum of payoffs they receive. We followed the growing trend of evolutionary modeling of strategic competition. The focus of evolutionary modeling is usually on the success of strategies and not on the success of players. A strategy is regarded as succcessful if the number of players that play it increases with time or stays above some pre-defined level. A strategy is unsuccessful if it is adopted by fewer and fewer players or by some pre-defined small proportion of all players. The success of a strategy in turn depends on the payoffs it receives. In general, a monotonical relationship between the rank of a strategy as measured in payoff gains and its chances of propagation is assumed. In our model, the relevant ranking parameter is the sum of payoffs in a series of n consecutive games (we assumed n = 30). Thus, evolutionary selection does its work after every n-th round of a contest. The 10 percent of the players which score worst according to the ranking give up their strategy and are assigned a new one out of the listed fifteen with equal probability. Of course, the memory of those players is emptied, too. The whole procedure therefore may reflect the bancruptcy on the part of some players and the entrance of new players into the market who do not know in advance which strategy will suit them.

Our selection mechanism is markedly different from the one used in dynamical game-theoretical models in biology. There, the evolutionary rise or decline of strategies is governed by a system of difference or differential equations known as game-dynamical equations. In social science applications, however, it remains an open question which selection mechanism is appropriate. It is to be remembered that we no longer talk about the chances of sexual reproduction of players but about the diffusion of strategies due to learning or trial-and-error. In the absence of empirical proof that some mechanism depicts social reality truthfully we have to proceed on grounds of theoretical considerations. The most important consideration seems to be, given the uncertainty about the right selection math, that the qualitative results of evolutionary investigations should be robust in the face of a change of selection equations. Consequently, we have experimented with several equations and our results proved astonishingly robust, indeed. For the sake of simplicity, however, we will only discuss the results of one mechanism—and a very simple one, too, since we have no apriori reason to prefer a more complicated rival.

3. Results

On the basis of the described modeling assumptions several computer simulations were conducted with a population of 4500 players and sufficient rounds for the evolutionary process to get a grip on the strategies. The focus of our interest lay on changes in the number of players that play a certain strategy. Picture 1 shows how this number changed in the evolutionary process when all 15 strategies competed for success. All strategies started with the initial share of 1/15 of the overall population (= 300 players). Picture 1 is split up into three parts following the ranking of strategies for reasons of perspicuity:

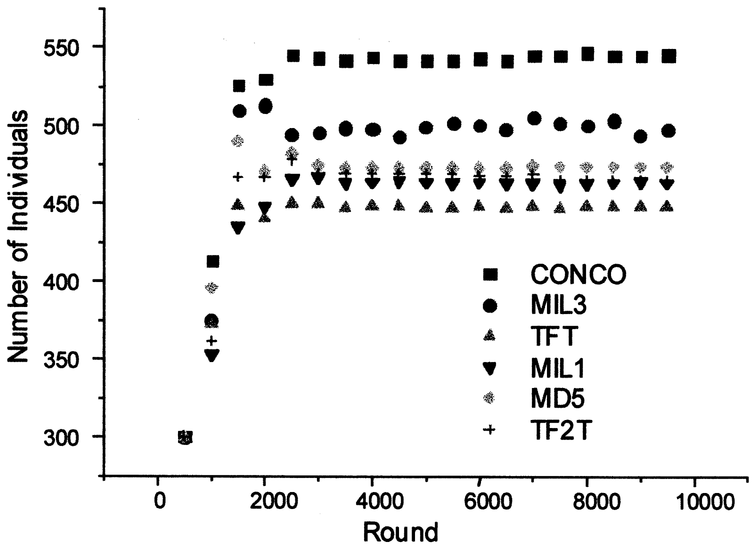

Picture 1a

A first notable feature of our simulations was that a group of six top ranking strategies (Picture 1a) emerged, each of which finally had more than (roughly) 450 players. CONCO, the simple exit-strategy, stands out as winner of the contest with 546 players. It is followed by MIL3 with 498 players. This means that the good guys win, but a not so nice, clever opportunist comes in second and well in front of the cooperative hero TFT. Remember that MIL 3 starts a new partnership with an attempt at exploitation and leaves a partner cooly if the environment seems suitable for preying on naive cooperators. On the other hand, it abandons its exploitative behavior after the first round of a partnership and even accepts an atonement payment. This docility accounts for its thriving among the pugnacious cooperators which dominate our evolutionary scenario. Next in rank are MIL1, another opportunistic milieu-strategy, MD5, a blend between CONCO and TFT, and finally TF2T and its cousin TFT. Note that TFT scored worst relative to the members of the leading group but still may count on the respectable number of 449 adherents.

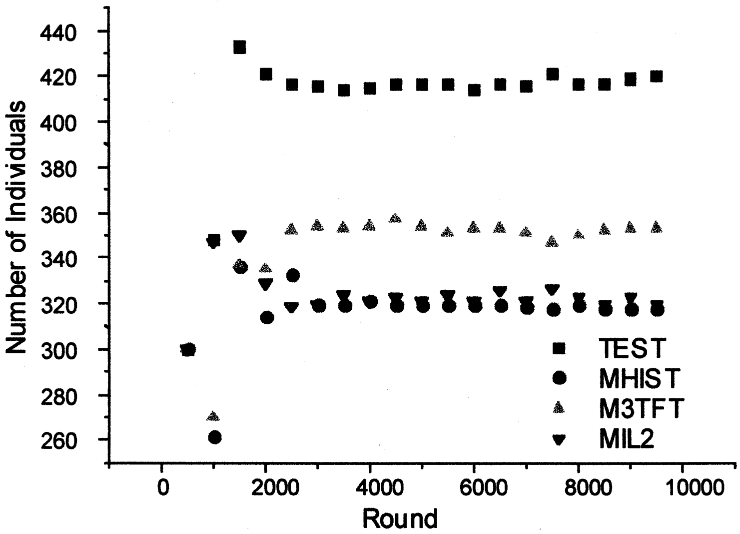

Picture 1b

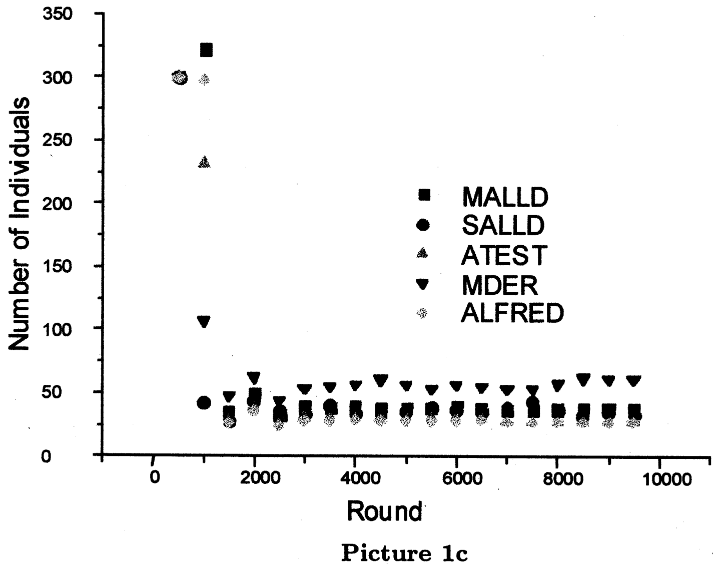

Picture 1c

Behind the leading group a payoff-gap opens up. With the exception of TEST (420 players) the next cluster of strategies can be found in the range between 300 and 350 players. It includes MHIST, which is the only strategy that keeps track of all intercation with a partner, and, thus, confirms that behavioral history does not count much in contests of rational actors. Even lower on the list, there is a bottom cluster of strategies with less than 100 players, which comprises mainly rude defectors of the ALLD variety. No strategy dies out completely because of the random assignment of strategies to new competitors.

For an overall interpretation of this picture the following points seem to be important. Often, game-theoretical analyses fail to distinguish clearly between TFT, as embodying the retaliatory principle of 'an eye for an eye', and strategies which rely on an exit option. If TFT is modelled as a retaliatory stayer and brought together with serious rivals it is simply not the best strategic choice. CONCO, although it is also a very simple cooperative strategy, is a superior alternative in situations where exit comes easy. Given the background of Axelrod's contest, there is some interest in a direct comparison of CONCO and TFT, since both are simple, nice strategies. In order to have a closer look at these two, we arranged for a special test, which confronted them singly or together with the pure defectors SALLD and MALLD and the opportunistic milieu-strategy MIL1. As can be seen from picture 2, CONCO clearly won the joint competition with roughly 2200 over the 1400 players of TFT.

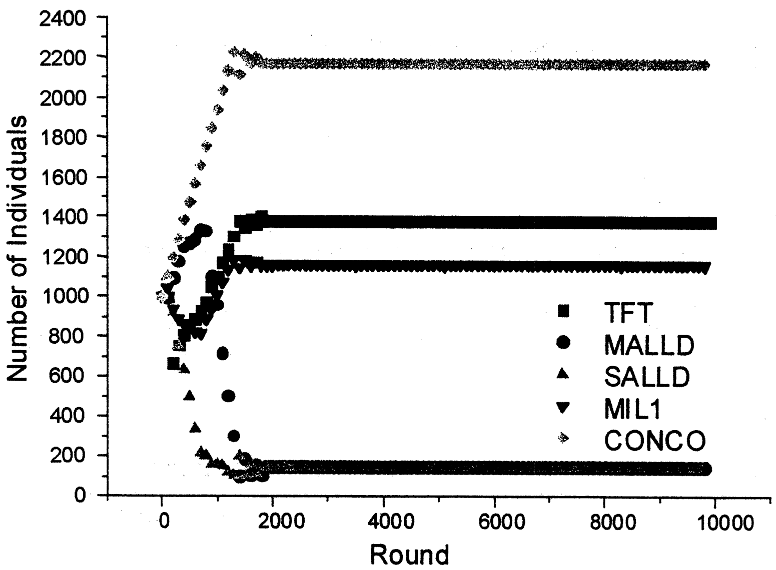

Picture 2

When CONCO and TFT competed separately with SALLD, MALLD; and MIL1 both managed to win, although CONCO did a much better job in keeping MIL1 in check (Picture 3a,b).

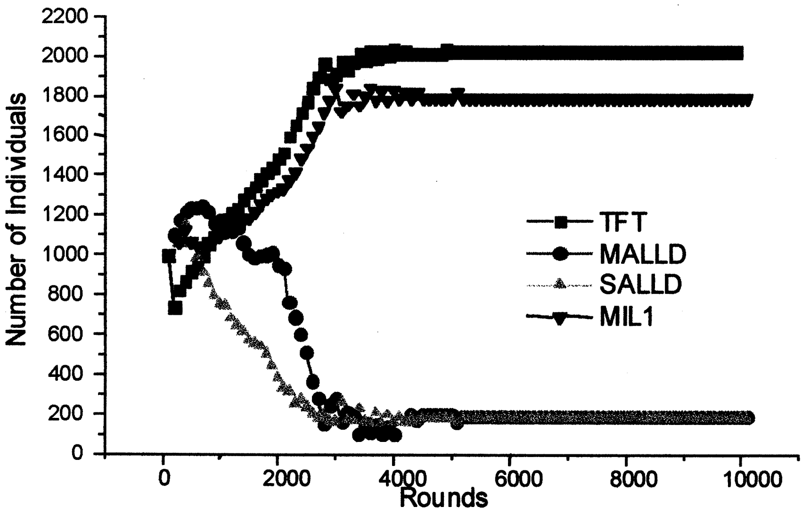

Picture 3a

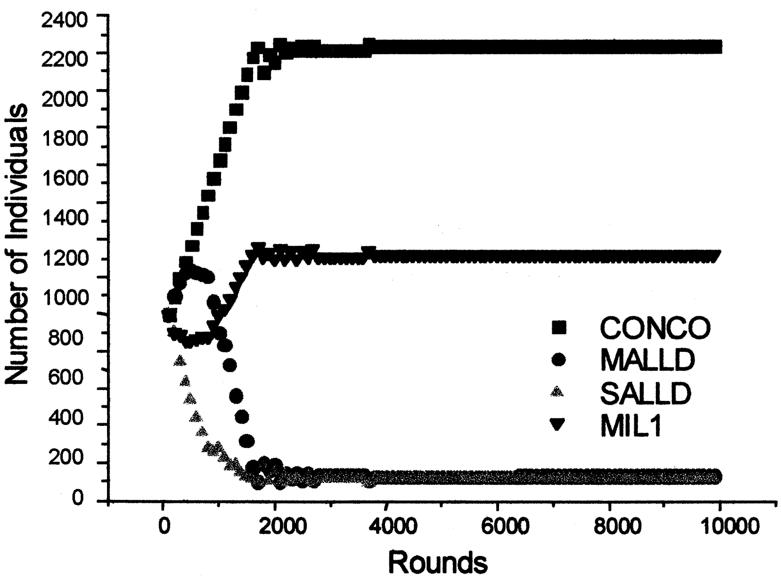

Picture 3b

4. Discussion

The last section dealt with the results of our computer tournament in some detail. If seen from a very general angle its results confirm the main findings of earlier research. The presence of an exit-option in iterated PD-games facilitates egoistical cooperation. The now well documented force of adverse selection against uncooperative actors in iterated PDs with exit-option suffices to outweigh the benefits of exploitative hit-and-run behavior. One important factor on which this selective effect is based is the mobility of the actors. In our model, all actors were highly mobile. Actors looking for a new partner could potentially mate with every other actor who is also a single. Majeski et al. (1997, 164) arrive at a similar judgement when they stress the importance of mobility as a cooperation generating factor in their spatial model, in which actors can move in and out of neighborhoods. One should, however, distinguish between potential and actual mobility in all models. This means, that the chances to move and the amount of observed mobility in a model world play different roles concerning the emergence or non-emergence of egoistical cooperation. Whereas no increase per se in potential mobility seems to damage egoistical cooperation, too much actual mobility will destroy it.

Despite the success of the cooperators in the present model it should also be noted that clever opportunism does pay. Thus, cooperators have to be on their guard and fit for action if cooperation is to remain stable. In an environment of provocable cooperators, however, the opportunists tend to act nicely, having no chance to profit from exploitative behavior. This result calls for second thoughts on the side of theorists who believe that there is an evolutionary bonus on genuine rule-guided morality as compared to opportunistic moral-like behavior. Such a bonus can only form a basis for ethical argument if it prevails under a variety of social conditions. It should therefore presumably show up in our model. The success of opportunists of the MIL-kind, however, speaks against the existence of a bonus for rule-guided morality. In our model it pays to assume the face of morality if threats of retaliation are strong but to be willing to exploit others if things should change. Of course, it also pays to be cooperative and provocable. But this just means that provocability is the decisive factor in generating cooperation and being cooperative as a rule does not count for much.

This finding is in tune with others that predict only limited cooperativeness in Iterated-PD-tournaments with an exit-option and random matching, whereas preferential partner selection leads to improved cooperativeness (cf. Ashlock et al. 1996). Note, however, that although sophisticated opportunists can be successful in our model, a large part of the actor population practises egoistical cooperation. In addition, their cooperation is very robust. It prevails under different sets of strategies and a wide range of parameter values. We found it in differently programmed models, under different evolution equations, and even when the tournament started with a high amount of uncooperative actors. The present model ist especially designed to test this robustness of egoistical cooperation, since it contains opportunistic strategies which try to exploit the expected evolution of the whole model. Nevertheless, they fail to curb egoistical cooperation decisively. Things may look different if 'exit' in an iterated PD means that players do not play the game but receive a 'waiting payoff' instead. Macy/Skvoretz (1998; 651) state that under this condition the value of the 'waiting payoff' determines largely whether egoistical cooperation will evolve or not. As noted, the anonymity of new partners in our model allows us to exclude the option of not playing. It is the possibility of preferential partner selection which provides a rationale for the tactics of waiting for an attractive partner. This shows how complex the strategic situation in an iterated PD with an exit-option is. The possibility of preferential partner selection tends to support egoistical cooperation as Majeski et al. have found. It also makes it worthwhile to wait for attractive partners. The latter option, however, makes egoistical cooperation dependent on 'waiting payoffs' according to Macy/Skvoretz. Success in reducing the anonymity of actors will, therefore, not necessarily improve the chances of egoistical cooperation. This is so because, as our model proves, anonymity in itself is no deadly threat to egoistical cooperation.

Bibliography

Ashlock, D./M. Schmucker/E. Stanley/L. Testfatsion (1996), Preferential Partner Selection in an Evolutionary Study of Prisoner's Dilemma, in: *BioSystems 37*, 99–125

Axelrod, R. (1984), *The Evolution of Cooperation*, New York

Durkheim, E. (1977), *Über die Teilung der sozialen Arbeit*, 1. Aufl. 1893, Frankfurt

Macy, M./M. Skvoretz (1998), The Evolution of Trust and Cooperation Between Strangers, in: *American Sociological Review 63*, 638–660

Majeski, S./G. Linden/C. Linden/A. Spitzer (1997), A Spatial Iterated Prisoners Dilemma Game Simulation with Movement, in: R. Conte/R. Hegselmann/P. Terna (eds.), *Simulating Social Phenomena*, Berlin, 161–167

Orbell, J./R. Dawes (1991), A Cognitive Miser Theory of Cooperators' Advantage, in: *American Political Science Review 85*, 515–528

Parsons, T. (1937), *The Structure of Social Action*, Glencoe

Rapoport, A./A. Chammah (1965), *Prisoner's Dilemma*, Ann Arbor

Schuessler, R. (1989), Exit Threats and Cooperation Under Anonymity, in: *Journal of Conflict Resolution 33*, 729–749

— (1990), *Kooperation unter Egoisten*, München

Tönnies, F. (1935), *Gemeinschaft und Gesellschaft*, Leipzig

Vanberg, V./R. Congleton (1992), Rationality, Morality and Exit, in: *American Political Science Review 86*, 418–431

Appendix

SALLD:	defects always and never opts for leaving a partner.
MALLD:	defects always and opts for exit after every round.
TFT:	begins playing C, answers C with C and D with D in the next round, never opts for leaving a partner.
TF2T:	begins playing C, answers C with C and two subsequent Ds with D in the next round, never opts for leaving a partner.
TEST:	begins playing D and cooperates in round 2 and 3 of a partnership. If the partner chooses C in the third round, TEST will play D in round 4 and alternate between C and D in all later rounds. If the partner defects in round 3, TEST will continue as TFT.
CONCO:	cooperates always but leaves a partner immediately after his first defection.
M3TFT:	begins as TFT but opts for leaving a partner after his third defection.
MD5:	cooperates always and opts for leaving only if the partner has defected five times or more.
ALFRED:	defects always, tries to stay as long as the partner cooperates but leaves a parter after receiving a D. [No insult to bearers of that name intended.]
ATEST:	plays two times D and opts for leaving after the second round with a partner.
MHIST:	cooperates always and opts for leaving only if the ratio of defective to cooperative choices of a partner is lower than a value k^* (here: $k^* = 1.5$).
MIL1:	acts like TFT if the search pool has not grown by a factor f^* (here: $f^* = 1.1$) relative to the previous average of all played rounds in the whole simulation. Otherwise, MIL1 will defect and leave its partner, playing MALLD in subsequent rounds until the search-pool-size is again lower than f^*-times its average. In this case MIL1 tries to stay with a new partner and plays TFT again.
MIL2:	acts like MIL1 except in cases where MIL2 leaves a shrinking search-pool. In this cases it starts a new partnership with a D-choice and subsequently plays TFT trying to stay with the partner unless the search-pool satisfies the condition stated for MIL1.
MIL3:	acts like MIL2 but with the following changes. MIL3 reacts by playing ALL C to a retaliatory defection by a partner.
MDER:	acts like TFT unless the search-pool increases by a factor f^{**} relative to the last round (here $f^{**} = 1.05$). In this case MDER defects and leaves its partner in order to play MALLD unless the search-pool stops growing with factor f^{**}. Then, MDER accepts a new partner like TFT again.