CONCLUSIONS

Chapter 11 What can be Learned?

The present work had several goals. The first was to explore the potential effect of individuals' motivation for fair outcomes on behavior in an asymmetrical social interaction, described with the repeated investment game. Second, the performance of various strategies should be assessed by analytical exploration and evolutionary simulations. The third goal was to find decision strategies that are appropriate in describing individuals' decision processes in the repeated game.

Special interest was devoted to simple strategies, and different methods were used to investigate them: on the one hand analytic exploration and evolutionary simulations, and on the other experimental investigation of individuals' behavior. Simple strategies can be used for coordinating on a particular equilibrium, thereby solving an equilibrium selection problem. Simple strategies are also psychologically plausible descriptions of the decision process. Their fit in predicting individuals' behavior was compared to other models of decision making, in particular to a learning model. Besides being psychologically plausible models of bounded rationality, simple strategies can coincide with the "rational" game-theoretical predictions. People are smart when they use these strategies since they require only little memory and computational power yet still lead to good outcomes.

11.1 Trust, Reciprocity, and Fairness

In all experimental studies substantial trust and returns expressing motivation for fair outcomes were observed. For the conducted repeated games this is not surprising given that repetition allows one to withdraw trust, and thereby, to punish the opponent. Individuals in a social interaction such as the investment game observe the possibility to increase their payoff and seize this opportunity. Participants in a one-shot investment game may also observe the monetary incentives offered by trusting their opponent, but the risk of losing a great deal is also much higher. This is a marked difference to the repeated game, where the payoffs of only one single period are negligible. Given this high risk, observed trust without repetition may be a more impressive demonstration of cooperative individuals, but it may also be less representative of common relationships in which the risk of being exploited is moderate because of repetition, allowing high trust that is not blind.

Consistent with the equity theory that predicts that individuals are motivated to reach fair outcomes, the most frequent behavior of participants in all experiments was a return rate that led to equal final payoffs. The influence of a motivation for fairness could be nicely demonstrated with the second study in which the endowment for player B was varied in two conditions. This treatment had a strong effect on the return rates, which were higher in the condition with no endowment, consistent with the equity principle. The omission of an endowment for player B changes the interpretation of what is considered a fair distribution and, thereby, substantially reduces the return rates.

However, it is reasonable to argue that the decision to return a fair amount in an ongoing interaction is also motivated by self-interest, since even self-interested individuals may wish to maintain good relationships, which would mean allowing both interaction partners to profit. If one takes this perspective, fairness does not necessarily express unselfish motives but is a mechanism used to coordinate a relationship. Fairness considerations appear in a broader sense as a device for coordination between different equilibria, and therefore, they solve equilibria selection problems. Following this argument, fairness motives are not in conflict with an evolutionary perspective on human behavior.

In contrast, from an evolutionary perspective reciprocity and fairness in anonymous situations without ongoing interactions are difficult to explain (Caporael, Dawes, Orbell, & Van de Kragt, 1989; Elster, 1989; Nowak et al., 2000). This does not mean that fair returns, which have been observed in the investment game without repetition, are not also expressions of a fairness motive. However, this motive might vanish if all our real-life interactions became anonymous, without repetition.

Study 2 also demonstrated that the investments can be influenced by a motivation for fairness. The results weakly support the prediction that participants invest more, on average, if player B has no endowment than if player B has an endowment. This result is surprising if one argues that participants trust their opponents only because they want to increase their payoffs. This opportunity is reduced if one assumes that player B will return less if not provided with an endowment. However, if participants in the role of player A are motivated to reach fair outcomes they have to make an investment in order to make equal payoffs for both players possible. This result is in line with a study by van Dijk and Vermunt (2000), which shows that individuals occasionally make higher offers to opponents in the dictator game, in which the opponent is totally powerless, than in the ultimatum game, in which the opponent has the power to reject offers made by the first player. However, Forsythe et al. (1994) found conflicting results.

A broader definition of reciprocity implies that beneficial behavior of another individual is positively reciprocated with beneficial behavior and uncooperative behavior is negatively reciprocated with punishing behavior. From such a broad definition a positive correlation between the investment rate and the return rate follows. If low investments are regarded as small contributions to the social interaction, such a correlation is implied by the equity principle. On the contrary, the equality principle (i.e. all get the same regardless of their contributions) predicts the opposite. Whereas no substantial correlation between the investment and return rates could be observed in the first study using the one-shot investment game, in the repeated games substantial correlation were observed. This supports the interpretation that reciprocity, including negative and positive reciprocal behavior, is promoted by repeated interactions. In general, the observed association between the investment and return rates supports the argument that a motivation for fairness depends on the behavior and intentions of others (see also Falk & Fischbacher, 2000; Rabin, 1993).

11.2 Evolutionary Simulations

Do strategies that perform well in asymmetrical social relationships differ from those that work well in symmetrical relationships? This question is of particular importance if social relationships in which the roles of individuals are not interchangeable are very common. The relationship between an employer and an employee is a prototypical example. There are many other relationships where the consequences of the decisions for the interaction partners differ substantially (e.g. patient vs. physician, child vs. father). Therefore, the prisoner’s dilemma cannot be used as a general model of social interactions. However, the majority of evolutionary research explaining cooperative behavior has focused on the prisoner’s dilemma. For this reason, an extension to asymmetrical games is required.

In more detail, the conventional game-theoretical prediction for the indefinitely repeated investment game is quite unsatisfying because it predicts too many outcomes and leaves open how they are reached. Finite state automata as repeated game strategies could explain how particular payoff combinations are obtainable. To investigate how different finite automata as strategies for the game compete against each other the Nash equilibrium concept and the evolutionary stability concept were applied. It was shown that even for a small set of strategies a large number of equilibria exist. Therefore the evolutionary stability concept appears to be more appropriate to evaluate strategies. However, for the indefinitely repeated investment game no strategy could be classified as an evolutionary stable strategy. This means it was not possible to find a strategy such that if a population of agents uses only this strategy it is not possible that agents using another, mutant strategy could invade the population. Given that it is plausible that individuals make unsystematic errors when they use a strategy the concept of Limit evolutionary stability was additionally applied for evaluating strategies. It turned out that the Never-Invest and Return-Nothing strategy combination and the Min-Grim and Return-Min strategy combination were Limit ESS for the selected set of strategies. This result was supported by the evolutionary simulations in which all finite automata with a restricted number of states were allowed as strategies. It turns out that the Limit ESS strategy combination was frequently observed in the evolutionary processes when no application errors occurred. However, if application errors were introduced, the Limit ESS strategy combination of Min-Grim and Return-Min was no longer observed. In contrast, if a maximum of four states were allowed, the strategies that incorporate mechanisms that tolerate errors by the opponent and produce efficient outcomes were frequently obtained. However, no single strategy could be found to be the predominant outcome of the evolutionary process.

In sum, moving the attention of evolutionary simulations away from symmetrical games (e.g. prisoner’s dilemma) to asymmetrical games like the investment game influences the type of strategies that evolved. “Kind” strategies like Tit-for-Tat that enable the players to return to a cooperative state after uncooperative decisions were seldom found for the asymmetrical game. Additionally the observed strategies differed for the different roles in the asymmetrical game. To obtain efficient outcomes in the asymmetrical game, it is necessary to implement some kind of punishment mechanism for player A’s strategy and, thereby, prevent exploitation from player B. In contrast, for player B a constant return rate implemented with, for instance, a Min-Return strategy is sufficient. Most strategies for the investment game producing efficient outcomes could be classified as variations of the Grim strategy. From these results it follows that trust and reciprocity will emerge less frequently in asymmetrical social interactions than in symmetrical ones, and a profitable interaction will be more vulnerable to mistakes. If only one individual in an interaction has the possibility of exploiting the other individual, exploitation becomes more destructive. If an individual exploits by accident, the caused loss for the other individual cannot be compensated. Retrieval of an efficient outcome requires a risky disbursement by the individual who was exploited.

11.3 Comparing Simulations With the Experimental Results

For study 4 the same automata used in the evolutionary simulations were used to model individuals' decision processes. This method appears very fruitful. The representation of the strategies as finite automata makes an interpretation very amenable. Although the strategies turned out to be rather simple they could have been more complex as the number of states for both player roles was restricted to four states. This maximum number of states allows a large number of automata and combinations of automata to be selected as models for the decision process. Because of the incredibly large search space it is difficult to find the optimal model. Hence, a genetic algorithm, which is known for its high effectiveness in solving different optimization problems (Goldberg, 1989), was used for selecting good-fitting models. It may also be valuable to follow this approach in other areas for deriving models of cognitive processes. Although the tools we use influence the way we think about psychological processes (Gigerenzer, 1991), it is reasonable to have a variety of tools to prevent building theories that are too strongly influenced by one single approach.

How are the strategies found in the experiment of study 4 for player A related to the strategies found in the evolutionary simulations of study 3? The strategies, which could predict a substantial proportion of participants' decisions, establish an efficient outcome by investing the entire endowment if substantial returns are made repeatedly. Individuals that apply a strategy like Moderately-Grim express initial confidence in the other player by investing the entire endowment. The Min-Grim strategy found in the evolutionary simulations shares important similarities with the Moderately-Grim strategy. It also starts with the investment of the entire endowment and repeats this investment if substantial returns are made. Both strategies also share the terminal state; if the strategy reaches this state no investments are made in all following periods. Moderately-Grim differs from the pure Grim strategy by an additional state that incorporates a forgivingness component. If player B once makes a substantial return then in any following period Moderately-Grim tolerates a low return once. Only after repeated low returns, unless a low return is made in the first period, will the strategy move to the terminal state with no investment. In sum, both strategies share main features, but the strategy developed to predict individuals' behavior is more tolerant toward low returns.

A strategy similar to Hesitant, which was selected to describe participants' decisions in the role of player A, was found in the evolutionary simulations. This strategy, called Cautious, was obtained in the evolutionary simulation in which execution errors occurred with low probability. Cautious starts with an investment of 0% in the first period, increases it to 100% in the second period, and for all following periods maintains the investment at 100% if a substantial return is made. If a low return is made Cautious moves to the third state with no investment for all following periods. There are only two differences between Hesitant and Cautious. First, Hesitant starts with an investment of 50% compared to the 0% of Cautious, and second, if Hesitant reaches the third state with no investment it always returns to the first state in the next period, whereas Cautious stays in the third state for all following periods. Therefore, once again the strategy selected for the participants is more forgiving, as it is possible to return to an efficient situation after low returns are made.

In sum, the strategies selected for predicting individuals' behavior are similar to strategies that perform well in evolutionary simulations. The main difference between the strategies selected as models for individuals' behavior and the strategies found in the evolutionary simulations consists of higher forgivingness concerning low returns incorporated in the strategies for the participants. The most striking similarity between the strategies selected for predicting individuals' behavior and those in the evolutionary simulations is the aspiration level for player A strategies. All strategies use an aspiration level that guarantees a payoff that is only above player A's endowment. This is again an aspect that shows how tolerant the participants appear when confronted with low returns.

However, whereas the strategies for player B in the evolutionary simulations were adapted to the low aspiration levels of player A's strategies, so that they made only sufficient low returns, the strategies selected to predict participants' decisions in the role of player B make much higher returns. The first strategy, Reactive, selected for the participants in the role of player B makes a return of 70% if player A makes a substantial investment. The second strategy, Half-Back, also makes a high return of 50% if a substantial investment is made. Only if player A makes an investment lower than 12% (respectively 17%) no return is made. Consistent with the results of the evolutionary simulations, the strategies selected for the participants in the role of player B turned out to be simpler (i.e. having fewer states) than the strategies selected for player A, as they consist of only two states. Compared to the Min-Return strategy, which was frequently obtained in the evolutionary simulations and led to an efficient outcome, the strategies for the participants incorporated a second state. The second state for both strategies incorporates a punishment mechanism, as it leads to no return if only a very small investment is made. The most striking difference is the amount of return that is made if an efficient relationship is established. Whereas the Min-Return strategy only makes the minimum return that is necessary to give player A a payoff above the endowment, the strategies for the participants incorporate high returns, which for one strategy leads to equal payoffs for both players.

In sum, a comparison of the strategies developed to describe peoples' decisions and those found in the evolutionary simulations shows some striking similarities, but the strategies developed for the individuals incorporate some extra mechanisms. The strategies for player A incorporate a forgivingness mechanism and the strategies for player B incorporate a punishment mechanism. The forgivingness mechanism will tolerate eventual single exploitations, which could be the result of mistakes. The punishment mechanism leads to no return if a very low investment is made and thereby could be used for initiating high investments.

The evolutionary simulations have shown that simple strategies are sufficient for obtaining efficient outcomes and are not outperformed by more complex strategies. Although individuals presumably do not deliberate on the game-theoretical analysis of the folk theorem and its predictions, they may be equipped with strategies that are consistent with the theorem's predictions. They frequently yield efficient outcomes with equal final payoffs for both players. In comparison to the stable fairness principles of how monetary payoffs should be allocated (e.g. Bolton & Ockenfels, 2000; Deutsch, 1975; Loewenstein, Thomson, Bazerman, 1989; Wagstaff, 1998), the strategies developed for study 4 are dynamic and take the decisions, and thereby the intentions, of other individuals into account.

11.4 Payoff Distribution

There is a substantial difference in how the payoffs are distributed between the two players if one compares the experimental results with those of the evolutionary simulations. Although in the experiments participants in the role of player B obtained a higher payoff than participants in the role of player A, the difference was not as large as one would predict from the evolutionary simulations, in which player B obtained almost all of the surplus. A closer look at participants' decisions in each period of the game shows that the modal distribution of payoffs was an equal payoff for both players.

The result of equal payoff for both players obtained frequently in the present experiments is in line with many other experimental studies, like those of one-shot dictator or ultimatum games (Güth et al., 1982; Ochs & Roth, 1989). Experimental results in these games deviate from the game-theoretical prediction that concedes to the proposer in the ultimatum game or the dictator game all or almost all of the endowment. Most proposers offer 30-50% of the endowment, and most responders reject offers lower than 30%. Although this shows a tendency of proposers to make use of their bargaining power, the results strongly deviate from the game-theoretical prediction. The investment game shares some similarities with the dictator game: One could interpret the investment game as a modified dictator game in which player A's decision is a preliminary stage in which player A makes a decision about the amount of money with which he wants to play the dictator game with player B. Player B then functions as the proposer in the dictator game. One could also interpret the investment game as an ultimatum game in which player A takes the role of the responder. An investment of player A is similar to unconditionally accepting a proposal of player B. The similarities between the investment and ultimatum game become more apparent with repetition of the game. If the game is repeated, player A will always make an investment of 100% if both players agree on splitting the trebled investment according to the proposal of player B. If player A rejects player B’s proposal, this implies that no investment will be made in the following period. In the one-shot ultimatum game, it is the proposer who has the higher bargaining power. Accordingly, one could argue that it is player B who has the higher bargaining power in the repeated investment game, which is consistent with the results of the evolutionary simulations.

How can the difference between the evolutionary simulations and the experimental results concerning the payoff distribution be explained? One could point to individuals’ representation of the social interaction of the investment game. This representation in an experiment might deviate from the abstract characteristics of the investment game. If individuals interpret the experimental interaction of the investment game as a “reciprocity domain” in which individuals regard themselves as equal (see Bugental, 2000), then they will allocate resources according to the allocation principles they are equipped with for this domain. There might be a strong tendency to allocate resources according to the equity principle or the equality norm if individuals represent a social interaction as a reciprocity domain. This interpretation can be supported by the fact that the social context in which a social interaction is usually embedded was removed in the experiments. The investment game, for instance, resembles the relationship between an employee and employer, or between a patient and a physician. The problem for participants in an experiment is that by eliminating the context, it gets more difficult to understand the different bargaining powers and less easy to justify them. One might argue that participants in an experiment regard and treat themselves as equal if they are not provided with reasons to think otherwise, which could be provided with a particular context.

The present work illustrates how individuals are motivated by fairness and how fair allocation in ongoing interactions is attainable with simple strategies. I agree with Bugental (2000) that people use different "social algorithms" to solve the problems they are faced with. Bugental (2000) and similarly Fiske (1992) give only rather vague descriptions of these algorithms. In contrast, in the present work precise strategies were provided. These strategies were able to predict a substantial proportion of individuals' decisions in a particular social interaction. They could explicate the underlying dynamic process that is involved in a situation of trust, reciprocity, and fairness. These simple and precise strategies can be though of in the traditional notion of essential cognitive tools for decision making. Simple strategies enable us to make reasonable decisions and behave adaptively in our social environment—humans would be lost without them.