Chapter 4 Behavior in Repeated Games

Behavior in repeated games has been extensively studied with the prisoner's dilemma. In general, repeated games can be distinguished as finitely and indefinitely repeated games.[3] Finitely repeated game with complete information do not fundamentally differ from games without any repetition, as by the backward induction argument the game-theoretical prediction is the same as for the one-shot game. However, in experimental games in which a game is played for a very long period people often do not follow the backward induction argument; instead they might look forward a few periods (Camerer, 1997).

Rapoport and Chammah (1965) describe their study of the repeated prisoner's dilemma in their well-known book called "Prisoner's dilemma." In their experiments 140 participants played the prisoner's dilemma for 300 periods but were told they would play the game "a large number of times."[4] In their experiments the influence of different payoffs was studied, which all incorporated the payoff structure of a prisoner's dilemma. Additionally they varied whether participants played the prisoner's dilemma only with one single underlying payoff matrix or with changing payoff matrices.

It turned out that the proportion of cooperative decisions varied substantially between the different experimental conditions. The smallest magnitude of 19% of cooperative choices was found in a condition in which a high incentive to exploit the other player or a high risk of being exploited existed (the defecting player gained 50 and the other player lost 50 points; mutual cooperation (defection) yielded only a gain (loss) of 1 point for both players). Additionally in this condition the payoff matrix was not presented to the participants and the game was played for only 50 periods; thereafter another payoff matrix was used. The largest magnitude of cooperation with 77% was found in a condition in which participants played the prisoner's dilemma 300 times with one single-payoff matrix, which was presented to them and which yielded a payoff of 9 points (1 point) for mutual cooperation (defection) and a gain (loss) of 10 for diverse decisions. Rapoport and

[3] Some authors prefer the term infinitely repeated game. However, as this would imply that the game does not end at any time I prefer the term "indefinitely," which emphasizes that the game ends at some time but the exact period in which it ends is not defined. For the repeated prisoner's dilemma the term "iterated" prisoner's dilemma has also been established.

[4] This method has the problem that it is difficult to derive the equilibrium prediction for the game. The game has to end at some time and participants will form beliefs about the probability that the game will end in the next period. Only if these probabilities, which are unknown to the experimenter, are sufficiently large, the equilibrium prediction will allow cooperation between the players.

Chammah (1965) emphasized the dynamic process of the repeated game and pointed out that in the initial phase of the game participants extended their defective decisions but after some time they "recovered" to more frequent cooperative decisions again. They also were the first to describe the dynamic process with a stochastic learning model and compared it to a Markov chain model, and they proposed that future research should seek the "mathematical model which among many tested models describes most accurately" people's decisions. This proposal has recently received intensive interest by several authors. Roth and Erev (1995; see also Erev & Roth, 1998), for instance, have proposed a reinforcement learning model for zero-sum games and in Erev and Roth (1999) they suggest how the model could be used for social dilemmas. Camerer and Ho (1998) propose a combination of belief-based and reinforcement learning models that takes individuals' expectations about the behavior of other individuals into account. Both models will be described in more detail in chapter 10.

If individuals play a game for more than 100 periods one could question how seriously each single decision is taken. In Rapoport and Chammah's (1965) experiments, for instance, the payoff for the equilibrium prediction varied between minus and plus one cent, hence the incentives for one single decision were negligible. Given these low incentives and the fact that the game was played for 300 periods it is reasonable to assume that participants did not defect in the prisoner's dilemma only to increase their payoff but also to make the game more interesting. Additionally it can be questioned what kind of real-life interactions are represented with a repeated game of 300 periods, in which all decision are performed in a relatively short period of time (i.e. 1 hour, yielding 12 seconds per decision). In the present work I will study interactions that last for a relatively short period, so that higher incentives for each decision can be provided, and which might lead to deeper thinking about the strategic situation.

Deutsch and Lewicki (1970) investigate a short-term interaction with the "trucking game," which is similar to the "chicken game." The well-known chicken game is a two-person symmetrical bargaining game with two Nash equilibria consisting of different decisions of the players (see Figure 4). The chicken game is quite interesting as it is not a "social dilemma" with an inefficient Nash equilibrium. Instead the chicken game has two efficient equilibria. However, if an equilibrium is reached one player obtains a higher payoff than the other. Therefore the important aspect of this game and of similar "coordination games" is whether and how the players can coordinate on one equilibrium.

		Player B	
		A	B
Player A	A	3,3	*2,4*
	B	4*,2*	1,1

Figure 4. Chicken game.

The game has two equilibria indicated by the two cells in which both players' payoffs are marked with asterisks.

Deutsch and Lewicki (1970) studied the trucking game under two conditions, one without repetition and another with 20 periods of repetition (this number was not told to the participants). It turned out that participants were much more efficient in the repeated game, as it enabled them to cooperate by alternating between the two equilibria.

Short-term interaction has also been studied by Roth and Murnighan (1978). They conducted an experiment in which the indefinitely repeated prisoner's dilemma was played with different continuation probabilities of 0.1, 0.5, and 0.9, yielding a small number of 1.1, 2, and 10 expected periods of the game. The equilibrium prediction includes no cooperative decisions for the first continuation probability contrary to the second and the third continuation probabilities that allow cooperation. Consistent with the equilibrium prediction the number of cooperative choices was much higher in the second and third conditions, even though the proportion of 36% cooperative decisions was quite small in the condition with the largest continuation probability.

Selten and Stoecker (1986) investigated a finitely repeated prisoner's dilemma of 10 periods. However, participants played the repeated game 25 times against different opponents, thereby making it possible to investigate whether the experience participants gained would change their behavior in later games. Note that because the length of the games was known to the participants the equilibrium includes no cooperative decision in any period. It could be observed that participants in the first periods cooperated but this cooperation broke down after one participant initiated defection, which was then followed by mutual defection for the rest of the game. As the game got closer to the end, participants in general started to defect. This behavior is called an "end effect." Interestingly, through experience, participants gained when playing the repeated game repeatedly, this end effect occurred earlier and earlier, hence initial defection occurred in earlier periods of the game.

I assume that in real-life social interactions it is the rule that individuals involved in an interaction expect that they will meet the other person again. At least individuals are not sure about whether they might meet each other again in future and it is seldom the case that one expects that the interaction is the definite last one. If individuals encounter in an experiment a repeated game with a definite end this is a rather artificial situation and therefore it is not surprising that with experience individuals will adapt their behavior to the novel situation, which explains the change of behavior in Selten and Stoecker's experiment.

Van Huyck, Battalio, and Walters (1997) studied the indefinitely repeated investment game (which they called the peasant-dictator game) with different endowments and different rates of returns of the investments. They showed that when the investment game was repeated indefinitely with the same opponent the investment rates increased substantially up to an average of 80% of the endowment in a situation in which any investment of the endowment was quintupled, and when the participants had already gained a lot of experience. It also became clear that the return rate depended on the endowments. An equal split of the final payoff was frequently observed if the payoff for player A obtained by an equal split exceeded player A's endowment. In contrast, participants did not agree on an equal payoff if this payoff would be lower than their guaranteed endowment. This result demonstrates the primacy of "self-interest" above "fairness." In the case of low rates of returns of the investments the investments decreased substantially, leading to Van Huyck et al.'s (1997) conclusion that reputation (as possibly established through repetition) is only an imperfect substitute for binding contracts. This conclusion seems to be quite questionable for two reasons. First, in their experiment the rate of returns of the investments were on average quite low, giving only low incentives for investment. Second, the expected number of periods for each game was small, which makes it quite difficult to build up a reputation.

In summary, the research reported above focuses on behavior in repeated games and in particular on the magnitude of cooperative behavior. Most studies investigated how the magnitude of cooperation varied between different conditions. Several studies followed the goal of increasing the proportion of cooperative behavior with a particular treatment. In their review of experiments using the prisoner's dilemma Pruitt and Kimmel (1977) summarized altogether 12 aspects of a social interaction that make the goal of reaching mutual cooperation more likely. They also point out several aspects that increase the expectation that the opponent will cooperate in future interactions. Hence, past research could point out several conditions that support or undermine cooperation. However, what is still rare are studies that describe individuals' decision processes in ongoing social interactions.

One exception is a study by Wedekind and Milinski (1996) in which participants in an experiment played a 20-period prisoner's dilemma against an opponent who played a predetermined strategy. The participants did not know the number of periods, nor that the opponent used a predetermined strategy. Two predetermined strategies were used: The first strategy is to cooperate in all periods except the 16^{th}, in which a choice to defect is made. The second strategy is to cooperate in the first period, then to do what the opponent did in the previous period, and then starting from the sixth period, always to defect. By using the predetermined strategies it was possible to investigate participants' behavior in response to a fixed behavior and thereby, as the authors argue, to infer the kind of strategies the participants followed. They showed that for most individuals (70%) the decisions were more consistent with the so-called *Pavlov* strategy than the Tit-for-Tat strategy. The Pavlov strategy is to cooperate if both players made the same decision in the previous period, otherwise to defect. Both strategies and their advantages have been investigated in evolutionary simulations reported in chapter 6. Although the study is quite restricted by investigating only a small set of potential strategies individuals could use, it provides insight into the decision process of individuals. For instance, if a person is classified as a user of the Pavlov strategy, this implies that if the person made a decision to defect in the previous period and the opponent made a cooperative decision then the person will repeat the decision to defect in the present period. This means that the person does not hesitate to exploit the opponent but does hesitate to initiate cooperation.

Another example of a study that aims to investigate the decision process is a recent study by Engle-Warnick and Slonim (2001). The authors examined a repeated trust game in finite repeated and indefinitely repeated conditions (for the trust game see Figure 2 on page 10). Player A can either trust or distrust player B, who subsequently decides to exploit or reciprocate a trusting decision.

The authors proposed a set of strategies to model participants' behavior. The criteria for proposing the set of strategies included the maximum complexity of the strategies, past experimental evidence on repeated games, and post-experimental protocol responses of participants in the experiments, so that finally the set appeared a bit arbitrary. However, they showed that one single strategy was appropriate in explaining participants' behavior in the role of player A in the indefinitely repeated trust game. This strategy starts with a trusting decision in the first period and repeats this decision as long as player B reciprocates trust; if player B exploits player A, the strategy chooses to distrust player B in all following periods. For player B a larger set of three strategies was found as appropriate to describe participants' behavior. The strategy with the highest fit basically always reciprocates a trusting decision of player A. For the finitely repeated games it turns out that special strategies were necessary to describe the end effect of the game, which occurred in the experiment and were similar to the end effect founds in Selten and Stoecker (1986). These strategies incorporated distrust and exploitation for later periods of the game.

As done in the previous studies, I also aim to investigate the decision strategies people use in a bargaining game, in particular in the indefinitely repeated investment game. In contrast to the studies reported above, I will investigate a larger set of strategies and I will use only a complexity criterion to restrict the set of strategies. The strategies that are best in predicting participants' behavior are selected as the model for the individuals' decision processes.