CHAPTER 1

Introduction

> "Againe, men have no pleasure, […] in keeping company, where there is no power able to over-awe them all. For every man looketh that his companion should value him, at the same rate he sets upon himselfe: And upon all signes of contempt, or undervaluing, naturally endeavours […] to extort a greater value from his contemners, by dommage; and from others, by the example."
>
> Thomas Hobbes, Leviathan

How can cooperation flourish among egoistic individuals? Thomas Hobbes argued that only a powerful institution, the omnipotent Leviathan, can sustain cooperation! Indeed, many collective actions humans engage in are organized by institutions that determine and enforce rules: States provide education and build streets, companies produce and sell products, and clubs entertain and support their members. However, many interactions we engage in are not supported by institutions, or existing institutions do not determine and enforce rules of cooperation. For instance, a group of collaborating students cannot always exclude a shirking group member, or a project manager cannot easily replace a dodging project member. The question is, therefore: How do groups achieve and maintain cooperation when they cannot rely on institutions?

When institutions are excluded as a means for insuring cooperation, only characteristics of the situation—such as the frequency of interaction, the efficiency gain through cooperation, or the structure in which individuals interact—or properties of individuals—such as their motives, personality, or the decision rules they use—remain as determinants of cooperative behavior. This dissertation focuses on the mechanisms of cooperative decisions. Additionally, I will investigate how interaction structure affects cooperative outcomes and the choice of decision strategies.

When using the term mechanism, I refer to the cognitive process underlying cooperative decisions. While the literature on cognitive processes in individual decision making is numerous (siehe z.B., Betsch, Haberstroh, & Hohle, 2002; Fellows, 2004; Hastie, 2001; Holyoak & Spellman, 1993; Payne, Bettman, & Johnson, 1992; Rieskamp, Busemeyer, & Mellers, in press; Shafir & LeBoeuf, 2002), less is known about the cognitive processes underlying interdependent decision making, including cooperative decision making. Nevertheless, established approaches exist that differ in their assumption about players' individual rationality and in the generality with which they can be applied to various decision domains. (With "individual rationality" I refer to rationality defined as maximizing individual payoffs by applying rules of logic.)

Game theory, as a descriptive theory, assumes unboundedly rational players, who have complete knowledge about the structure of the game (modeling the interaction) as well as rational beliefs about others' behavior and their beliefs (Colman, 1999). According to game theory, players logically infer optimal decisions in a game, based on their knowledge about the game and their beliefs. While game theory proved to be a useful tool to analyze interdependent decision making, its value as a descriptive theory is undermined because people frequently violate the predictions of game theory (Colman, 1999, 2003a). One important reason for the failure of game theory as a descriptive theory is its unrealistic assumptions about decision-makers' individual rationality. When comparing models of decision making in cooperative groups, I will therefore focus on boundedly rational mechanisms of decision making.

A prominent approach to boundedly rational decision making in games is learning (z.B. Fudenberg & Levine, 1998). Learning, as I use the term in this dissertation, means to adapt behavior based on one's own and on others' experience. I will not examine how individuals achieve insight in the strategic structure of an interaction, or how individuals invent new strategies to cope with the incentive structure and others' behavior. I will rather assume that individuals already dispose of a set of alternative actions (e.g., cooperate or defect), and learning then determines which of these alternatives is chosen.

An alternative approach to bounded rationality is that people use decision rules or heuristics (Axelrod, 1984; Gigerenzer, Todd, & the ABC Research-Group, 1999; Komorita & Parks, 1999; March, 1996) to make cooperative decisions. While the learning approach assumes that individuals learn within a game which choice is the best, the heuristic approach assumes that players enter situations with a decision rule that is applied in the same manner for repeated decisions. While the learning models I consider are domain general—they can be applied to different types of decisions, such as, for example, individual or interdependent decisions—the decision rule tested will be a domain specific social heuristic.

Simple decision rules often perform astonishingly well (e.g., Axelrod & Hamilton, 1981; Martignon & Hoffrage, 1999; Nowak & Sigmund, 1993; Payne, Bettman, & Johnson, 1993), and can also describe the behavior of many individuals (e.g., Rieskamp & Hoffrage, 1999). An underexplored question is how individuals acquire decision rules. One general claim is that decision rules are adapted to the decision environment (Gigerenzer et al., 1999) and to performance criteria, such as accuracy and efficiency of decisions (Payne et al., 1993). These criteria constrain the set of possible decision rules, but do not describe the mechanism by which individuals acquire specific rules. Mechanisms that can explain how people acquire decision strategies are evolutionary adaptation (e.g. Cosmides, Tooby, & Barkow, 1992), individual learning (e.g., Rieskamp & Otto, submitted for publication), and social learning (e.g., Joseph Henrich & McElreath, 2003). Probably all three mechanisms partly explain how individuals acquire decision rules. Some decision mechanisms or their building blocks might be innate, whereas others are learned individually or from others. Interestingly, research on these mechanisms usually examined different learning mechanisms in isolation, neglecting the possibility that, for instance, a combination of individual and social learning might explain how people acquire decision rules. Hence, I will examine learning mechanisms that combine individual and social learning.

How do people maintain cooperation in groups? Starting with this question, this preface introduced individuals' decision mechanisms and the decision environment as determinants of cooperative behavior. Accordingly, the main goals of my dissertation are to identify the decision mechanisms behind individuals' cooperative decisions in groups, to show how they relate to different aspects of the decision situation, and to examine learning mechanisms with which people can acquire decision rules. More specifically, I will first test if general learning models or a reciprocity-based heuristic describe cooperation in groups better. In a second step, I will further examine if reciprocators imitated others' intentions or the consequences of others' behavior. Finally, I will test if social learning improves individual decision making, and examine different models combining social and individual learning. Before examining these questions in detail, I set the stage by introducing the public goods game used to examine cooperation in groups, by describing existing models of interdependent decision making, and by depicting accounts of how people learn to make decisions.

1.1 Cooperation in Public Goods Games

Cooperating in its broader sense means "to act or to work with another or others" or "to associate with another or others for mutual benefit" (*Merriam-Webster Online Dictionary*, 2005). I examine cooperative behavior in a narrower sense, namely, cooperation in social dilemmas (Dawes, 1980), where individually rational behavior leads to a suboptimal outcome for the group. Cooperating in a social dilemma means sacrificing some of one's own benefit in order to achieve a cooperative outcome that is beneficial to every group member. Cooperation is difficult to maintain in social dilemmas because noncooperative group members cannot be excluded from the benefits of others' cooperation.

Social dilemmas are a useful tool to examine cooperation because cooperation in these dilemmas are costly. Only if cooperation is costly and can increase one's payoff, can alternative explanations of "cooperative" behavior be excluded. For instance, contribution to a charity should be considered as altruistic because the contribution cannot make the contributor better off. Any interaction that is organized by a legal contract should not be regarded as cooperative because the parties can enforce others well-behavior. Typical examples for social dilemmas are public broadcasting (in the US, but not in Germany), worker unions, working groups with variable remuneration based on group performance, or syndicates developing industry standards. All these examples have in common that a contribution to the public good is voluntary and beneficial to the whole group, while no member of the group can be excluded from the benefits of others' cooperation.

1.1.1 Public Goods Games

The best examined social dilemma is the Prisoners' Dilemma game, in which two players have to decide if they cooperate or defect. As in most social dilemmas, both players are better off if they defect. However, if both defect, the worst possible outcome actualizes. The divergence of individual rationality ("I am always better off when I defect") and collective rationality ("When we both defect, we have together the worst possible outcome") has made the Prisoners' Dilemma game the standard framework for research on cooperative behavior (for a review, see Pruitt & Kimmel, 1977). While research on Prisoners' Dilemma games has generated a great amount of knowledge about cooperative behavior, this game is a two-person game and hence allows only limited insight into cooperation in groups. To examine cooperative behavior in groups, the Prisoners' Dilemma game was generalized to more than two persons. In the so-called n-person dilemma game (Dawes, 1980), three or more players decide about their contribution (i.e., cooperation) to a public good. N-person dilemmas are also called public goods games. As a public goods game is the most frequently used term across disciplinary boundaries, I will henceforth use this term[1].

Four main characteristics distinguish the types of public goods games: (1) the number of players who decide about their contributions to the public good and who will benefit from its provision, (2) the production function determines how individual contributions are transformed into payoffs from the public good, (3) the contribution mechanism describes the possible

[1] Usually two types of social dilemmas are distinguished. Those in which players have to decide how much they contribute to a public good, and those where players have to decide how much use they make of a common resource (here, cooperation means to use little). The former games are called n-person dilemmas, give-some games, or public goods games, the latter are called take-some games, common dilemmas, or resource dilemmas. While in psychology, economics, and sociology the general term for cooperative situations is social dilemmas, the term collective action is more frequently used in political sciences. The common feature of all social dilemma/collective action situations is the conflict between individual and collective rationality.

contribution levels, and (4) the protocol of play describes the order of contributions to the public good.

The public goods game I use in my dissertation research is a four-person public good with a voluntary contribution mechanism. Specifically, all four players decide simultaneously how much of their endowment they contribute to the public good. The production function multiplies players' contributions and then divides the product equally among the four players. Formally, N = 4 players are endowed with *e* points. Every player decides about the size of his or her contribution *c*. The payoff of player *i* is defined as

$$\pi_i = e_i - c_i + MPCR \cdot \sum_{i=1}^{N} c_i \ . \tag{1.1}$$

Here, $MPCR \cdot \sum_{i=1}^{N} c_i$ is the production function with $\sum_{i=1}^{N} c_i$ as the sum of all players' contributions and MPCR as the marginal per capita return. The game is a social dilemma when contributing is costly, that is, the MPCR has to be smaller than 1, and when contribution is efficient, that is, the MPCR has to be larger than 1 divided by the number of players. Table 1.1 depicts the payoff matrix for a four-person public goods game with endowments of 5 points and a MPCR of .5. [2]

The cells of the matrix depict payoffs for player A and mean payoffs for the other players, given the contribution of player A and the mean contribution of the other players. The possible strategies (i.e., the possible contributions) of player A are in the columns of the matrix. Each row stands for possible mean contributions of the three other players in the game.

Because games are an abstract representation of human interaction, some question their external validity. Appendix B discusses which insights about cooperative behavior in groups can be gained, by using public goods games.

---

[2] An alternative contribution mechanism is a dichotomous (yes/no) contribution, alternative protocols of play are sequential contributions or cumulative sequential contributions, alternative production functions are step level functions, where no public good is created if a minimum of players do not (e.g., Rapoport & Eshed-Levy, 1989), and a concave contribution function, where the size of the public good first increases with contributions, but decreases again for high contribution levels (e.g., Plott & Laury, 1998).

Table 1.1. *Payoff matrix for a four-person public goods game with a MPCR of .5.*

row player's contribution	Column player's contribution					
	0*	1	2	3	4	5
0*	5 5	6.5 4.5	8 4	9.5 3.5	11 3	12.5 2.5
1	4.5 6.5	6 6	7.5 5.5	9 5	10.5 4.5	12 4
2	4 8	5.5 7.5	7 7	8.5 6.5	10 6	11.5 5.5
3	3.5 9.5	5 9	6.5 8.5	8 8	9.5 7.5	11 7
4	3 11	4.5 10.5	6 10	7.5 9.5	9 9	10.5 8.5
5	2.5 12.5	4 12	5.5 11.5	7 11	8.5 11.5	10 10

Note. The first column and row depict possible contributions. Numbers in the upper (lower) left corner depict the row (column) player's payoff when the other players contribute on average the value in the first row (column) of the (row) column. Contributions marked with an asterix are equilibrium strategies. That is players cannot increase their payoffs when individually deviating from these strategies.

1.1.2 "Rational" Contributions in Public Goods Games

How much should a player contribute to a public good? The answer to this question is usually provided based on a game theoretic analysis. In order to derive optimal behavior in a game, the game theory makes assumptions about the players. Generally, players are assumed to be individually rational, that is, they "act according to their preferences and relative to their knowledge and beliefs at the time of acting" (Colman, 2003a). Hereby, *rational preferences* are complete, transitive, and their ordering is context independent (i.e., preferences have a weak order), *rational beliefs* are internally consistent, and *rational arguments* follow the rules of logic. In the framework, a decision is rational when it is instrumental, that is, it leads to the goal of the decision maker. These assumptions are implemented in the expected utility theory (von Neumann & Morgenstern, 1947), which assumes that players choose actions maximizing the anticipated outcome. As expected, utility theory is a general theory, additional assumptions need to be made to apply it to interdependent games. The first assumption (common knowledge of rationality *I*) is that players know the parameters of the game, which includes players' feasible strategies, the payoff function, as well as everything that can be logically deduced from strategies and payoffs. The second assumption (common knowledge of rationality *II*) is that players are instrumentally rational as described by the expected utility theory. It is part of the common knowledge of rationality *I*, that players know about others' instrumental rationality.

Of course, all these assumptions can be questioned (and tested). It might well be that a player does not try to achieve the best possible outcome in a given payoff matrix, for instance, because he or she is altruistic (Fehr & Fischbacher, 2003)[3]. It is also questionable if—for games that are more complex than social dilemmas—players have complete knowledge about all possible strategies and associated payoffs. It seems reasonable to question that players have higher order beliefs (I believe that you believe that I believe…), and use these beliefs to derive individually rational decisions (Colman, 2003b; Stahl, 1996). Finally, many have questioned if players are able to apply the rules of logic as expected from individually rational players.

Even though many assumptions of game theory have not withstood the empirical test, a game theoretic analysis remains valuable. First, while the assumptions do not hold for individuals, they might hold for groups (see, e.g., Bornstein, Budescu, & Zamir, 1997; Bornstein, Kugler, & Ziegelmeyer, 2004), and predict behavior of interacting groups adequately. Second, while individuals might not infer the same conclusions as a game theoretic analysis, repeated play of the same game might assist people to learn and to converge to the game theoretic solution of the game.

The most important concept to predict the outcome of a game is the Nash equilibrium (e.g. Colman, 1999). The Nash equilibrium states that individually rational players will choose the best response to the best possible strategy choice of the other player(s). To derive the Nash equilibrium for the public goods game of Table 1.1, player A first sees that—ceteris paribus his or her own contribution—the payoff for the other players increases when they decrease their own contributions. As this is true for every contribution level of player A, the rational contribution level for the other players is zero. The same holds for player A; regardless of others' contributions, his or her own payoff increases when he or she decreases their contribution. Hence, the Nash equilibrium of this game, the set of mutually best responses, is that all players contribute nothing, leading to a payoff of 5 for all players. Note that this outcome, which follows

[3] In this case, however, the monetary payoffs in the payoff matrix would no longer reflect players' utilities.

from individual rationality, is the worst possible outcome for the *group*. That is, while defection is individually rational, it is collectively irrational because other combinations of strategies (i.e., contributions) would lead to a better outcome at the group level. Again, this conflict of individual and collective rationality is the key characteristic of social dilemmas.

While the analysis above holds when players play only one round of a public goods game (a one-shot game), one might argue that cooperation is rational when the game is repeated because players should cooperate in order to convince others that mutual cooperation is possible. If the game is repeated, for instance, for 100 rounds, it seems to be reasonable to encourage cooperation to achieve the collectively superior cooperative outcome. But the motive of encouraging cooperation is rationally not sustainable if the game is repeated for a finite number of repetitions. If players are in round 99, there is no reason to motivate others' cooperation, hence, everyone will defect. If cooperation cannot be rationalized in round 99, the same holds for round 98, and then also for round 97, etc., and, finally, also for the first round. Hence, a rational player will infer, by backward induction (Luce & Raiffa, 1957), that defection is the best choice in finitely repeated public goods games.

The result changes if one considers infinitely repeated games, where after each round the game is continued with a certain probability. For infinitely repeated games, the Folk Theorem (Colman, 1999) shows that cooperative strategies can be the best response strategies, that is, they can be individually rational. For instance, the grim trigger strategy, which contributes everything in the first round and continues contributing everything as long as others did the same in the preceding round, but otherwise defects for ever, can be an equilibrium strategy and is thus individually rational for the infinitely repeated public goods game. Appendix A gives a numerical example for rational cooperation in public goods games. But even though the Folk theorem shows that cooperation can be individually rational in infinitely repeated games, it does not solve the problem of cooperation. This is because the equilibrium always consists of a set of best response strategies. While I illustrated individually rational cooperation with the grim trigger strategy, there are many more possible equilibrium strategies that players can choose from. Then the problem arises that players need to choose *compatible* repeated game strategies, they need to coordinate. This coordination problem has not been solved so far.

1.1.3 Empirical Contributions in Public Goods Games

The preceding section showed that cooperation in public goods games should be observed only rarely and only in indefinitely repeated games, if players are rational in the game theoretic sense. This section briefly reports on the results of public goods experiments examining players' contributions, and summarizes what is known about factors influencing contribution levels (for comprehensive reviews, see Kollock, 1998; Komorita & Parks, 1995; Ledyard, 1995).

Game theory makes the clearest prediction for one-shot public goods games with linear production functions, where the only rational behavior is not to contribute anything[4]. In an early test of this prediction, Dawes, McTavish and Shaklee (1977) found that, depending on specific conditions, between 27% and 84% of the players cooperated. The same basic result was found by Marwell and Ames (1979), who controlled for a number of factors, such as the distribution of endowments, group size, heterogeneity of benefits, provision points, and the economic training of the participants. In their experiments, players contributed, on average, 57% of their endowment (41% if one excludes players whose endowments were higher than the provision point in cases where a step level public good was played). Hence, early results, which have been replicated many times, indicate that players' contribution decisions in public goods games cannot be described by the standard game theory.

The same general result holds for contributions observed in finitely repeated public goods games, for which the backward induction argument, introduced above, predicts that individually rational players never cooperate. Beginning with Isaac, Walker, and Thomas (1984) and Isaac, McCue, and Plott (1985), researchers examined contributions to public goods that were usually repeated for 10 rounds. The general finding for finitely repeated public goods games is that contribution levels in the first round are usually approximately 50% and then decline to approach the Nash equilibrium of zero contributions toward the end of the supergame (Ledyard, 1995). Dependent on the theoretical preference, these results are interpreted in favor of game theory because players learn to be rational, or they are interpreted to the disadvantage of game theory because players do not follow the equilibrium inferred by backward induction (which is never to contribute). To my knowledge, no Experiment exists that reports contribution levels over time in indefinitely repeated public goods games.

Experimental results show that game theory (alone) cannot explain participants' contribution decisions in public goods games. Therefore, researchers examined many different aspects of public goods games and players in order to understand contribution behavior. The next paragraphs will report on how different aspects of public goods games, namely, game parameters, repetition, interaction structure, information, communication, and framing influence contributions. Properties of players are discussed in section 1.2.

---

[4] In the public goods game, with a step level production function, cooperation can be rational.

The important game parameters influencing cooperation are the MPCR, N, and—in the case of step level public goods—the provision threshold. Obviously, higher MPCRs lead to higher cooperation rates because they reduce the negative effect from being exploited and, at the same time, increase the efficiency of contributions (see also Komorita, Chan, & Parks, 1993; Rapoport, 1967). Experimental evidence for the positive effect of higher MPCRs was provided, among others, by Isaac, Walker, and Williams (1994) and Komorita et al. (1993). Results on the effect of group size are mixed. While some argue that smaller groups have higher contribution levels (e.g. Bagnoli & McKee, 1991; Chamberlin, 1974; Kerr, 1989), others found that larger groups contribute at the same or at a higher level (Isaac et al., 1994). The reason for these contradicting results lies in the interaction of group size and MPCR for the determination of payoffs. Holding the MPCR constant while increasing the group size leads to higher payoffs when all group members cooperate; working against possible effects of lower perceived self-efficacy in larger groups. Hence, one should adjust the MPCR when changing the group size. This approach was applied by Isaac et al. (1984), who found contribution rates of 75% for N = 4 and MPCR = .75, and of 33% for N = 10 and MPCR = .3. However, reducing the MPCR for the larger group, in order to align maximum payoffs, increases the negative effect of being exploited in larger groups, so that it remains unclear if, in this case, lower contributions of large groups were forced by this fact or by reduced self-efficacy. The last game parameter showing a clear influence on contributions is the provision threshold in step level public goods.[5] The general finding here is that while players increase their contributions when the provision threshold is higher, the public good is provided less frequently because the increase in contribution does not correspond to the increased threshold (Isaac, Schmidtz, & Walker, 1989).

Repetition in public goods games can take on different forms. One approach is that the same participants are reinvited to the laboratory to play another public goods game, but with different group members. Isaac et al. (1984) and Andreoni (1988) show that experience decreases contributions. Another possibility is that players participate in several public goods games within one session, either in the partners' design—where players repeatedly participate in the same group—or in the strangers' design—where players are randomly rematched for one-shot public goods games. While the first Experiment, comparing the partners' and strangers' result, found—counterintuitively—that contributions decrease more for partners (Andreoni, 1988), later studies found higher cooperation levels in the partners' design (e.g. Croson, 1996).

[5] Note that public goods with linear production functions are n-person Prisoners' Dilemma games, hence, the individually rational choice is not to contribute. In contrast, step level public goods are n-person chicken games, hence, cooperation can also be rational.

While most experimental public goods games treat public goods as if they are created isolated from other environmental aspects, real public goods are often provided in a larger context. For instance, a group generating a public good might do so to compete with another group, or players might be able to choose to which public good they want to contribute. In the intergroup Prisoners' Dilemma (Bornstein, 2002), two groups contribute to public goods, and individual payoffs depend on the participants' own contributions to the public good and on the payoff of their own group in the Prisoners' Dilemma. The Prisoners' Dilemma has a payoff structure, so that higher contributions to the public good from the groups are equivalent to more competitive choices (i.e., defecting), which were "individually" rational. Bornstein (2002) found in several experiments that embedding the public goods game in a Prisoners' Dilemma structure increases cooperation beyond simple group identity effects. Others examined the effect of "partner selection" on cooperation in the public goods game. In an experiment by Coricelli, Fehr, and Fellner (2004), participants could vote with which other person, out of 15 other players about whose contribution history they were informed, they wanted to form a four-person public goods group. Contributions in their partner selection condition were higher than in groups that were randomly matched. In a similar experiment by Ehrhart and Keser (1999), nine participants started in three-person public goods groups, and could then decide to which group to switch. In this experiment, a process developed where free riders switched to groups of cooperators, thereby reducing the success of cooperative groups, leading cooperators to leave their old group and to reassemble new cooperative groups, which were again "invaded" by free riders, etc. In sum, partner selection increases cooperation in the public goods game.

To further examine variables influencing cooperation in the public goods game, contributions can be observed if one relaxes the standard setup of public goods experiments. Among others, Orbell, Dawes, and van de Kragt (1990) could show that communication increases cooperation in one-shot public goods games. Isaac and Walker (1991) showed that this result also holds when communication is costly, and the game is played repeatedly. Sell and Wilson (1991) provided their participants either with no information, aggregate information, or individual information. Results showed higher contribution levels in the individual information condition, compared to the other conditions. Croson and Marks (1998) provided their participants with individual information, so that individual participants could either be tracked or not tracked across rounds. Interestingly, results showed that anonymous individual information leads to a rapid decline of contributions, whereas identifiable individual information maintains higher levels of cooperation.

1.2 Models of Cooperative Decision Making

The previous section introduced the public goods game, explained which behavior is individually rational, and contrasted this with observed behavior in these games. The most apparent conclusion drawn from experimental results on public goods games is that players' choices cannot adequately be described by standard game theory. Hence, this section introduces alternative accounts of behavior.

Before introducing prominent accounts of cooperation, I briefly give an overview about some alternative theories of cooperation. Pruitt and Kimmel (1977) proposed the expectancy value theory of cooperation which states that players cooperate if they have the goal to cooperate and if they believe that others will cooperate. While this approach is supported by some evidence, the problem remains that the elicitation of beliefs about others' cooperativeness changes behavior. Another explanation put forward was altruism, that is, the idea that players cooperate to improve the others' payoff. While a small fraction of altruists is consistently identified in games, their number is not sufficient to explain the level of cooperation usually observed in repeated public goods games (Andreoni & Miller, 1993). All these theories were developed to explain the cooperation of anonymous individuals, often in single interaction. Other theories explain cooperation that is embedded in larger social environments. Among these theories are those of reputation (Gunnthorsdottir, Houser, McCabe, & Ameden, 2001; Wilson & Sell, 1997), signaling (Kevin A. McCabe, Rassenti, & Smith, 1996), inclusive fitness (Hamilton, 1964), stochastic collusion (Flache & Macy, 2002; Macy, 1991), and indirect reciprocity (e.g. Yamagishi & Kiyonari, 2000). The next sections introduce prominent accounts of cooperation, most of which can be classified as explanations based on decision rules, on learning, and on social preferences.

1.2.1 Decision Rules for Cooperation

The decision rule approach assumes that players use simple rules or heuristics to decide about their choices in social dilemmas. The important difference to the game theory is that players are not assumed to thoroughly examine the incentive structure of a game and then infer which action maximizes utility. Rather, it is assumed that participants come into a situation equipped with a set of rules and then choose one of theses rules to make decisions.

The most popular rule for cooperation in Prisoners' Dilemmas is the Tit-For-Tat rule (Axelrod, 1984), which starts by cooperating, and then always imitates the choice of the other player in the previous round. While Tit-For-Tat—TFT, a specific implementation of the more general principle of reciprocity—was identified early as a good model for players behavior in public goods games (Rapoport & Chammah, 1965; Rapoport & Dale, 1966), its prominence is due to its success in a tournament among strategies for the repeated Prisoners' Dilemma game, where it outperformed all other strategies (Axelrod, 1984). Following up on Axelrod's insights, Samuel Komorita and colleagues (for a review see Komorita & Parks, 1999) examined which properties of TFT promote cooperation. Their general result was that TFT increases cooperation, especially if it rewards the others' cooperation, and punishes their defections immediately. Komorita (1965) found that not all players reciprocate, and that reciprocity is stronger if it leads to higher payoffs. While TFT was successful in many simulations, it has a weakness that led others to propose alternative strategies. Two "weaknesses" of TFT are that it does not forgive defection, which can lead to vicious cycles of mutual defection if two TFT strategies interact, and that it cannot exploit unconditional cooperators. One alternative strategy without these two weaknesses is PAVLOV (or Win-Stay Lose-Shift), which cooperates after mutual cooperation and after it could exploit, and defects otherwise. The inventors of this strategy showed that PAVLOV is superior to TFT (Nowak & Sigmund, 1993; but see Wu & Axelrod, 1995), and Macy (Macy, 1995) showed that players' behavior in a Prisoners' Dilemma game is best described by either TFT or PAVLOV.

These results on TFT were obtained with Prisoners' Dilemma games, but will they also hold for public goods games? Accordingly, Komorita, Parks, and Hulbert (1992) found that (simulated) reciprocal strategies induce participants' cooperation in social dilemmas experiments. However, they also found that the effect of reciprocity decreased in larger groups, and that without simulated reciprocal strategies in a group, reciprocal behavior did not emerge among players. In contrast, Sudgen (1984) reports that general contribution patterns to public goods are in line with his rule of reciprocity. In sum, while there seems to be clear evidence for reciprocal decision rules in Prisoners' Dilemmas, the evidence from public goods games is less conclusive.

Most research on decision rules assumes that people are already equipped with decision rules and, therefore, tries to identify which decision rule the players use. This leaves open, however, the question of why people are equipped with certain rules and not with others, and how they choose among the rules. One approach addressing the first question is research on fast and frugal heuristics (Gigerenzer et al., 1999), which assumes that decision rules are adapted to the decision environment and exploits individuals' cognitive capabilities. Applying this logic to decision rules for social dilemmas, a decision rule for cooperation in public goods games should have realistic assumptions about players' cognitive abilities in games, it should be adapted to other players' behavior, and a decision rule should lead to good outcomes (i.e., high payoffs) in public goods. As reciprocity has only low demands on players' cognition and has proved to be successful in Prisoners' Dilemmas and, importantly, also in public goods, I will suggest and test a reciprocity heuristic for cooperation in groups, and examine its success in different environments.

Decision rules for social interaction are usually domain specific. That is, dependent on the payoff structure of the situation (Kelley & Thibaut, 1978), or the organization of the interaction (Fiske, 1992; Kollock, 1994; Molm, Peterson, & Takahashi, 1999), different decision rules might be applied. While this remains a problem for heuristic approaches to decision making and seems unsolved for social dilemmas (but see Messick, 1999; J. M. Weber, Kopelman, & Messick, 2004), domain-general models do not have the problem of strategy selection. The next section introduces learning approaches to behavior in games, which—in stark contrast to the game theory—make minimal assumptions about players' individual rationality, and can predict the same behavior in the long run.

1.2.2 Learning to Cooperate

Reinforcement learning models are built on the law-of-effect, which suggests that individuals repeat behaviors that led to good outcomes (compared to alternative outcomes). Reinforcement learning is attractive for game theorists (Fudenberg & Levine, 1998) because it is consistent with a common finding in many experimental games that participants do not play Nash strategies from the outset of a game, but gradually converge to it. This finding is consistent with reinforcement learning because players who experiment with different strategies will receive higher payoffs when (by chance) choosing a Nash strategy, for which the law-of-effect then predicts a higher likelihood in the future. Beyond this general insight, formal models of reinforcement learning were tested experimentally and proved to be good models for decision making in experimental games (e.g. Camerer & Ho, 1999b; Erev & Roth, 1998).

While reinforcement learning models are based on the law-of-effect, they differ in the extent to which they incorporate additional cognitive processes, such as, for instance, belief updating and best response play. In the reinforcement learning model of Erev and Roth (1998), only strategies actually chosen by a player can be reinforced. In contrast, the belief learning model of Camerer and Ho (1999b) assumes that players also imagine which payoff they would have made had they chosen other strategies, and that this imagined payoff can also be used to reinforce strategies. A further variant of reinforcement learning models assumes that nonchosen alternatives are reinforced in a directional manner. According to directional learning models (Rieskamp, Busemeyer, & Laine, 2003; Selten & Buchta, 1999), players evaluate if their change in behavior from the second-to-last to the last period led to a higher or lower payoff. In the case of a higher payoff, players are assumed to change their behavior in the same direction again (and vice versa).

Can reinforcement learning models predict behavior in social dilemmas? Similar to research on decision rules, most research was conducted with Prisoners' Dilemma games. An early test was made by Rapoport and Dale (1966), who found some evidence in favor of reinforcement learning. More recently, Goren and Bornstein (1999) reported that it is mainly learning that changes behavior in the repeated intergroup Prisoners' Dilemma, Selten and Stöcker (1986) found that the directional learning theory can explain how players change their behavior across repeated supergames, and Erev and Roth (2001) found that behavior in Prisoners' Dilemmas is consistent with the assumption that players learn among the three strategies, defection, cooperation, and Tit-For-Tat. Learning was also examined in public goods games. Andreoni (1988) found that learning can explain some, but not all, of the decreasing contributions in public goods games, and Gunnthorsdottir and Rapoport (in press) found that reinforcement learning explains contributions in public goods on the group level, but not on the individual level.

Summarizing the results on learning in social dilemmas, reinforcement learning seems to be a good model to describe behavior in Prisoners' Dilemmas, although there are only a few tests on reinforcement learning in public goods. The same holds for the directional learning theory, which has not been tested in public goods games so far. Therefore, I will test how simple reinforcement learning and directional learning can account for cooperation in groups.

1.2.3 Social Values and Preferences

The decision rule approach and the learning approach retain the assumption that players want to achieve the best possible outcome for themselves, but suggest that they apply simpler cognitive processes to achieve their goals. In contrast, theories of social value orientation and social preferences assume that players do not maximize payoffs as given in the payoff matrix, but derive their utilities by weighing their own and others' payoff.

More specifically, theories of social value orientation (e.g. van Lange, 1999) assume that players derive an effective matrix (Kelley & Thibaut, 1978) from the given payoff matrix by adding or subtracting their own and others' payoff in every cell of the payoff matrix, and choose the strategy that maximizes their own payoff in the effective matrix. For instance, cooperative players are those who add their own and the other's payoff in a cell to derive the effective matrix, and competitive players subtract the other's payoff from their own payoff. Theories of social value orientation make no explicit statement about how people proceed to make decisions, given the effective matrix, beyond the assumption that players will make choices that maximize the effective payoff. While measures of social value orientation tend to predict behavior in one-shot games well (Au & Kwong, 2004), they were tested only once in repeated public goods games (Parks, 1994). Parks' results showed a low convergence of different measures of social value orientation, and also mostly low predictive value.

While psychologists use the term social value orientation, economists speak of social preferences or utilities. The dominant theories of social preferences are Fehr and Schmidt's (1999) theory of fairness, competition, and cooperation which basically assumes that players (in varying degrees) try to minimize inequity between their own and the others' payoffs, Bolton and Ockenfels' (2000) ERC theory of equity, reciprocity, and competition which assumes that players are motivated by their own payoffs and their relative standing to others, and Rabin's (1993) theory of fairness and reciprocity which assumes that players reciprocate the other's intentions. While psychologists' theories of social values and economists' theories of social preferences agree in their fundamental assumption that values in the payoff matrix need to be transformed to reflect players' utilities, they differ in their statements about players' rationality and in the way theories are tested. First, and most importantly, economists maintain the game theoretic framework, that is, they assume that the assumptions of individual rationality hold (given social utilities), and players thus play equilibrium strategies (given social utilities). Secondly, they use no independent measures of social preferences, rather, researchers examine if a set of parameters for the social utility function can explain behavior in a large variety of games, including public goods games.

While the theories of Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) are rather similar to psychological theories of social value orientation, Rabin's (1993) approach of reciprocity toward intentions has counterpart in psychological research. Hence, I will first examine how players' social value orientation can predict cooperation in public goods, and later examine the role of others' intentions for reciprocity.

1.3 Learning to Decide

Having identified decision rules, further questions are immediately posed. How did people acquire these decision rules? Also, why, or how, do they choose among decision rules? This section introduces approaches aiming to explain how people learn to make decisions. Specifically, I will briefly sketch the evolutionary approach which assumes that decision mechanisms evolve as solutions to adaptive problems, the individual learning approach which assumes that individuals discover decision rules through their experience with the decision problem, and the social learning approach which assumes that individuals learn from others.

The take of evolutionary psychology on cognition, in general, and also on decision making, in particular, is that organisms have a set of mechanisms (or modules) that evolved as a response to adaptive problems. Adaptive problems are problems directly or indirectly linked to the reproductive success of an organism. Decision-making mechanisms evolve through the interplay of random mutations and selection processes (Cosmides & Tooby, 2000; Cosmides et al., 1992). While evolutionary approaches to decision making are valuable in providing a framework that investigates the ultimate reasons for behavior, it seems difficult establish that a given human behavior (including cooperative behavior) can primarily be considered as a result of evolutionary adaptation (e.g. Heyes, 2003). Therefore, this dissertation focuses on processes of individual and social learning. Appendix C introduces evolutionary approaches to the explanation of cooperation.

1.3.1 Individual Learning

The idea of learning introduced here differs conceptually from the kind of learning described in Section 1.2, in that it is assumed that people do not learn among different possible actions (e.g., cooperate or defect in the Prisoners' Dilemma), but among different decision rules that can be applied repeatedly (e.g., Tit-For-Tat in the Prisoners' Dilemma). Models of individual learning describe and explain how people invent decision rules, and how they decide which rule to choose in a specific situation. Individual learning is different from social learning, as it assumes that a person discovers a decision rule, and also chooses independently which decision rule to apply.

Siegler and colleagues (1998; 2005) proposed a model of strategy choice and strategy discovery (SCADS) describing how children increase their performance in a simple addition task, by first improving their performance with the basic strategy, and later discovering new strategies that further improve performance. More specifically, they assume that a strategy consists of several operations, which are executed more efficiently through experience. As in associationist models, the propensity to choose a strategy grows with the success of this strategy. More efficient execution frees cognitive resources that can be invested in the discovery of new strategies, which are generated by rearranging the operations of the basic strategy. Siegler and colleagues assume that children internally generate many rearranged strategies, but because invalid strategies are generally not observed, they also assume that newly generated decision rules have to pass a filter, checking their validity before they are actually applied. Further, the switching from one strategy to the other is assumed to be gradual because the high propensity for the basic strategy (which was successful in the past) maintains high choice probabilities of the basic strategy, even in the presence of a better new decision rule. While the SCADS model could accurately describe the process of children's learning of addition rules, it also has drawbacks. One key drawback is that it was developed post hoc, in order to produce the results of a specific experiment, and has not been retested in its original form. The other drawback is that it is not clear if the key assumption of the models (an associationist process underlying the selection of decision rules, and a metacognitive process underlying the discovery of new rules) mainly contribute to its performance, or if additional assumptions (i.e., assumptions about basic operations) that need to be made to implement a specific task, such as addition, contribute more to the performance. Indeed, when applying the model to a new task, Siegler and Araya (2005) had to modify the original model in many respects. Given these difficulties and the lack of alternative (formal) models of strategy discovery, I will focus on simpler models that can describe how people learn *among* strategies.

Models of strategy learning assume that a reinforcement learning process can describe how people learn and choose among alternative decision rules (Erev & Barron, 2005; Hanaki, Sethi, Erev, & Peterhansl, 2005; Rieskamp & Otto, 2004; Stahl, 1996, 2000). The general idea of all these models is that decision makers have a set of strategies, at first choose randomly among them, and then prefer strategies leading to better outcomes. More specifically, Rieskamp and Otto (2004) tested how individuals learn among two decision rules for inferential decision making. For this aim, they examined if participants' repeated choices and their information search coincided with a decision rule, and if the decision rule made a correct decision. If all three conditions were met, the expectancy of the rule was increased according to the payoff for a correct decision. If participants searched for information, and decided according to the information rule, but made a wrong decision, the expectancy for that decision rule decreased. Rieskamp and Otto found that participants learned which rule made more correct responses in a task, and that this learning process was described well by their *strategy selection theory* (SSL). Hanaki et al. (2005) used a reinforcement learning model to examine how individuals learn among strategies for different repeated two-person games—Stag Hunt, Prisoners' Dilemma, Chicken, Battle of Sexes—empirically examined by (McKelvey & Palfrey, 2001). For this aim, they conducted a simulation Experiment that tried to predict the average choices of action in the four games. To generate the strategies among which the simulated players learn, they constructed all possible two-state finite automaton (26), which represent strategies, such as Tit-For-Tat or "always cooperate". Given these strategies, learning occurred in two steps. In a first simulation, automata were repeatedly and randomly rematched with other automata to play a repeated game. In this “preexperimental” phase, the initial propensities of the automata for the simulation of the experiment were determined according to the success of the automaton. Propensities of the automaton changed according to the law-of-effect. The second part of the simulation replicated the experiments run by McKelvey and Palfrey (2001), by randomly drawing automata from the population and allowing them to play the repeated games, in which propensities for strategies were again updated. The comparison of payoffs in the second part of the simulation with empirical payoffs from McKelvey and Palfrey indicated that the simulation successfully mirrored the learning process of real participants.

Beyond the work of Hanaki et al. (2005), experiments showed that reinforcement learning models can successfully model how people learn among repeated game strategies (Erev & Barron, 2005; Stahl, 1996, 2000), but all these models assume an autonomous decision making of players. In contrasts, players can often observe the others’ choices, or decision makers receive advice about what to choose. The next section therefore introduces models of social learning.

1.3.2 Social Learning

Social learning can take the form of imitation learning (e.g. Heyes, 2001), observational learning (Bandura, 1977), and advice-taking (e.g. Budescu, Rantilla, Yu, & Karelitz, 2003). In imitation learning, individuals learn by replicating the observed behavior of other people. Observational learning, in the sense of Bandura, goes beyond simple imitation in that learners infer the goal of the others’ behavior, and learn to achieve the same goal instead of imitating the same behavior. In advice-taking, individuals receive a hint on how they should behave or decide, and integrate this with their own information or preferences to make a decision.

Theories of advice-taking try to model the way in which decision makers integrate advice from different sources or integrate advice with their own information. For instance, Budescu et al. (2003) examined how people integrate advice from different advisors with different competences. Their results show that advisors who had more information or made more correct decisions in the past receive a higher weight. In a series of experiments, Schotter (2003) observed if players in one-shot games use the advice they received from another player who played the game just before them. They found high adherence to advice in social dilemmas, such as the Prisoners’ Dilemma and the chicken game, but also in coordination games, such as the Battle of Sexes. Interestingly, players also adhered to advice when it meant choosing a strategy leading to an inferior outcome for the player, compared to the other player. The common theme of these experiments is that decision makers in individual and interdependent decision making are influenced by advice from experienced decision makers.

While experiments on advice-taking preselect the advisor for participants, theories of imitation describe who will be imitated. A prominent theory of imitation suggested by Boyd and Richerson (1985; cited in Joseph Henrich & McElreath, 2003) predicts that decision makers should imitate more when a task is difficult and when the environment frequently changes. In an experiment designed to test this prediction, McElreath et al. (2004) could not find the predicted relation between task difficulty or variability of the environment and reliance on social learning, but they identified a model that best describes the social learning process. The best model assumed that players conform to majority choices that they observed in preceding rounds, and that learning among the choice options follows the law-of-effect. Social learning was modeled so that the reinforcement of options was contingent on the payoff that players received from choosing an option and from a number of other players who also chose the option. Apesteguia, Huck, and Oechssler (2003) tested the effect of information on imitation in interdependent decision making. In their experiment, participants in one condition could observe the decisions of other players with whom they played, whereas in the other condition they could observe another player with the same role as themselves, but playing in a different game. Their results show that players rather imitated others with whom they interact than others with the same role. Also, the likelihood to imitate another more successful player increased in the difference in payoffs.

In sum, theories of advice-taking and imitation correctly predict that decision makers are willing to use social information in individual or interdependent decision making. However, a common feature to all research on social learning is that decision makers receive advice, or can observe other players, before every single decision. In contrast, people often receive advice only once, and then have to make decision on their own. The last chapter of my dissertation proposes and tests models of social learning when people first receive advice, and then repeatedly make choices.

1.4 Organization of the Dissertation

The main aim of my dissertation is to explain how cooperation can be maintained without formal institutions. To set the stage for this examination, Chapter 1 has introduced the public goods games and shows that cooperation in social dilemmas is individually rational only in indefinitely repeated interactions. Nevertheless, participants in experiments commonly contribute to one-shot public goods or in finitely repeated public goods games. The second part of the introduction discussed models of cooperation in social dilemmas, the most important of which are decision rules for cooperation, learning, and social preferences. While research on these three approaches has been extensive in the past, the different approaches have not been compared directly. Therefore, Chapter 2 compares how well these approaches, especially a reciprocal decision rule, reinforcement learning, and directional learning can explain cooperation in groups. Beyond this, the first chapter also examines the influence of interaction structure on cooperation by proposing a new game, the Social Dilemma Network, which aims to exploit people’s reciprocal tendencies.

Chapter 3 takes a closer look at reciprocal decision rules by integrating two traditions of research on reciprocity: psychological research on reciprocity which implicitly models reciprocity as reciprocating others’ observed behavior, and economic theories of reciprocity which assumes that people reciprocate the others’ intentions. On this basis, I first examine how adaptive it is to consider the others’ intentions in public goods games, and then suggest and test two simple reciprocity rules, reciprocating either the others’ behavior or intentions.

While Chapters 2 and 3 examine the nature of people’s decision rules for cooperation, Chapter 4 aims to answer the question of how we learn to make good decisions. Specifically, this chapter looks at the effects of social learning in repeated decision making, and proposes models describing the learning process. Going beyond existing research and models on social learning, I will examine and model a situation in which decision makers receive a single piece of advice, and then have to make decisions repeatedly.

Finally, Chapter 5 summarizes the results of the empirical chapters, and critically evaluates their contribution to research on social dilemmas, decision rules, and learning, in general, and to the question of cooperation in groups, in particular.