CHAPTER 2

The Golden Rule in Groups I:

Using Reciprocal Tendencies to Improve Cooperative Outcomes

2.1 Introduction

Cooperative groups realize projects a person cannot realize alone. Production teams assemble cars, a group of shop owners join forces to advertise shopping in their street, or neighbors maintain a tidy neighborhood. However, the benefits of groups' goods are often non-excludable, that is, every group member can benefit from the established goods regardless of his or her contribution, and non-rival, that is, consuming the goods does not diminish the goods' value for other group members. Non-excludable and non-rival goods are considered public goods (Ostrom, Gardner, & Walker, 1994). Public goods represent a social dilemma situation (Dawes, 1980; Komorita, Chan, & Parks, 1993): If no one contributes, everyone is worse off than if all had contributed. However, every group member benefits most if he or she does not contribute to the production of the good. Consequently it is difficult to maintain cooperation in groups providing public goods.

Experimental results show that in iterated public goods games cooperation usually starts at an intermediate level and subsequently declines (Ledyard, 1995). Although different factors have been examined that influence cooperation in public goods games (for reviews see Dawes, 1980; Kollock, 1998; Komorita & Parks, 1995; Ledyard, 1995; Weber, Kopelman, & Messick, 2004; Zelmer, 2003), I argue that the underlying cognitive processes of cooperative behavior are rather unexplored. To remedy this, I tested three approaches to describing cognitive processes that determine decisions in social dilemmas.

Game theory predicts that self-interested payoff maximizers should not cooperate in finitely repeated social dilemmas. Since people's behavior often deviates from the game-theoretical prediction, different approaches were proposed to explain cooperative behavior. The first approach, social motivation, explains cooperative behavior with other-regarding, that is, social preferences. The theory of social value orientation (McClintock, 1978; van Lange, 1999) assumes that people are not purely motivated by narrow self interest, but by more complex considerations about their own and others' payoff. According to this theory, people's social value orientations are classified into one of three categories: prosocial, individualistic, or competitive. Individuals with a prosocial orientation prefer outcomes that realize a maximum joint benefit; individuals with an individualistic orientation prefer outcomes with a maximum individual benefit; and individuals with a competitive orientation prefer outcomes with higher payoff for themselves combined with maximum difference between their own and others' payoffs. Among others, van Lange (1999) showed that social value orientation can predict cooperation in social dilemmas. Recently, experimental economists have utilized similar approaches to incorporate non-selfish preferences in standard expected utility models. Prominent examples are theories by Bolton and Ockenfels (2000) and Fehr and Schmidt (1999). According to these theories, the utility of an outcome is not purely determined by self interest; instead, additional utility results from social preferences, such as, for instance, equality. Consequently people cooperate because cooperation simultaneously leads to high payoffs and equality.

The second approach, reciprocity, explains cooperation in groups by assuming that people use heuristics for social interactions (Gigerenzer et al., 1999; Messick, 1999; J. M. Weber et al., 2004). Accordingly, people are equipped with a repertoire of strategies they apply for decision situations they encounter, including decisions in social interactions (Gigerenzer, 2001; Todd, Rieskamp, & Gigerenzer, 2002). Depending on the decision situation, people select different strategies. A prominent rule for interacting with other people is the golden rule, or reciprocity, which prescribes to do onto others as they did onto you. In repeated interactions, cooperative strategies such as "Tit-For-Tat" are more likely to be selected (Fehr & Henrich, 2003), since they can outperform non-cooperative strategies (Axelrod, 1984; but see also Chapter 3 in Binmore, 1998). Likewise, Weber et al. (2004) suggested that interactions in social dilemma situations trigger reciprocal behavior, and Fiske (1992) stated that "equality matching," describing relationships in which people ensure positive reciprocity, "is a common blueprint for connecting people." Komorita (1965) has shown that reciprocity, implemented as a Tit-For-Tat strategy, is in fact a good model to describe people's decisions in iterated prisoner's dilemmas.

Social motives and reciprocity are not mutually exclusive explanations of cooperative behavior. One way to connect reciprocity with social motives is to assume that people with a preference for maximum joint outcome can be influenced by others' behavior, so that they may not cooperate when others defect. For instance, in Rabin's (1993) model of social preferences people preferred "cooperative outcomes" only when others reciprocated cooperation. A second way to connect reciprocity with social motivations is to assume that people with selfish preferences reciprocate, because a reciprocal strategy maintains cooperation that in the long run maximizes individual payoffs (van Lange, 1999). While investigating motivations for cooperation is important, I focus specifically on the decision processes involved in cooperative behavior. Decision processes serve the role of realizing the aims of cooperative motives and are a relatively unexplored realm of research.

The third approach I consider explains cooperation in groups from a learning perspective. The learning approach does not make assumptions about people's social motivations. Instead it suggests that people's behavior is mainly a function of past experience, such that behavior becomes more frequent when it has led to positive consequences; thus behavior is a function of its reinforcement. Explaining cooperation in social dilemma situations by reinforcement learning mechanisms has a long tradition in psychology (e.g. Rapoport & Chammah, 1965). Recently developed reinforcement learning models are able to explain behavior in a variety of experimental games (Erev & Roth, 1998) and have been used to model dynamics in iterated prisoners' dilemma games (Flache & Macy, 2002). In contrast to simple reinforcement learning models, directional learning models predict that people keep track of the direction of their behavioral changes, so that after a successful change, behavior that results from a change in the same direction becomes more likely (Rieskamp et al., 2003; Selten & Stöcker, 1986). For instance, if a person decreased the contribution to a public good and received a larger payoff compared to the previous payoff, it is most likely, according to directional learning, that she will again decrease her contribution in the next round. In contrast, simple reinforcement learning predicts that it is most likely the she will repeat the same contribution that led to the highest payoff so far. The learning approach and the reciprocity approach make different predictions in regard to an individual's reaction to the high payoffs resulting from free riding on another's cooperation. Whereas the reciprocity heuristic predicts cooperation as a reaction, reinforcement learning predicts repetition of defection and local adaptation learning predicts decreasing one's contribution to the public good.

My main goal was to determine which of the three approaches—social motivation, learning, or reciprocity—is most suitable for predicting behavior in social dilemmas. All three approaches have been successfully applied in the past, making the comparison worthwhile and necessary. I will first propose a formal model of reciprocity to represent the reciprocity approach. Second, to represent the learning approach two learning models will be specified. Third, to represent the social motivation approach, I will assess individuals' social value orientation. I will then report on the comparison of the approaches in two *n*-person dilemma games, which differed in the interaction opportunities they gave to group members.

2.2 Reciprocity in Groups

Although many researchers agree on the basic understanding of reciprocity—that individuals do unto others as others have done unto them—reciprocity has been defined in various ways, such as a behavioral strategy (Axelrod & Hamilton, 1981; Komorita & Parks, 1999; Trivers, 1971), an external norm (Gouldner, 1960), an internal norm (Gallucci & Perugini, 2003), a social preference (Bolton, Brandts, & Ockenfels, 1998; Rabin, 1993), or an evolved ability (Trivers, 1971). In the iterated prisoner's dilemma game, which is the classic paradigm for Experimenting cooperation in dyads, reciprocity has been implemented as the Tit-For-Tat strategy (Axelrod & Hamilton, 1981; Rapoport & Chammah, 1965), in which a player imitates the other player's behavior in the previous interaction. It was shown that people deliberately reciprocate in the iterated prisoner's dilemma and cooperate more when playing against programmed reciprocal strategies (see e.g. Bixenstine & Gaebelein, 1971; Sermat, 1967). Reciprocal behavior was also observed in other two-person interactions such as the trust game (Kevin A. McCabe, Rigdon, & Smith, 2003; Pillutla, Malhotra, & Murnighan, 2003), or the gift exchange game (Fehr, Kirchsteiger, & Riedl, 1998).

Since reciprocity strongly influences behavior in two-person interactions, it seems natural to conjecture that people also reciprocate in groups. Early investigations of reciprocity in groups found that programmed reciprocal strategies successfully maintain cooperation in groups (e.g. Komorita, Chan, & Parks, 1993); absent these programmed strategies, however, little reciprocity in groups could be found (e.g. Bornstein, Erev, & Goren, 1994). In contrast, recent experiments provide evidence of reciprocal behavior in groups (e.g. Ehrhart & Keser, 1999; Fischbacher, Gächter, & Fehr, 2001; Kurzban, McCabe, Smith, & Wilson, 2001), although the proportion of reciprocating participants in public goods games varies. Kurzban and Houser (2001) reported that 28% of the participants in their experiment were classified as reciprocators, compared to an estimated proportion of conditional cooperators of 50% in Fischbacher et al. (2001), and 76% classified as reciprocators in Houser and Kurzban (2005).

Empirical evidence suggests that people reciprocate in public goods games, but reciprocating cooperative behavior might not be sufficient to maintain a high level of cooperation. There are two arguments why cooperation is difficult to maintain in groups. First, suppose a group consists of ten members, one single selfish individual who never cooperates, and nine reciprocators who always contribute as much as the others' preceding average contribution. Since the others' average contribution will always be lower compared with the reciprocators' contributions, reciprocators will repeatedly decrease their contributions, so that after some time no one will contribute anymore (see also Fehr & Fischbacher, 2003). Following this argument the larger the group size, the higher the chance of including a selfish individual who deters cooperation.

A second argument for why cooperation is difficult to maintain in public goods games is the different effects of defection in groups and dyads: In the iterated prisoner's dilemma defection has a severe consequence for the exploited player, whose payoff always falls below the level he could guarantee to himself through defection. Hence a player can anticipate that defection in the iterated prisoner's dilemma will most likely trigger the other player to defect, too. In contrast, if in a public goods game a few players make contributions and the others do not contribute, the cooperating players can still receive a larger payoff compared to a situation where all players make no contributions, the acceptable number of defectors being dependent on the group size and on the efficiency gain through cooperation. Since the probability of getting along with defecting seems higher in a public goods game, defecting is more tempting in a public goods game than in the iterated prisoner's dilemma. In sum, it can be predicted that even in the presence of many cooperators cooperation in groups will decline over time and that cooperation is negatively correlated with the group size.

These predictions have gained empirical support: Apart from the general finding of declining cooperation in public goods games (Ledyard, 1995), Dawes (1980) and Komorita, Parks, and Hulbert (1992) demonstrated the negative effect of increasing group sizes on cooperation. Marwell and Schmitt (1972) found that cooperation declined more over time in a three-person-dilemma compared to the iterated prisoner's dilemma.

To address the low cooperation rates typically found in public goods games, researchers have studied contribution that could increase cooperation. For example, Erev and Rapoport (1990) allowed for sequential contributions, Bornstein and Ben-Yossef (1994) embedded the public goods game in an intergroup conflict, Fehr and Gächter (2000; Fehr & Gächter, 2002) allowed for costly punishment, and Coricelli, Fehr, and Fellner (2004) provided the opportunity to select with whom to produce a public good. I argue that cooperation is easier maintained in iterated prisoner's dilemmas compared to iterated public goods games because in the former game individuals can directly reciprocate, whereas in the latter game individuals must react in the same fashion to several others who might have behaved differently. The ability to treat group members differently is an important factor influencing cooperation in groups. Consequently it can be predicted that cooperation will increase when a public good is divided among its members and dyadic interdependencies are enhanced. To test this prediction I will examine behavior in a standard public goods game and compare it to behavior in a modified public goods game that I call the social dilemma network game (SDN). In the SDN the public project is split into multiple public projects such that every individual can cooperate simultaneously in several two-person public projects with every other member of the group (see Figure 2.1). Similar network-structures of groups have been examined by Flament and Apfelbaum (1966) and Feger and von Hecker (1998), who reported that reciprocity guided participants' behavior in such interactions.

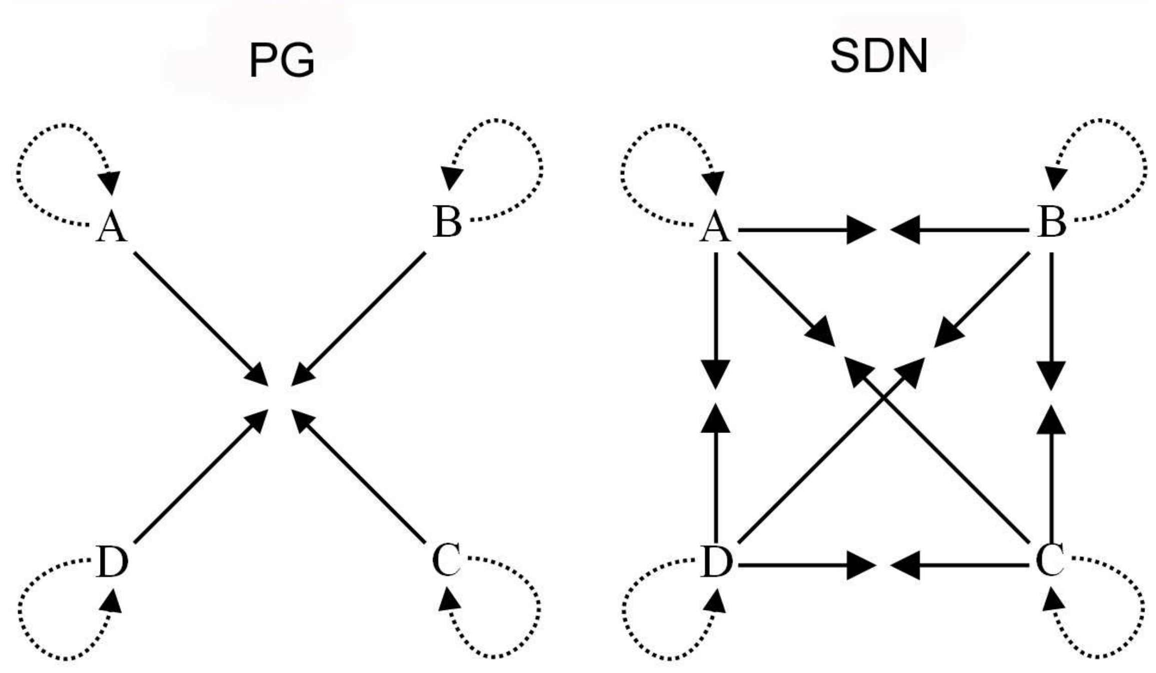

Figure 2.1. Interaction structure of the public goods game (PG) and the social dilemma network (SDN). Solid lines symbolize possible contributions to public projects, and dotted lines symbolize possible contributions to private projects. A player in the SDN can freely allocate his endowment to the four projects he has access to. In the experiment, players were endowed with 10 points in every round, which could be allocated to any allocation option available to a player.

I define the public goods game used in Experiment 2.1 as follows: Each of $N = 4$ players in the public goods game has an endowment, E, which she can allocate to a private project or to a public project (I use the term *public goods* to refer to the game and the term *project* to refer to allocation options in a game), where the contribution to the public project is c and the investment to the private project is $E - c$. Investments to the private project lead to a payoff equal to the investment. Contributions c to the public project are multiplied by a constant and then equally split among all members of a group. The efficiency gain of a contribution to the public project is expressed as marginal per capita return (MPCR), which is the quotient of the multiplication constant and the number of players. A player's payoff, π_i, in the public goods game is defined as $\pi_i = E_i - c_i + MPCR \cdot \sum_{i=1}^{N} c_i$.

The SDN is defined as follows: Every player has an endowment of E that can be allocated to a private project and three two-person public projects, one with each other player (compare Figure 2.1). As in the public goods game, an investment into the private project leads to a payoff equal to the investment, whereas contributions to each public project are multiplied by a constant and then equally split among the two members of the public project. A player's payoff, π_i, in the SDN is defined as $\pi_i = E_i - \sum_{m=1}^{N-1} \mathbf{c}_i + MPCR \cdot \sum_{m=1}^{N-1} C$, where m is an index for the N-1 public projects a player can contribute to, and C is the contributions of the player and her respective partner to the public project. Players profit from any contribution to one of their public projects, regardless of their own contribution.

In the public goods game a player only decides about the contribution to one public project, so that it is not possible to direct contributions to specific other players. In contrast, a player in the SDN separately decides about the contributions to each of her three public projects. This allows the player to make contributions conditional on the other player's contribution to the shared public project. This should foster the effectiveness of reciprocity, leading to the prediction that cooperation should be greater in the SDN compared with the public goods game. This *selection mechanism prediction* is tested in the experiment. Likewise I predicted that the reciprocity approach is better at predicting people's decisions in the SDN compared to the public goods game.

2.2.1 A Reciprocity Heuristic for Cooperation in Groups

How can the cognitive process that leads to reciprocal behavior in groups be specified? Formal models that define the reciprocal decision mechanisms for behavior in groups are, to my knowledge, very rare (but see Parks & Komorita, 1997; Rapoport & Chammah, 1965). To overcome this deficit, I will formulate a model of a reciprocity heuristic (REH) for repeated interactions. In general, the proposed reciprocity heuristic describes the following decision process for public good situations. In the first round of a game the endowment is equally allocated among the available projects. In the following rounds players contribute as much to the public project(s) as the others did in the preceding round. For the SDN the reciprocity heuristic predicts that players will contribute as much to a public project as the respective other did in the preceding round. If the endowment is not sufficient to reciprocate others' contributions, the endowment is allocated proportional to others' contributions. This general reciprocity principle is assumed to be mediated by individuals' generosity, which is the tendency to give more than a strict reciprocity principle predicts. For instance, if others in a public goods game contribute 50% of their endowment to the public goods game, then a generous individual might contribute 60% of her endowment (e.g. to encourage cooperation). Figures 3 and 4 display flow charts illustrating the process of the reciprocity heuristic for the SDN and the public goods game.

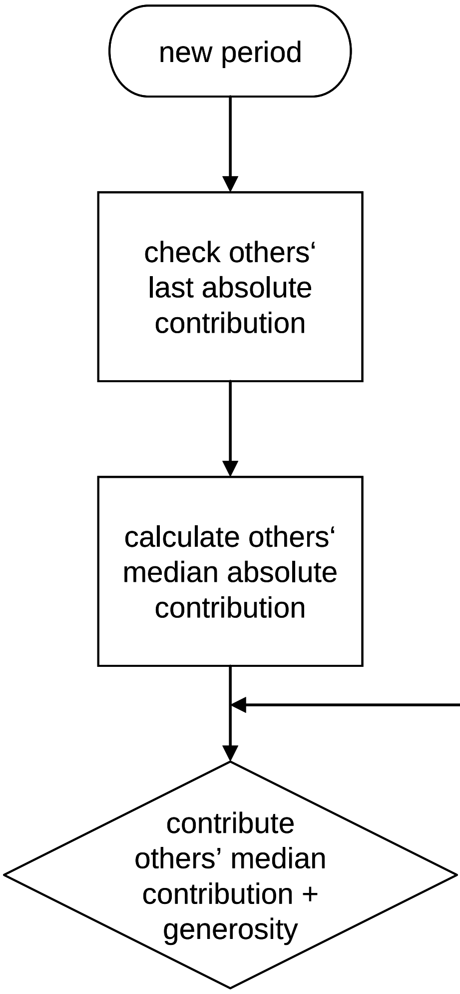

Figure 2.2. Process of the reciprocity heuristic in the public goods game. A rounded box signals the beginning of a process. Boxes stand for cognitive processes of gathering and manipulating information. Squares with round corners are alternative processes that also influence the outcome of a decision. Diamonds are decision nodes in the process, which determine how the process continues or have the final decision as an output. The reciprocity heuristic assumes that individuals use the median of others' contribution to represent their aggregate contribution. Individuals generosity or greediness determines how much they over- or under-contribute compared to exact reciprocation.

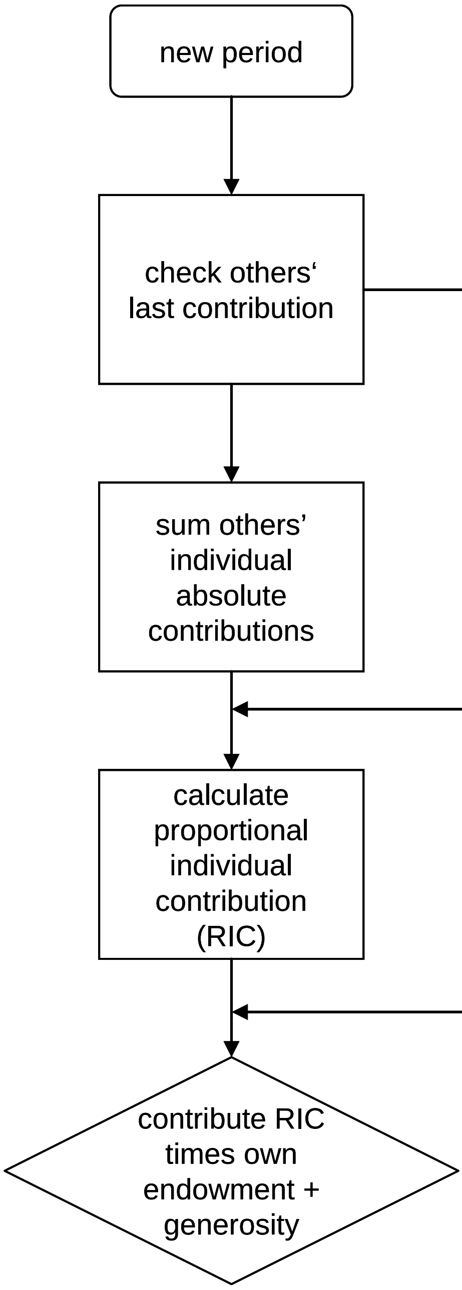

Figure 2.3. Process of the reciprocity heuristic in social dilemma network. The reciprocity heuristic assumes that individuals contribute to public goods with other individuals according to others proportional contributions. The proportional contribution is the contribution of the respective individual divided by the some of all other individuals' contributions.

When constructing models for predicting the behavior in the two games it is helpful to conceptualize the decision situation as an allocation problem, in which the decision of making contributions consists of allocating an endowment to M allocation options (projects). Thereby the decision of a player can be represented by a multi-dimensional allocation vector $\mathbf{c}^m$, where the dimension m denotes the different allocation options. In the case of the public goods game two dimensions of the allocation vector result (M = 2, m = 1 as private project, m = 2 as a public project), and in the case of the SDN four dimensions result (M = 4, m = 1 as private and m = 2, m = 3, m = 4 as three public projects). When representing the decision problem in a multi-dimensional space one can characterize the similarity of two allocations by their Euclidian distance. The following reciprocity models will first determine a *most likely allocation* and then determine the probability of other allocations.

Mathematically, REH is defined as follows: The most likely allocation l in the first round to the allocation option m, is an equal split of the available endowment E among the M options, so that $\mathbf{c}_l^m = E / M$, with M as the overall number of projects a player can invest in, m = 1 as the private project and m > 1 as the public project(s). The probability p_j that one of the possible allocations j is chosen is defined as:

$$p_j = \exp\left(-x_{jl} / 2\sigma_C^2\right)/U\ , \tag{2.1}$$

where x_{jl} is the Euclidean distance between any possible allocations j and the most likely allocation l, the free parameter σC is a standard deviation, defining to what extent similar allocations to the most likely allocation are also chosen with a substantial probability, and U is a constant that normalizes the probabilities so that they sum to 1. According to Equation 2.1, the probability of choosing an allocation increases with its similarity to the most likely allocation l.

In the second and all following rounds contributions depend on the other players' behavior in the preceding round. For the public goods game, the most likely allocation is determined by contributing as much as the other players did in the preceding round,

$$c_l^2 = median\{o_1, o_2, \ldots o_{n-1}\} + \gamma \cdot (E - median\{o_1, o_2, \ldots o_{n-1}\}) \tag{2.2}$$

with o_i as the other players' contributions in the preceding round and with $\gamma \in [0,1]$ as a generosity parameter determining how much a player contributes above the others' median contribution in the previous round. The contribution to player i's private project is $\mathbf{c}_l^1 = E - \mathbf{c}_i^2$. In the case of the SDN, the most likely allocation is determined by contributing as much as the other players contributed to the respective public projects in the preceding round:

$$\mathbf{c}_l^m = \frac{o_m}{O} \cdot [O + \gamma \cdot \max\{(E - O), 0\}], \qquad \mathrm{m} > 1 \tag{2.3}$$

with o_m as the other players' contributions in the preceding round to project m and $O = \sum_{m=2}^{M} o_m$. Finally, the contribution to a player's private project is $\mathbf{c}_l^1 = E - \sum_{m=2}^{M} \mathbf{c}_l^m$. The choice probabilities for all possible allocations are determined according to Equation 2.1.

2.3 Learning and Cooperation in Groups

In the following, two learning models are specified, representing the learning approach introduces above.

2.3.1 Reinforcement Learning

Reinforcement learning models have a long tradition in describing behavior in social dilemma situations (see for example Rapoport & Mowshowitz, 1966) and recently have regained prominence in describing behavior in experimental games (Camerer & Ho, 1998; Erev & Roth, 1998; Macy, 1995; Stahl & Haruvy, 2002). Derivates of Erev and Roth's (1998) reinforcement learning model have been tested for a variety of games, including public goods games. For instance, Gunnthorsdottir and Rapoport (in press) found that Erev and Roth's reinforcement learning model predicted cooperation accurately on a group level. I now define a reinforcement learning model (RLM) that is a modified version of Erev and Roth's model.

When applied to the public goods game or the social dilemma network game the RLM models the following decision process: All possible allocation options are assigned subjective expectancies, which are assumed to be equal in the first round. After payoffs are known the expectancies of all possible allocations are updated according to the received payoff[6], so that the chosen allocation and allocations that are similar to the chosen one obtain larger reinforcements. In the second and all following rounds new allocations are selected proportional to the updated expectancies. In general the model predicts that the probability of choosing allocations increases for those that led to higher payoffs. Figure 2.4 shows a flow chart illustrating the general process of the reinforcement learning model.

[6] Alternatively, reinforcement could be a function of others' contributions. Because this would make the learning model very similar to the reciprocity heuristic and I instead wanted to compare payoff-based reinforcement learning with reciprocity, I maintained the original reinforcement function as defined by Erev and Roth (1998).

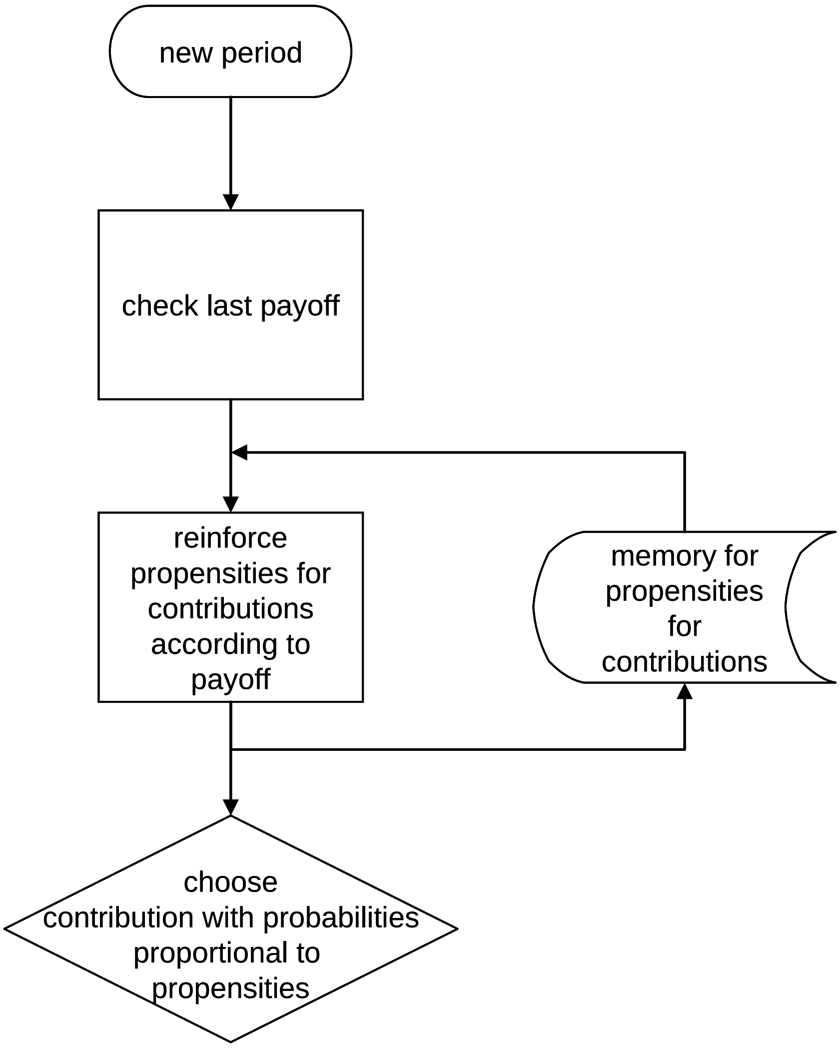

Figure 2.4. Schematic representation of the reinforcement learning model. Stylized paper sheets are symbols for memory, in which information from previous interactions is stored. When propensities are updated, the propensity for the last option is updated with the full payoff. Other propensities are updated according to their Euclidean distance to the last contributions, with more distant contributions receiving less reinforcement.

Mathematically, the RLM is defined as follows: The preferences for all possible allocations j are expressed as expectancies q_{tj} for each round t. The probability p_{tj} that an allocation j is chosen in round t is defined by (cf. Erev & Roth, 1998)

$$p_{tj} = \frac{q_{tj}}{\sum_{j=1}^{J} q_{tj}}. \qquad (2.4)$$

For the first round, all expectancies are assumed to be equal and determined by the average payoff that can be expected from random choice, multiplied by w, which is a free initial attraction parameter and is restricted to $w > 0$. After the choice of a particular allocation k in round t is made, the expectancies are updated by the corresponding reinforcement r_{tk}, which is a function of the received payoff π_{tk}. It is assumed that allocations similar to the chosen one are also reinforced. Therefore, to update the expectancies of any allocation j, the reinforcement r_{jk} is defined by the generalization function:

$$r_{tj} = \pi_{tk} \cdot \exp(-x_{jk}^2 / 2\sigma_R^2)\,, \tag{2.5}$$

where x_{jk} is the Euclidean distance of any allocation j to the just-chosen allocation k and with the standard deviation σ_R as the second free parameter. This function was selected so that the reinforcement r_{tj} for the chosen allocation is equal to π_{tk}. Finally, reinforcements are used to update the expectancies by the following updating rule (cf. Erev & Roth, 1998):

$$q_{tj} = (1-\phi)\ q_{t-1\,j} + r_{tj}\,, \tag{2.6}$$

where $\phi \in [0,1]$is the third free parameter, the forgetting rate, which determines how strongly previous expectancies affect new expectancies. If the forgetting rate is large, the just-obtained reinforcement has a relatively large influence on the expectancies for the following round, compared to reinforcements of earlier rounds. All allocations are chosen at least with a minimum probability, so that the minimum expectancy is restricted to v = 0.0001 (according to Erev & Roth, 1998). The choice probabilities are again determined according to Equation 2.4.

2.3.2 Local Adaptation Learning

The local adaptation learning model (Rieskamp et al., 2003) incorporates the idea of a directional learning process. The model, called LOCAD, was successfully applied in the context of an individual decision problem of allocating a resource among different financial assets. LOCAD assumes that people change their decisions in specific directions to improve payoff. LOCAD has similarities to the learning direction theory (Selten & Stöcker, 1986), which—besides iterated prisoner’s dilemmas—was applied successfully to experimental games like the repeated ultimatum game (Grosskopf, 2003), or the repeated beauty contest game (Nagel, 1995).

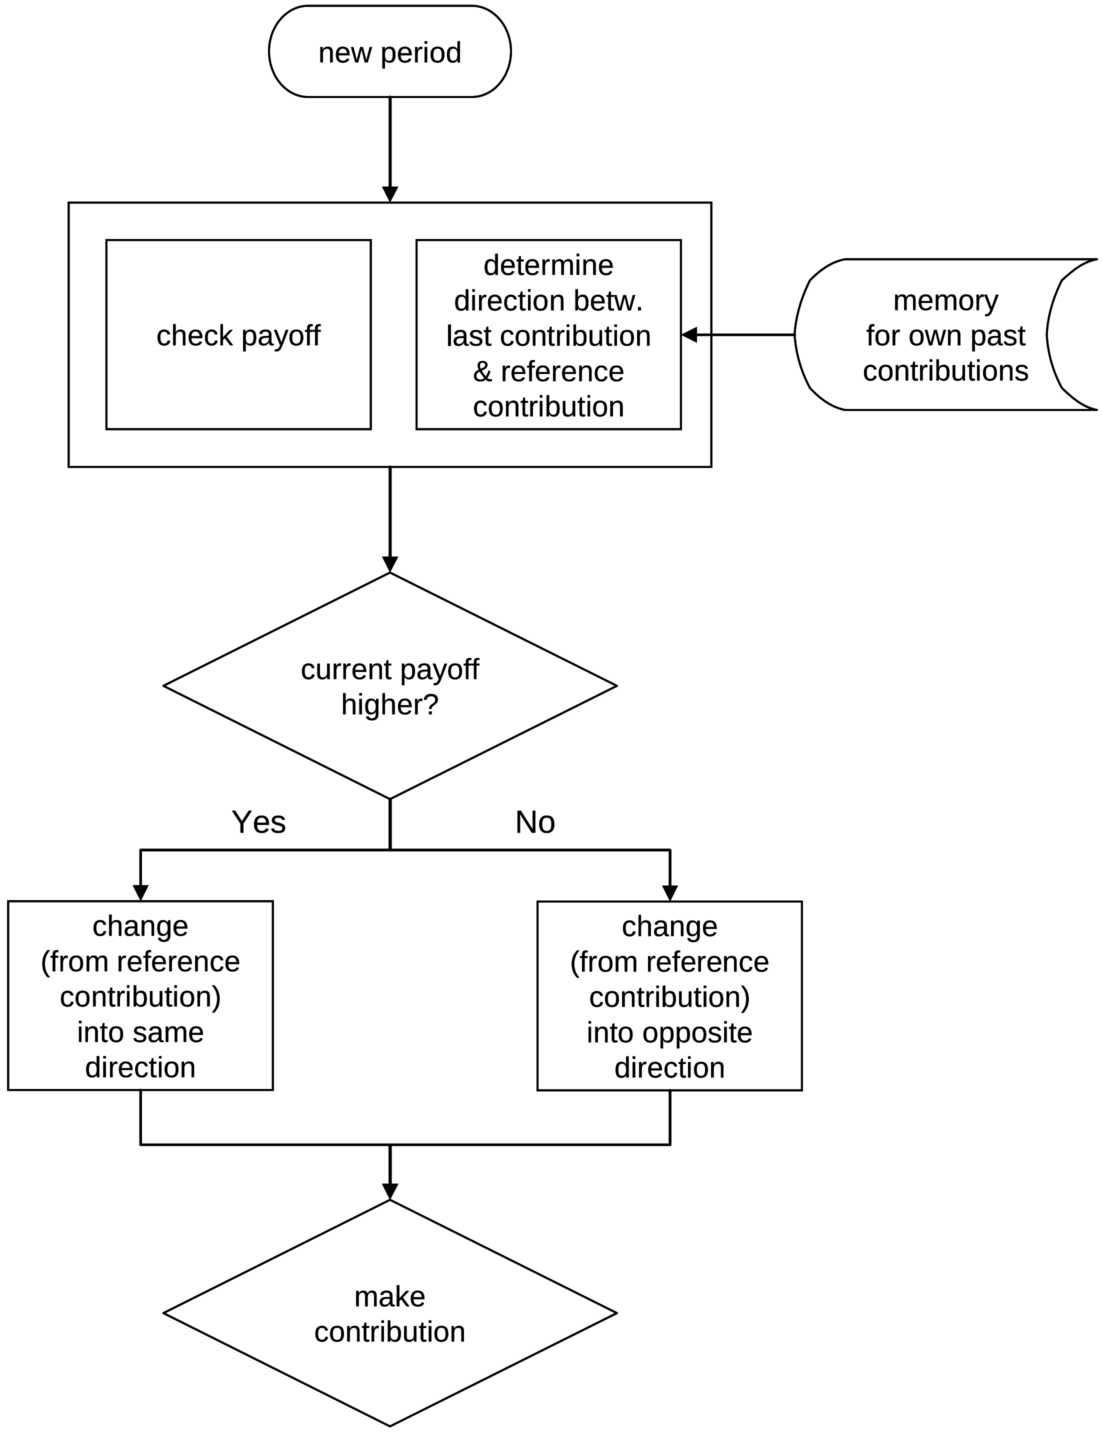

Figure 2.5. Schematic representation of local adaptation learning model (LOCAD). The reference contribution is the contribution leading to the highest payoff so far.

When the local adaptation learning model is applied to the public goods game or the SDN the following decision process results: The first allocation is selected randomly. Then a new, slightly different allocation is chosen. If the last allocation led to a better payoff compared to the second to last, the allocation of the current round is changed in the same direction as previously.

If the last allocation led to a lower payoff, the allocation of the current round is changed in the opposite direction. The distance (step size) from old to new allocations declines with every change. For larger payoff changes, larger step sizes result. For illustration, imagine an individual first contributing 50% of his resource to the public project and after that 60%. If the second decision led to an increased payoff, he will increase his contribution with the next decision to, for instance, 70%. Figure 2.5 shows a flow chart illustrating the general process of LOCAD.

Mathematically, LOCAD is defined as follows: At the beginning of a game, the first allocation is selected with equal probability. In the second round the probability of selecting an allocation is determined by

$$p_{2j} = f_S(x_{jk})/U = \exp\left[-(x_{jk} - s_2)^2 / 2\sigma_S^2\right]/U\ , \tag{2.7}$$

where x_{jk} is the Euclidean distance of any allocation j to the last chosen allocation k with the standard deviation σ_S as the first free parameter, and U as a constant that normalizes the probabilities p_{tj}, so that they sum to 1.

The step size, s_t, changes from round to round as follows (cf. Rieskamp et al., 2003):

$$s_t = \frac{s_1}{2}\frac{|\pi_{t-1} - \pi_{t-2}|}{\pi_b} + \frac{s_1}{t}\ , \tag{2.8}$$

where $s1$ is the second free parameter that determines the size of the initial step, $\pi_{t-}1$ is the payoff of the preceding round (with $\pi0 = 0$), and πb is the payoff of the reference allocation b, which is the allocation that produced the best payoff so far. The definition of the step size integrates two factors. First, smaller differences between the last two payoffs and larger reference payoffs lead to smaller step sizes for the current round. The second factor is time: The more rounds that have been played, the smaller is the step size. Note that for round 2 the step size $s2$ reduces to the initial step size, s_1.

For the third and all following rounds the probability of selecting any particular allocation is the product of two processes, one that selects the step size and one that selects the direction of change. The probability of selecting an allocation for any round $t > 2$ is defined by

$$p_{tj} = f_S(x_{jb}) \cdot f_A(y_{jk})/U\ , \tag{2.9}$$

where $f_S(x_{jb})$ determines the likelihood of the step size according to Equation 2.7, with x_{jb} as the Euclidean distance from any allocation j to the reference allocation b. The likelihood of the direction is determined by $f_A(y_{jk}) = \exp\left[-(y_{jk} - a_t)^2 / 2\sigma_A^2\right]$, with y_{jk} as the angle between the direction that led to the allocation k in the preceding round and the direction leading to any possible allocations j in the present round, and a_t equal to 0° if the previous allocation resulted in a higher or equal payoff than the reference allocation; otherwise *a*t equals 180°. The direction leading to the allocation $\boldsymbol{c}_k$ in the preceding round is defined as the normalized vector from the allocation at *t*-2 to the allocation at *t*-1. The direction of an allocation *j* in the present round is defined as the normalized vector from the reference allocation *b* to any allocation *j*. The angle[7] between two directions ranges from 0 to 180°. If the direction of the preceding round was successful (unsuccessful), then similar directions have higher (lower) probabilities. The standard deviation for the angle, σ_A, is the third free parameter of LOCAD and determines the probability of a direction as a function of its similarity (dissimilarity) with the direction as suggested by the evaluation of the last change.

The original LOCAD model of Rieskamp et al. (2003) was proposed for an individual allocation problem where a particular allocation always led to the same payoff. In such a stable environment the assumption of a reference allocation with which new allocations are compared appears reasonable. However it is uncertain whether the model also represents a good description of people's decision processes in a dynamic environment with interdependent payoffs. Here the assumption of a reference allocation, which always leads to the same payoff, might be less plausible, since the payoff of the reference allocation might change if others behave differently (i.e. contribute less or more). Thus, it is an empirical question whether the observed advantage of the LOCAD model compared with RLM for the individual decision-making problem as demonstrated in Rieskamp et al. (2003) can be generalized to group interaction.

2.4 Predictions of the Models

My goal was to test how well the reciprocity approach, represented by the reciprocity heuristic, the learning approach, represented by RLM and LOCAD, and the social motivations approach, represented by the theory of social value orientation, can predict individuals' decisions in two social dilemmas. According to the theory of social value orientation people's social preferences determine their contributions in a public goods game (see for example De Cremer & van Lange, 2001; but see also Parks, 1994). The theory predicts that individuals with a prosocial orientation will contribute more to the public project compared to individuals with an individualistic or competitive orientation, particularly in the first round of a game.

According to the reciprocity approach, individuals base their decisions solely on others' cooperation, neglecting their own payoffs. In contrast, the learning approach assumes that individuals' behavior is contingent only on positive or negative reinforcement, that is, their

---

[7] Mathematically the angle is determined by the arccosines of the vector product of the two normalized direction vectors. When computing Equation 2.9 the angle is expressed as radians.

received payoffs, whereas others' behavior is neglected. The two learning theories differ, as the RLM predicts that individuals probabilistically choose allocations proportional to the received payoffs, whereas LOCAD predicts that people change their allocations locally in specific directions, depending on the success of preceding changes.

Although the three computational models make quite different assumptions about cognitive processes, they can predict similar decisions, due to their flexibility resulting from the models' free parameters. Therefore I conduct a cross-validation Experiment, in which the models' parameters are fitted with a calibration sample and afterward the models are tested by using a validation sample. Because the models also predict what information should be acquired before a decision is made, I will also test the models with respect to the predicted information search. The reciprocity heuristic predicts that people should only search for information regarding the other players' contributions. In contrast, the learning models predict that people will only search for information regarding their own payoff. To test the models' information search predictions a computerized information board was used in Experiment 2.1 to monitor participants' search (for the method of information boards see Payne, Bettman, & Johnson, 1993).

2.5 Method

To test the different approaches, first participants' social value orientation was assessed and second, participants played either the public goods game or the SDN.

2.5.1 Participants and Procedure

The 60 participants with an average age of 24 years were mainly students from different departments of the Free University Berlin. Experiment 2.1 had three conditions, one SDN, and two public goods games; 20 participants were assigned randomly to each condition. In the SDN I applied an MPCR of .75; that is, every point invested in a two-person public project was multiplied by 1.5 and then equally divided between the two participants who could contribute to this project. In the public goods game-high (public goods game-low) condition I applied an MPCR of .75 (.375); that is, every point invested in the public project was multiplied by 3 (1.5) and then equally divided among the group members. The public goods game-high condition equates the payoff of a single contributing participant for the SDN and the public goods game; that is, in both games the fear of being exploited is equal. When participants in the public goods game-low condition contributed their whole endowment to the public project they received the same payoff as participants in the SDN who also contributed their whole endowment; thus the incentive for cooperation was identical in the two games.[8]

The experiment was conducted in two parts. In the first part participants' social value orientation was assessed. This took on average about 30 minutes and participants earned on average 4.7 euros as payment for their participation. In the second part of the experiment, conducted about 3-10 days later, participants played the games in groups of four. Payment was again made proportional to received payoffs, resulting in average earnings of 19.5 euros (including a 2.5 euro show-up payment) for part two.

2.5.2 Measuring Social Value Orientation

Social value orientation was measured with the triple-dominance measure of social values following van Lange, De Bruin, Otten, and Joireman (1997), where individuals repeatedly had to choose between three possible payoff distributions for themselves and an anonymous other. After a participant completed the questionnaire, one of his or her choices was selected randomly and the participant was given the corresponding payoff. In addition, the participant received a payment dependent on the choice of another participant. To determine this additional payoff, the participant had to choose one questionnaire from the stack of others' questionnaires (participants of the same session were not in the stack and in the first session questionnaires from a previous experiment were used) and a choice of this questionnaire was randomly selected to determine the additional payoff. Following van Lange (1999), participants were classified according to the three social value orientations, if at least six of their nine choices corresponded with one of the three types of social value orientations.

2.5.3 Playing the Public Goods Game and the Social Dilemma Network

Participants arrived individually at the laboratory and were seated in private rooms, preventing any personal interaction. Participants were instructed that they would take part in a repeated group decision-making task together with three other persons, all being endowed with 10 points per round, which could be allocated either to a private project or to a public project (to three public projects in the case of the SDN). Then the instructions explained the respective game, first in text form and then with some numerical examples. It was further explained that for every collected point 0.03 euros would be paid afterward to the participants and that the game

[8] According to Rapoport (1967) and Komorita, Chan and Parks (1993) the payoffs for exploited players (who cooperate while others defect) and payoffs for cooperative outcomes (when all players cooperate) are important determinants of cooperation rates in social dilemmas. Low payoffs for exploited players decrease cooperation and high payoffs for cooperative outcomes increase cooperation.

would be repeated several times, but that no participants knew how often. The experiment was neutrally described as a decision task, terms like "cooperation" or "free riding" were omitted and the instructions did not tell participants to achieve any particular goal. To check whether participants understood the game correctly, they had to complete a test quiz, which asked participants to calculate payoffs for different configurations of contributions. If participants failed to answer a question, the experimenter clarified any misunderstandings, so that all questions could be answered correctly. Finally participants received a description of the computerized information board used to conduct the game, which explained the kind of information participants could access and how to make their decisions. The game started after all participants had correctly completed a second test quiz regarding the interface.

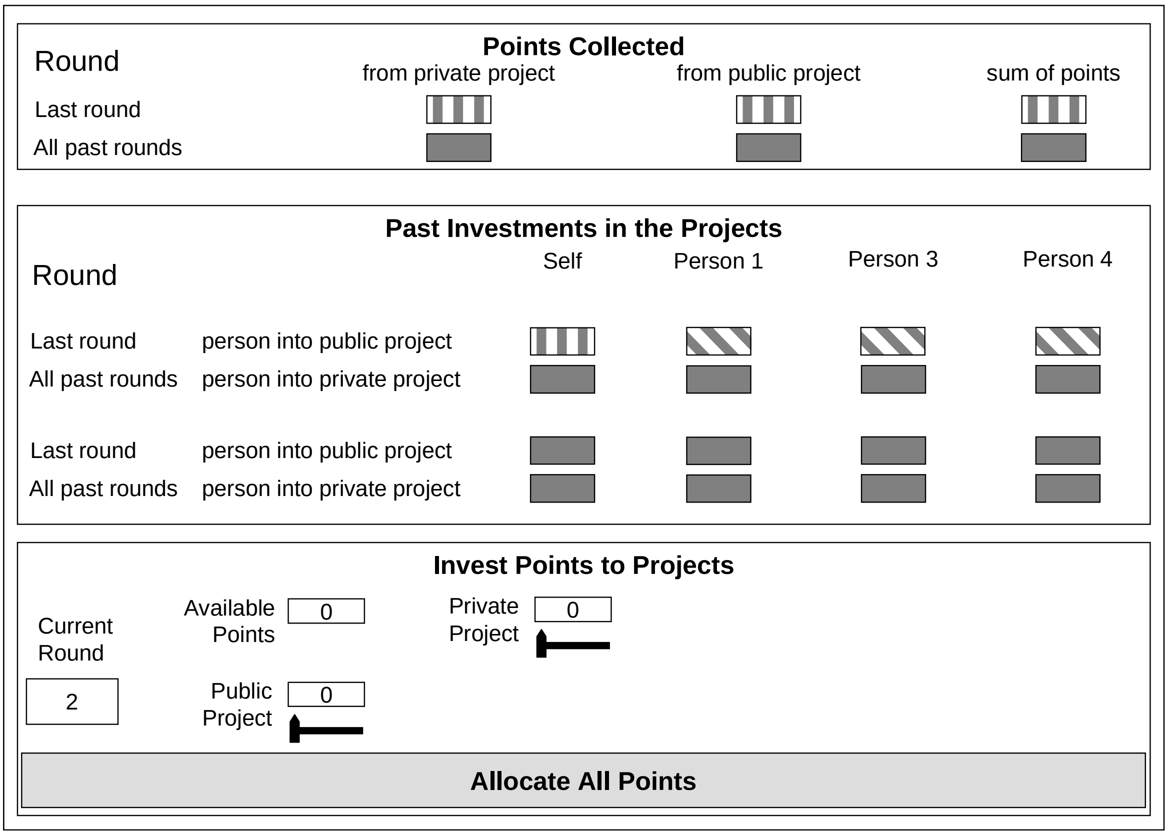

Figure 2.6. Player interface in the SDN. The figure shows a schematic picture of the interface participants used to play the public goods game (example for Player 2, original display was in German). Dark boxes are information boxes participants could click on. A box opened when it was clicked and stayed open until another box was clicked. Information behind boxes with diagonal lines was categorized as reciprocity search and information behind boxes with vertical lines was categorized as learning search.

The games were conducted on 4 desktop PCs connected via a LAN network. The software CING (Czienskowski, 2004) allowed participants to make their own choices as well as to access information about others' and own past choices. In the first round of a game, participants received their endowment and allocated it to the projects. After all participants had decided about their allocations in a round, the next round started automatically. From the second round on participants could access information on the computerized information board (see Figure 2.6), by clicking on the corresponding boxes. Participants could access information about their collected points, about others' contributions to the shared public projects and their private projects, and about their own allocation behavior; all information available was for the last round or for all past rounds cumulated. Clicking a new box automatically closed the previously opened box. Participants were allowed to open the boxes as often as they wanted and information search did not incur financial costs. After a decision was made, no further information could be acquired until the next round started. Participants could allocate their endowment at any time in a round; they were not obliged to search for information. Each game lasted for 30 rounds.

After the last round, participants were asked to complete the post-experimental questionnaire, which consisted of demographic questions, questions about participants' strategies, and a reciprocity questionnaire. Finally participants received their payments.

2.6 Results

I first describe participants' contributions in the games and test how contributions correspond to participants' social value orientation. Thereafter I report the comparison of the three models as they predict participants' decisions and their information search.

2.6.1 Contributions

First I compared the contributions to the public project(s) in the three different games averaged across 30 rounds. Median contributions were 9.2 in the SDN, 5.9 in the public goods game-high, and 5.8 in the public goods game-low condition. As a measure of effect size I computed $\tilde{\delta}$, which is the difference between the medians divided by the interquartile range of the pooled groups (see Grissom & Kim, 2001; Laird & Mosteller, 1990). Table 3.1 shows block-wise median contributions across the three games. As predicted, participants in SDN contributed more than those in public goods game-high (n = 10, U = 0, p = .009; $\tilde{\delta}$ = 2.00) and public goods game-low (n = 10, U = 0, p = .001; $\tilde{\delta}$ = 2.33). Contributions in public goods game-high and public goods game-low were not different (n = 10, U = 10, p = .668; $\tilde{\delta}$ = .75). To test for potential trends of contributions over time I aggregated the 30 rounds in 6 blocks of 5 rounds each and compared contributions over blocks. Only participants in the SDN showed increasing contributions from block 1 to block 6, $\chi 2$ (5, N = 20) = 17.66, p = .003, $\tilde{\delta}$ = .86.[9] Thus the contributions in the SDN were higher compared to both public goods games, and this difference is not only due to first round differences.

Table 2.1. *Median Contributions in the Social Dilemma Network (SDN) and public goods game (PG) conditions.*

Condition	Block						
	1	2	3	4	5	6	All blocks
SDN	8.3	8.8	8.8	9.3	9.8	10.0	9.0
public goods game-high	6.0	6.3	5.0	7.5	5.0	7.5	6.1
public goods game-low	5.3	5.5	5.8	5.5	5.0	6.0	5.5
All conditions	6.8	7.0	7.3	7.5	6.8	7.8	7.0

Note. Numbers are medians for the five groups per condition. Group values were computed as the groups' median contribution across five rounds.

To test the effect of social value orientation, I classified the 60 participants according to van Lange, De Bruin, Otten, and Joireman (1997): 36 participants were classified as prosocials, 17 as individualists, none as competitors, and 7 could not be classified as their choices were not sufficiently consistent. The effect of social value orientation on cooperation was tested for the first round (performing the same tests for the first block I found the same results) as well as across all rounds by comparing median contributions of prosocials with median contributions of individualists. Because the three games might moderate the effect of the social value orientation, each game was considered separately.

In no condition did contributions differ significantly between participants with different social value orientations. The median contributions in the first round of prosocials (individualists) were 6 (8) in the SDN, 6 (6) in the public goods game-high, and 5 (3) in the public goods game-low. The median contributions across all rounds of prosocials (individualists) were 9 (9) in the SDN, 6 (4) in the public goods game-high, and 6 (6) in the public goods game-

[9] When applying $\tilde{\delta}$ to dependent variables, I first calculated the difference values of the two variables, then calculated median and interquartile range of the difference values, and finally estimated the effect size as the quotient of median and interquartile range.

low. Analyzing the effect of social value orientation on contributions for all 60 participants together (i.e. pooling across the three games) also does not reveal an effect. As a further test, the number of prosocial choices in the triple dominance measure of social value orientation was correlated with participants' contributions, separately for the first contribution and the average contribution across the game (again pooling conditions). The correlations are $\tau = -.04$ ($p = .726$) for the first contribution and $\tau = .09$ ($p = .375$) for average contributions. In sum, participants' social value orientation did not substantially influence their contributions.

2.6.2 Model Comparison

The three models were tested with respect to how well they predict participants' decisions and participants' information search.

2.6.2.1 Predicting Contributions

All three models predict for each round the probability with which any possible allocation vector is selected, depending on individual decisions and payoffs in the preceding rounds and the models' parameters, which were fitted for each participant separately. To address the problem of over-fitting (Myung & Pitt, 2002), parameters were optimized by using only participants' decisions in the first 15 rounds and the models' predictions were then cross-validated for the remaining 15 rounds. As a goodness-of-fit measure the mean sum of squared errors (see e.g. Selten, 1998) was employed, by calculating for every possible allocation in a round the sum of squared differences of predicted probability and observed behavior (assigning the chosen allocation a value of 1 and all other allocations a value of 0). The mean sum of squared errors ranges from 0 to 2, with 0 as an optimal fit where the observed allocation is predicted with probability 1. To find the optimal parameters for each model and individual, first a grid search was performed to identify good start parameter values. These values were then optimized with the sequential quadratic programming method (Fletcher & Powell, 1963), as implemented by the program "Matlab". The parameters of REH were restricted to $0 \leq \gamma \leq 1$ for the generosity parameter, and to $.1 \leq \sigma\mathrm{c} \leq \sqrt{200}$ for the error parameter (the square root of 200 is the maximum Euclidean distance between two allocations). For the RLM the initial attraction was restricted to $1 \leq w \leq 30$, the standard deviation for the reinforcement to $.1 \leq \sigma\mathrm{R} \leq \sqrt{200}$, and the forgetting parameter to $.1 \leq \phi \leq 1$. For LOCAD the initial step size was restricted to $1 \leq s1 \leq \sqrt{200}$, the standard deviation of the step size to $.1 \leq \sigma\mathrm{S} \leq \sqrt{200}$, and the standard deviation of the direction angle to $6° \leq \sigma\mathrm{A} \leq 360°$. Table 3.2 shows medians of fitted parameter values for all models in the SDN and for the two public goods game conditions together.

Table 2.2. *Parameter Values and Model Fits for the Two Games.*

Game	REH	RLM	LOCAD
SDN	Model fit SSE .93 (.89, .99)	Model fit SSE .92 (.76, .97)	Model fit SSE .99 (.92, 1)
	SD allocation 1.26 (.84, 1.86)	Initial 1 (1, 11.28)	SD angle σ_A 95° (72°, 139°)
	Generosity γ .34 (.02, .75)	Forgetting ϕ .7 (.42, .88)	SD step size σ_S 1.28 (.69, 2.15)
		SD .7 (.62, 1.27)	Initial step size $s1$ .98 (.1, 2.44)
public goods game	Model fit SSE .88 (.84, .91)	Model fit SSE .86 (.79, .89)	Model fit SSE .92 (.91, .93)
	SD allocation 3.61 (2.15, 5.66)	Initial 1 (1, 6.44)	SD angle σ_A 205° (116°, 256°)
	Generosity γ .32 (0, .51)	Forgetting ϕ .27 (0, .64)	SD step size σ_S 8.89 (4.78, 14.14)
		SD 1.27 (.1, 1.86)	Initial step size $s1$ .1 (.1, 5.66)

Note. Medians of the mean sum of squared errors (SSE) were calculated based on the last 15 rounds of the games. Median parameter values were fitted for individuals to the first 15 rounds of the games. Numbers in parenthesis are first and third quartile values. REH: reciprocity heuristic; RLM: reinforcement learning model; LOCAD: local adaptation learning model.

For REH the best fits were achieved with a positive generosity parameter, γ, implying that most participants tended to contribute somewhat more than the others did in the round before. Whereas the smaller σC for the SDN, compared to the public goods game conditions, shows that the most likely predicted allocation was usually similar to the observed allocations, the higher σ_C for the public goods game conditions indicates that large standard deviations compensated for the divergence of observed and predicted behavior, by assigning relatively higher (lower) probabilities to allocation alternatives that deviated from (corresponded to) the most likely reciprocal allocation.

The parameters fitted for RLM resulted in, on average, low values for the initial attraction parameters, *w*, implying that already the first decision and the corresponding reinforcement will have a large impact on the updated expectancies and thereby on the following decision. The on average low values obtained for the standard deviation imply weak generalization of reinforcements; that is, the selected allocation receives a relatively large reinforcement and similar allocations receive only very small reinforcement. The forgetting rate, Φ, is smaller for the public goods game conditions compared with the SDN condition. The higher Φ for the SDN condition works in the same direction as the low *w* does at the beginning of the game; behavior is largely determined by the most recent reinforcements. The combination of a low initial attraction, little generalization, and a high forgetting rate for RLM, in particular in the case of the SDN, describes a decision process of simply repeating the allocation of the preceding round (see Figures 7 and 8 for illustration).

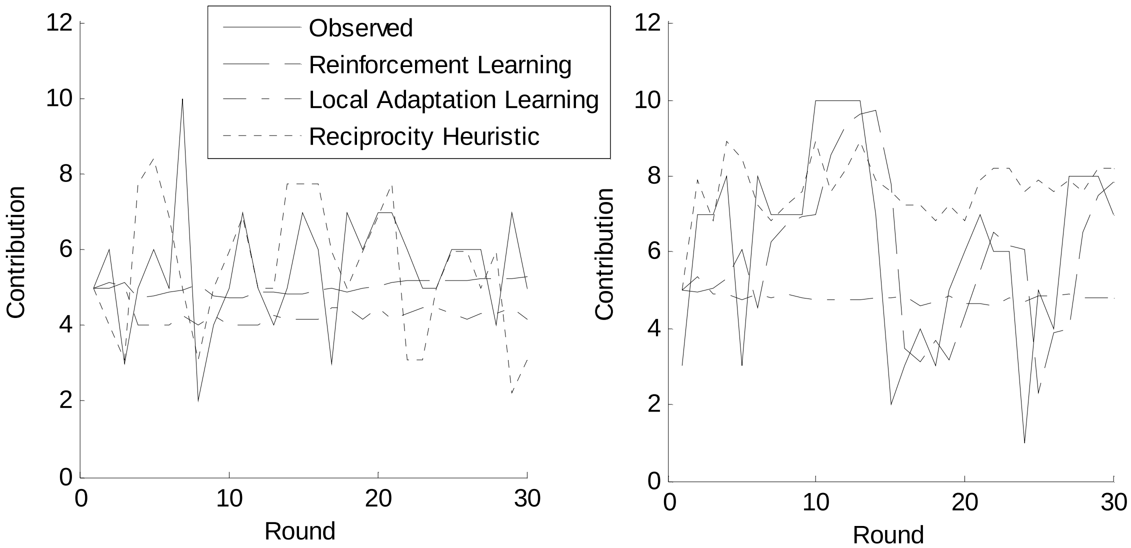

Figure 2.7. Observed and predicted contributions in the public goods conditions (examples). Contributions predicted by the models are computed as the weighted average of all possible allocations in a round, with predicted probabilities of allocation options as weights. Chart A depicts observed and predicted contributions, for a participant who could be modeled better with the reciprocity heuristic compared to reinforcement learning. In Chart B, reinforcement learning predicts behavior more accurately. RLM: reinforcement learning model; LOCAD: local adaptation learning model; REH: reciprocity heuristic.

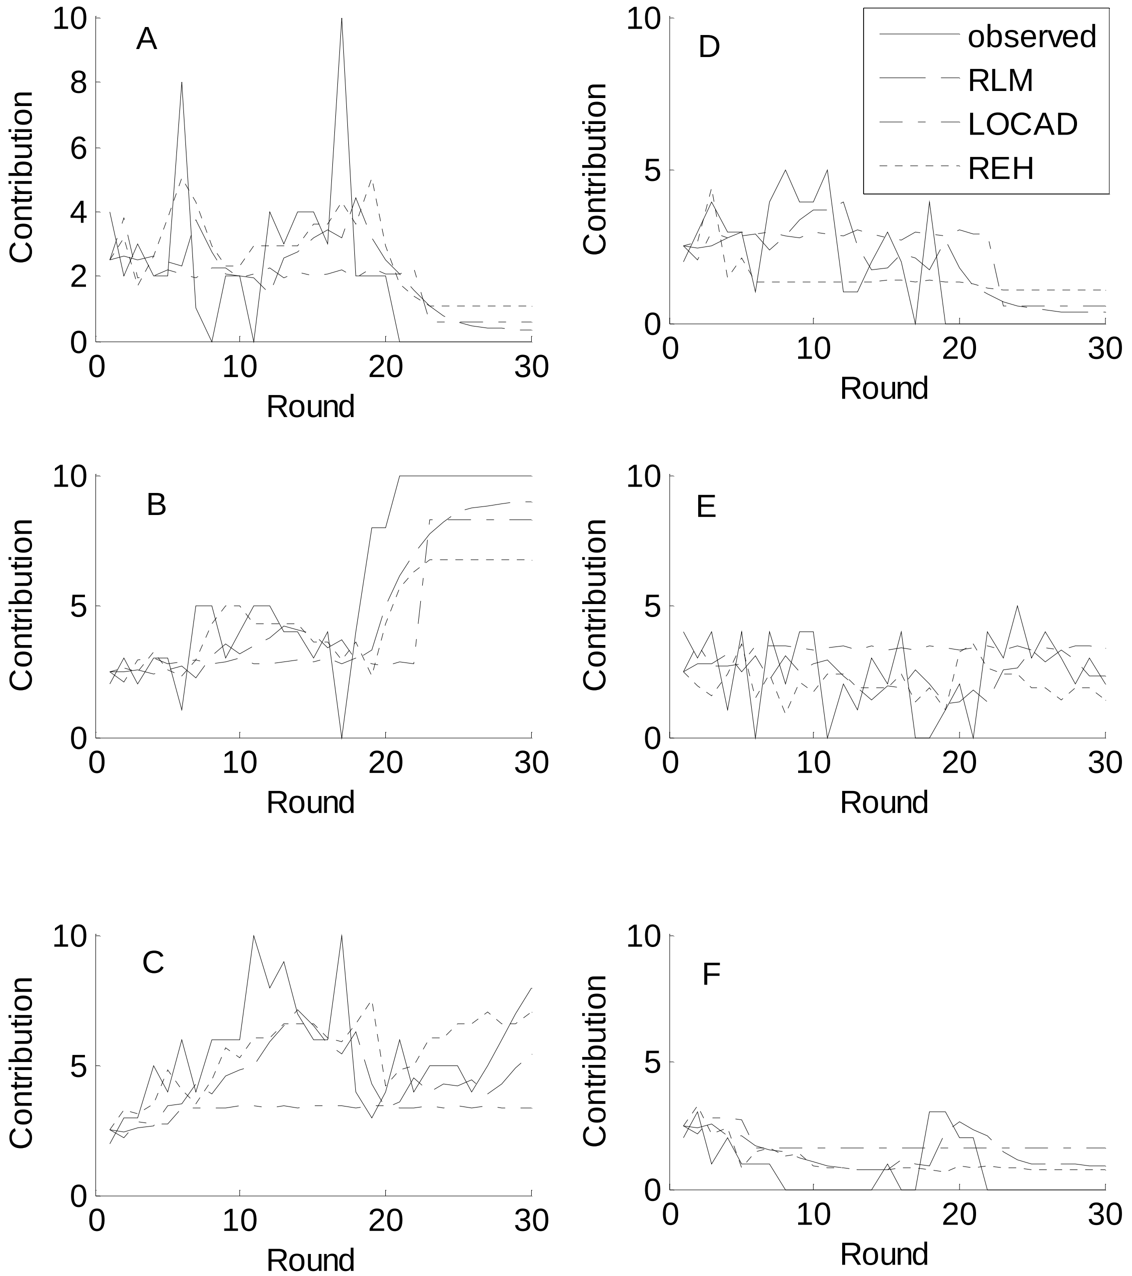

Figure 2.8. Observed and predicted contributions in the social dilemma network (examples). The charts depict observed and predicted contributions of two participants in the SDN to their private projects (Charts A and D) and to two of the public projects. For the first player (Charts A, B, C), reinforcement learning and the reciprocity heuristic can both predict contributions. The second player serves as an example where no model made accurate predictions.

For LOCAD medium to high values were best for the standard deviation of the angle (σA); thereby the model predicts large deviations from previously successful directions. In the case of the SDN this is combined with a small initial step size ($s1$) and a small standard deviation of the step size (σ_S), so that the model predicts that from the start of the game only small modifications of the allocations are made. In the case of the public goods game, the optimized parameter values for standard deviation of the step size are larger compared to the SDN, so that LOCAD predicts somewhat larger changes of behavior in the public goods game.

In sum, the optimized parameter values for three models indicate that RLM tends to mimic a strategy that always repeats the allocation of the previous round, especially in the SDN. The changes predicted by LOCAD are constrained by the relatively small step size. Finally, the REH compensates for non-reciprocal components of participants' behavior in the public goods game conditions by having a higher standard deviation and the fitted values for the generosity parameter suggest that participants use a generous reciprocal strategy.

Table 3.2 contains median model fits and Figures 7 and 8 illustrate the performance of the models by plotting their predictions together with observed contributions. As the models were used to predict individual behavior, I show the models' predictions for two individuals, one example where the models provide an adequate description and one where the models provide a less adequate description of the decision process. To select the best model in predicting the behavior in both games, I first compared the models' fits with a baseline model before they were compared with each other.

The static baseline model predicts a constant probability distribution across all rounds of the game with which the possible allocations are selected. Specifically, given the most likely allocation, which is a player's mean allocation in the first 15 rounds, the probability of choosing an allocation j is $p_j = \exp\left(-x_{jl} / 2\sigma_C^2\right) / U$ with x_{jl} as the Euclidean distance between the most likely allocation and any possible allocation, σ_C as a free parameter constrained to $.1 \leq \sigma C \leq \sqrt{200}$, and U as a constant that normalizes the sum of all probabilities to 1. Although, the baseline is static, it is already a strong competitor for the other models, since its predictions are based on participants' decisions. The competing models can only outperform the baseline if participants change their behavior dependent on others' behavior or their own payoff and when this conditional behavior can be predicted by the competing models.

The three computational models were compared with the baseline model (separately for the two games) by considering each model's prediction of participants' contribution in the last 15 rounds, that is, the crucial cross-validation sample. For each participant it was determined whether the computational model or the baseline predicted behavior better. The percentages of participants for which each of the three computational models made better predictions compared to the baseline are provided in Table 3. In the case of the social dilemma network, REH and

RLM outperformed the baseline in predicting participants' decisions for 80% (T = 4, p = .004) and 95% of the participants (T = 1, p = .001), respectively, whereas the LOCAD only outperformed the baseline model for 65% of the participants (T = 7, p = .111). In the case of the public goods game, only REH and RLM outperformed the baseline in predicting participants' contributions for 60% (T = 15, p = .009) and 63% (T = 16, p = .033), respectively. In sum, with the exception of LOCAD, the models demonstrated their ability to predict participants' decisions by taking the decision of other participants (for REH) or payoff information (for RLM) into account. To assess the effect size of the models' fits compared to the baseline's fit I computed for every participant and model a difference score between the model and baseline model. In the SDN condition I found a medium effect size of $\tilde{\delta}$ = .57 for REH and of $\tilde{\delta}$ = .40 for the RLM, and a small effect sizes of $\tilde{\delta}$ = .14 for LOCAD. In the public goods game effect sizes were smaller with $\tilde{\delta}$ = .24 for REH and $\tilde{\delta}$ = .08 for RLM, and the baseline was better than LOCAD with an effect size of $\tilde{\delta}$ = -.53.

Table 2.3. *Pair-Wise Comparison of Models with Baseline and with Each Other.*

Game	Model	RLM	LOCAD	REH
SDN (n=20)	Baseline	95% (T = 1, p = .001)	65% (T = 7, p = .111)	80% (T = 4, p = .004)
	RLM		5% (T = 1, p = .001)	45% (T = 9, p = .512)
	LOCAD			70% (T = 6, p = .199)
public goods game (n = 40)	Baseline	63% (T = 15, p = .009)	15% (T = 6, p = .001)	60% (T = 16, p = .033)
	RLM		8% (T = 3, p = .001)	43% (T = 17, p = .274)
	LOCAD			85% (T = 6, p = .001)

Note. Each cell shows the percentage of comparisons the column model was better than the row model. Results of Wilcoxon tests are presented in parentheses. All comparisons with the baseline were conducted as one-sided tests.

As a second step the three models were compared with each other. In the case of the social dilemma network the RLM outperformed LOCAD in predicting the decisions for 95% of the participants and REH performed better than LOCAD for 70% of all participants (see Table 3), whereas REH and RLM performed equally well. In the public goods dilemma, RLM and REH outperformed LOCAD in predicting the contributions for the majority of participants (RLM 92%, REH 85%), whereas RLM and REH again did equally well in predicting the contributions. In sum, the model comparison leads to the conclusion that REH and RLM were the best models for predicting participants' contributions. Both models were better than the baseline model and LOCAD in both social dilemmas. However, when comparing RLM with REH, the two models predicted participants' behavior equally well, so that based on participants' contribution decisions it is not possible to decide which model should be preferred.

2.6.2.2 Predicting Information Search

I used the models' predictions of participants' information search as a second criterion to test the models. The reciprocity model predicts that participants only need to search for information regarding other participants' contributions to the public projects while neglecting information about their own payoffs. In contrast, the two learning models, regardless of whether LOCAD or RLM, predict that participants will search for information regarding their own payoffs while neglecting information about others' contributions. Therefore, I classified participants' information search in each round according to these two categories: If a participant *exclusively* searched for information regarding the other players' contributions in the last round, then this search pattern was classified as consistent with the reciprocity model. In contrast, if a participant *exclusively* searched for information regarding his own payoff or behavior in the last round, this search pattern was classified as consistent with the learning approach. However, if, in any given round, participants searched for information belonging to both categories, this search remained unclassified.[10] Figure 2.6 shows how information search was classified.

To test whether participants searched for information according to the reciprocity model or the learning models I determined with which model participants' information search in the majority of rounds was consistent, ignoring all rounds in which the search remained unclassified. In the SDN, 80% of the participants primarily searched for information according to the reciprocity model, whereas the remaining 20% primarily searched for information according to the learning models, χ^2 (1, N = 20) = 7.20, p = .007. In the public goods games a similar result was obtained: 62.5% of the participants primarily searched for information according to the reciprocity model, whereas the remaining 37.5% primarily searched for information according to the learning models, χ2 (1, N = 40) = 2.50, p = .114. Thus, participants' search was consistent with the reciprocity model, in particular for the SDN compared with the public goods games.

For a more detailed picture of individuals' information search, I analyzed the time allocated to and the number of clicks for different categories of information. The available

[10] One might argue that by searching for some of the information, other information could be inferred. However, since information search did not cost anything, I think it is implausible to assume that participants did the cognitively demanding task of making these inferences, instead of directly searching for the information they wanted.

information was categorized as above; however, the information search remained unclassified only if it was consistent with neither the reciprocity model nor the learning models. Figure 2.9 shows how individuals allocated their attention, measured as the number of box openings and amount of time spent for the different types of information. The time spent for the two types of information and the number of clicks on information boxes of the two types are highly correlated and lead to the same conclusions. Therefore I focus on the time spent for the two types of information, which in relative terms remained constant across rounds.

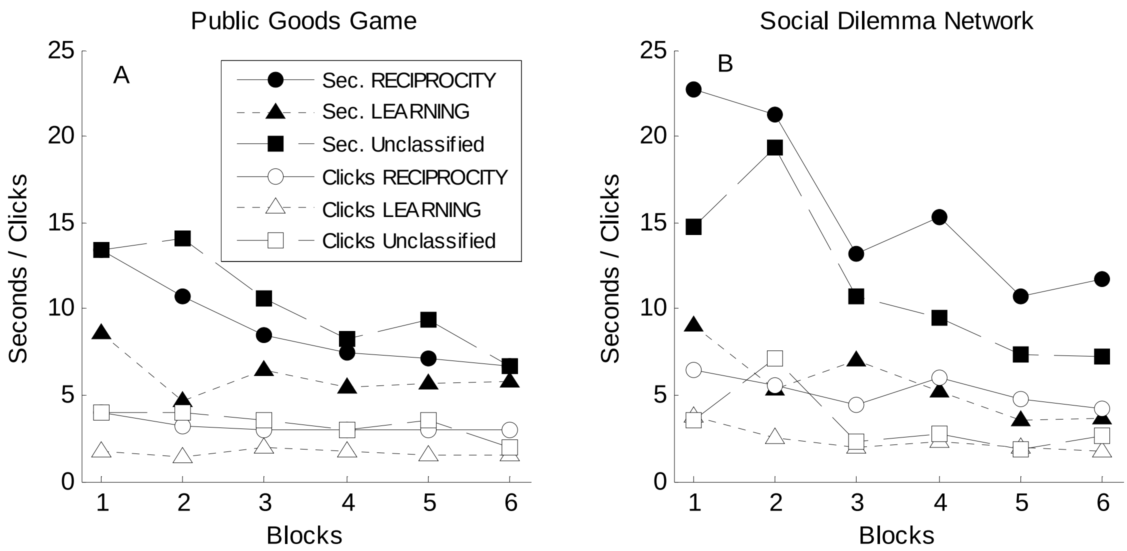

Figure 2.9. Participants' information search. Given are median seconds per category and median clicks per category for each block. Blocks consisted of six or five (first block) rounds.

In both games participants spent more time on information needed for reciprocity (SDN: *Mdn* = 14.73 s, public goods game: *Mdn* = 8.08 s) than for information corresponding to the learning model (SDN: *Mdn* = 3.38 s, public goods game: *Mdn* = 4.73 s); SDN: $T = 7$, $p < .001$; $\tilde{\delta} = .96$, and public goods game: $T = 101$, $p < .001$, $\tilde{\delta} = .61$. When taking into account that a larger proportion of information boxes are considered to be consistent with the learning models compared to being consistent with the reciprocity model (see Figure 2.9), this results provides even stronger evidence for a reciprocity-based decision process.

2.7 Discussion

The present experiment examined two main questions: Are individuals in groups more cooperative when they can select partners? What models are best in predicting individuals' contributions in social dilemmas in groups? As predicted, I found that participants cooperated more if they could select partners. Whereas participants in the public goods games mostly started with and maintained an intermediate level of cooperation, participants in the SDN started with slightly higher cooperation followed by increasing cooperation. This result is especially remarkable, because the efficiency gain through cooperation in the public goods game-high condition was twice that of the SDN condition. Participants' social value orientation did not correlate with their cooperation in the games. When comparing the different models, the reciprocity heuristic was best in predicting individuals' allocations *and* their information search, in particular in the SDN. The parameter values of the reciprocity heuristic show that individuals were more generous than a strict reciprocity principle would predict.

2.7.1 Cooperation

In contrast to the results of other public goods experiments (Ledyard, 1995) on average I did not observe declining cooperation in the public goods games and the average cooperation rate was relatively high. This can be attributed to three causes. First, decreasing cooperation rates are usually found in finitely repeated games where the number of iterations is known to the participants. This is an important point, because once participants are in the final round cooperation can no longer be reciprocated. In fact, by backward induction, the game-theoretical solution is derived, that one should not cooperate in any round when the game is finitely repeated. Although participants usually do not follow the game-theoretical solution in a finitely repeated game, cooperation decreases toward the end of the game (Selten & Stöcker, 1986). By not telling participants the number of rounds in the present Experiment, the games became indefinitely repeated and no "end game effects" occurred. Second, in the experiment, information about other players' individual contributions to the public projects was accessible, instead of only information about the average contributions. This could also have increased the cooperation rate, as other experiments have demonstrated (Croson & Marks, 1998; Sell & Wilson, 1991). Third, in the public goods game-high condition a relatively high MPCR was employed, which can also prevent the decline of cooperation, as shown for instance by Isaac, Walker, and Williams (1994).

Providing participants with the possibility to select their interaction partners as in the SDN substantially increased cooperation. This result is in line with previous research. For instance in an experiment of Riedl and Ule (2002) participants had to repeatedly select one strategy for simultaneously played prisoner's dilemmas and cooperated more if they had the possibility to selectively exclude players. In another experiment reported by Yamagishi, Hayashi, and Jin (1994), participants could repeatedly select a partner from a four-person group to play a continuous prisoner's dilemma. Here the possibility of partner selection led many participants to choose the same partners repeatedly—they formed a commitment relationship (Hayashi & Yamagishi, 1998; Kollock, 1998)—and to contribute their whole endowment. Coricelli et al. (2004) let participants endogenously select group members from a population of 16 participants to play a four-person public goods game and they also report higher contributions in the condition of partner selection. Additionally to Coricelli et al., I showed that partner selection increases cooperation even when combined with lower efficiency gains, thus highlighting the strong positive effect of partner selection. While I examined the effect of decomposing a four-person public goods game into multiple two-person public projects, I argue that this approach can also be used for larger groups, which can be split into several smaller groups (c.f. Kameda, Stasson, Davis, Parks, & Zimmerman, 1992).

2.7.2 Explaining Cooperation

Contrary to findings of van Lange (1999) and De Cremer and van Lange (2001), participants' social value orientation did not correlate with cooperation. People's behavior in a repeated game, as I have demonstrated by the superiority of the computational models compared to the baseline model, depends on others' behavior, explaining why participants' average cooperation rate in the games did not correlate with their value orientation. This result might be due to the method I used to measure social value orientation, which differed from previous experiments. First, participants' social value orientation was measured a few days before the participants interacted in the social dilemmas. Second, whereas in other experiments participants often made hypothetical choices to assess their social value orientation, in experimentthey were paid according to their choices (for the effect of performance-contingent payment on behavior see Hertwig & Ortmann, 2003). While I acknowledge that the specifics of the measurement procedure of social value orientation could explain why a correlation with participants' cooperation rate was not observed, I also note that others have also reported mixed results about the effect of social value orientation on cooperation in social dilemmas (Parks, 1994). In sum, given the results, behavior in iterated social dilemmas is less likely to be explained by the assumption that individuals simply maximize social utility.

Among the three models, LOCAD was least able to predict individuals' allocations. The concern that LOCAD might not predict behavior in a dynamic and interdependent environment, in which the same allocation can lead to different outcomes depending on others' decisions, appears justified. This does not necessarily negate the assumption that individuals adapt locally when learning. For instance, I tested a modified version of LOCAD, which does not consider all past choices and payoffs but simply chooses the better of the last two allocations as the reference allocation. This version outperformed the original LOCAD model, although it did not outperform REH or RLM. More important, the analysis of the information search demonstrates that participants did not search for information as it would be predicted by any learning model I considered. Thus, the idea that people's cooperation in social dilemmas can be described by a directional learning process received little support from the results.

The simple reinforcement learning model considered (RLM) predicted participants' contributions better than LOCAD and equally well as the reciprocity heuristic. However, when considering participants' information search the results speak in favor of the reciprocity heuristic. In contrast to the prediction of the learning model, participants did not mainly search for information regarding their own payoff, but more often only looked up information about how the other players behaved, as predicted by the reciprocity heuristic. Moreover, when examining the optimized parameter values for the RLM, in particular for the SDN, it becomes clear that for many individuals the model predicts in any round the observed behavior of the previous round. Such a model achieves a good data fit when behavior changes little from round to round, but it does not capture the general idea of reinforcement learning—that behavior depends on the history of past experience and the different outcomes associated with different behaviors. Thus, RLM's good fit should only cautiously be interpreted as supporting the hypothesis that reinforcement learning underlies participants' decision processes.

In contrast to both learning models, the reciprocity heuristic, REH, did well in predicting individuals' contributions *and* their information search, making REH the best model to predict participants' behavior in both games.[11] An examination of the optimized parameter values of the REH points out the specific reciprocity heuristic people use. The finding of an, on average, elevated generosity parameter value shows that people do not reciprocate others' contributions with exactly the same contributions, but instead are more generous than others. This also implies that in the case when others make no contributions at all, the REH still predicts some contributions. This finding is consistent with results from the simulations of To (1988) and Hayashi and Yamagishi (1998), who report that a nice reciprocal strategy, which does not retaliate immediately, performs best in social dilemmas (but see also Komorita et al., 1993), presumably because this allows people to maintain cooperation in the face of occasional defections. In the SDN the reciprocity heuristic predicted contributions much better compared to

[11] This decision is also supported by the analysis of the strategies described by the participants of the SDN condition, as those frequently described reciprocal strategies. When individuals described non-cooperative strategies, they still frequently included reciprocal moves that were aimed at maintaining others' cooperation. However, as the validity of introspective reports is debated (Ericsson & Simon, 1993; Nisbett & Wilson, 1977) I did not include self reports as a criterion for model selection.

the baseline, while this advantage was weaker in the public goods game[12]. This is in line with the prediction that a reciprocity mechanism works well and is thus applied more frequently or strictly when people have the possibility to select their partners.

The conclusion that the reciprocity model is a better description of individuals' decision processes for social dilemmas in groups compared with the competing learning models is subject to some limitations. There might be more sophisticated, alternative learning models that implement, for example, belief learning (Camerer & Ho, 1999a), which could perform better in predicting behavior. Also, I assumed that information search about others' contributions is consistent with the reciprocity model but inconsistent with a learning model. However, one could argue that by searching for information about others' contributions people are able to compute their own payoffs—the information needed for a learning mechanism. But this seems unlikely when payoff information is directly available. Also, the analysis of the information search needs to be limited due to the proportion of search that remained unclassified. Finally, since information search did not lead to any costs, participants might have looked up information that they did not really require for the decision strategy they used. One possible way to reduce variability in individuals' information search could be to induce search costs.

When compared with the results of other experiments (Fischbacher et al., 2001; Keser & van Winden, 2000; Kurzban et al., 2001), the evidence for reciprocal behavior seems less conclusive in the public goods games of the experiment. This may have resulted from the different approaches used to test for reciprocal behavior. Keser and van Winden tested whether players' contributions in a repeated public goods game changed in the direction of the other players' average contribution in the preceding round. The REH makes more precise predictions of the contribution magnitude and is thus a stricter test for reciprocation. This suggests that, had I used a less precise formalization of reciprocity, I could have found more evidence for reciprocal behavior. This seems unlikely in light of participants' descriptions of their strategies, as only a few participants in the public goods game conditions described a purely reciprocal strategy (see Footnote 12). There are also some important differences in my Experiment compared to Fischbacher et al.'s (2001), where participants played a one-shot public goods game and were confronted with all possible average contributions of other players in the same round and asked how much they would contribute conditional on the other players' average contributions. In contrast to my Experiment, participants in Fischbacher et al.'s Experiment only had information

[12] Participants' descriptions of their strategies suggest that this is due to more complex strategies, which, for instance, (a) change over time, or (b) try to maintain others' cooperation while sometimes exploiting them, or (c) try to "teach" others to cooperate by unconditional cooperation in the first rounds.

about others' *average* contribution and the game was played only once. In my view, reciprocal behavior was easier to realize for participants compared to my experiment, because others' average contribution was the only salient information available. Reciprocal behavior was also easier to employ for participants in Kurzban et al.'s experiment, because individuals could adjust their contribution *within* a round of an iterated public goods game, based on information about others' individual contributions in the *same* round (sequential contribution mechanism). The evidence from Fischbacher et al. (2001), Kurzban et al. (2001), Keser and van Winden (2000), and the present experiment indicates that at least some individuals act reciprocally in public goods games and many participants expect others to reciprocate. The extent to which reciprocal behavior is realized, however, seems to depend on an interaction structure favorable to reciprocity, like partner selection or sequential contribution with binding commitments. Accordingly, Komorita et al. (1992) and Bornstein, Erev, and Goren (1994), who conducted public goods games without partner selection or binding commitments, could not identify a reciprocal strategy, when looking at games where no programmed strategies were involved.

2.7.3 Conclusion

Different conceptualizations of reciprocity have been suggested. Similar to my approach, Kurzban et al. (2001) and Komorita and Parks (1999) conceptualized reciprocity as a strategy for social interaction. In contrast, Fischbacher et al. (2001) and Bolton et al. (1998) defined reciprocity as a distributional preference and Perugini et al. (2003) described reciprocity as a personal norm. Following my interpretation, reciprocity is a strategy individuals use to achieve maximum joint payoffs and to maintain others' cooperation, when they engage in sequential or repeated interaction. Additional goals that are compatible with the application of a reciprocity strategy are the stimulation of others' cooperation (see Komorita & Parks, 1999; Komorita et al., 1992) and punishment of non-cooperators (Fehr & Henrich, 2003). Alternatively, people might also use the strategy because they have a social preference for using it consistent with the interpretation that people have a social preference for reciprocity. However, reasoning about the motives behind reciprocal behavior is difficult, and existing work does not seem to provide an unambiguous answer to these questions. While many authors have argued that a necessary condition for reciprocity to function is repeated interaction (see e.g. Axelrod & Hamilton, 1981; Bixenstine & Gaebelein, 1971; Komorita & Parks, 1999; Trivers, 1971), specifically designed experiments with both repeated games and one-shot games will be needed to distinguish between selfish and social motivations behind reciprocity.