CHAPTER 3

The Golden Rule in Groups II:

Do Reciprocal Cooperators Care About Others' Intentions?

Chapter 2 compared different approaches explaining cooperation in groups and found that people's behavior is best explained by the reciprocity heuristic. When formalizing reciprocity, I assumed that people reciprocate others' contributions, that is, people contribute the same absolute amount as others did. While this is a straightforward way to implement reciprocity, in most experiments used to investigate cooperative behavior in groups, where endowments are constant between participants and over time, it seems less straight-forward to predict how reciprocators will behave, if endowments vary between persons and over time. Specifically, when endowments are equal, players who signal positive intentions, by contributing high amounts, also create positive consequences for others. Hence it is not possible to infer which aspect of others' contribution behavior reciprocators reciprocate when cooperating in groups. Is it the positive consequence that arises from the high contribution (in absolute terms), or is it the positive intention, or niceness, signaled by the high contribution. Therefore the aim of Chapter 3 is to investigate, which aspects of others' behavior reciprocators reciprocate.

3.1 Introduction

Cooperative groups are essential cornerstones of human culture. While this observation is common sense, it hides a puzzle: How do groups maintain cooperation? To examine this question, researchers have observed individuals in public goods games in which players decide if/or how much they will contribute to a public good (Dawes, 1980; Ostrom & Walker, 2003). All contributions to the public good increase in their value and are split equally among group members, regardless of individual contribution decisions. Withheld endowments maintain their original value for the player. Although a cooperative state in which everyone cooperates is beneficial to all group members, the fact that free riders cannot be excluded from the public good can deter cooperation. Specifically cooperators will always be worse off than free riders in the same group; hence self-interested individuals should not cooperate (in finitely repeated games) following this game-theoretic argument.

Yet contrary to the pessimistic game-theoretic prediction, cooperation in groups is often observed (Dawes, 1980; Kollock, 1998; Ledyard, 1995)[13]. Some people cooperate out of altruism, but the proportion of altruists is too small to explain the high levels of cooperation

[13] Strictly speaking, game theory predicts defection only for finitely repeated games. All experiments mentioned in this article used finitely repeated games.

frequently observed (Andreoni & Miller, 1993). Another explanation for the levels of cooperation seen—in repeated interaction—is reciprocity (Axelrod & Hamilton, 1981; Fehr & Fischbacher, 2003; Komorita & Parks, 1999; Kreps & et al., 1982; Rapoport & Chammah, 1965; Sugden, 1984; Trivers, 1971). Reciprocators maintain cooperation by imitating cooperative behavior, and they protect themselves from repeated exploitation by imitating defection. Experimental evidence supports the assumption of reciprocal cooperation in groups. For instance, Chapter 2 showed that a reciprocity heuristic predicted participants' behavior in a public goods game better than learning models (Erev & Roth, 1998; Rieskamp et al., 2003) and participants' social value orientation (e.g. van Lange, 1999). Kurzban and Houser (2005) demonstrated that 63% of their participants could be classified as reciprocators. Evidence for reciprocation has been found not only in further experiments (Fischbacher et al., 2001; Keser & van Winden, 2000; Kurzban & Houser, 2001), but also in field studies (Carpenter & Seki, 2005; Frey & Meier, 2004; Heldt, 2004).

This article examines which aspect of others' behavior reciprocators imitate when they cooperate in groups. Do people reciprocate based on the consequences of others' behaviors or do they reciprocate others' intentions? Whereas the role of intentionality for cooperative behavior has been examined in dyadic games such as the ultimatum game (Blount, 1995) and the trust game (Kevin A. McCabe et al., 2003), little is know about how perceived intentions influence cooperation in groups. Previous experiments examining reciprocity in groups (e.g. Fischbacher et al., 2001; Komorita et al., 1993; Kurzban & Houser, 2005) always examined public goods games in which players with positive intentions could contribute substantial amounts, but not situations in which participants with positive intentions could only make small contributions. Because these experiments did not allow intentions and absolute contributions to vary independently, they were not suited to derive conclusions about which aspects of others' behavior reciprocators imitate in public goods games.

I will first present the consequential and the intentional approach to reciprocal cooperation. Then I will formulate two computational models of reciprocity representing these approaches—and derive the models' different predictions. Experiment 3.1 examines these predictions and Experiment 2 elaborates upon Experiment 1's findings. I conclude with a discussion of the results.

3.2 Consequential and Intentional Reciprocity

At first glance the meaning of reciprocity seems to be obvious: do unto others as they do unto you. That is, one's own behavior is conditional on others' behavior. While it is clear that conditional means rewarding good actions (and punishing bad ones), research often neglects to examine on which aspects of others' behavior specifically actions are conditional on.

A prominent version of reciprocity is the Tit-For-Tat strategy (Axelrod & Hamilton, 1981; Rapoport & Chammah, 1965). Tit-For-Tat requires players to imitate other players' behavior of the preceding interaction. While research on the prevalence and effectiveness of Tit-For-Tat (for a review see Komorita & Parks, 1999) has not explicitly specified the role of intentions, Tit-For-Tat has usually been formulated in a consequential way. That is, players were assumed to reciprocate the observed behaviors of other players, not taking into account the circumstances (e.g. the possibility to cooperate) under which others acted.

Different from the Tit-For-Tat strategy, reciprocity-based theories of social preferences suggest that individuals evaluate and reciprocate outcomes based on others' kindness, which is assumed to signal players' intentions (e.g. Dufwenberg & Kirchsteiger, 2004; Falk & Fischbacher, 2000; Rabin, 1993). Specifically, player A judges player B as kind if she thinks that B acted to improve A's payoff, given B's belief about A's action. Hence, the evaluation of kindness depends, among other things, on the "third order belief" of player A about B's belief about A's action. It is important to note that this approach considers games in which an actor who wants to be kind can actually act kindly. For instance, Fischbacher, Fehr, and Gächter (2001) have argued that people reciprocate intentions in public goods games. However, in their experiments positive intentions always coincided with high contributions, so a final evaluation of the role of intentions seems not feasible based on their results.

Experimental findings suggest an important role of intentions for cooperative behavior in dyadic interactions. For instance, Blount (1995) reported that in the ultimatum game—where one player makes a proposal for how to split an amount that is then accepted or rejected by the second player—individuals accepted smaller offers when these came from computers, compared to offers from human participants. McCabe, Rigdon, and Smith (Kevin A. McCabe et al., 2003) reported that trustees in a trust game returned a greater amount of money if the truster could actually decide to trust or not, compared to a condition where the truster was forced to "trust" the trustee.[14]

[14] Additional evidence for the relevance of perceived intentions for the evaluation of behavior comes from diverse fields, such as judicial decision making and child development. For instance, Howe and Loftus (1992) reported that half of their participants used an "intention-only" rule when determining levels of blame in hypothetical court cases, and Behne at al. (2005) showed that even 9-month-old children reacted differently when they did not receive a toy because an adult did not give it to them deliberately, compared to when the adult could not give the toy.

How can individuals evaluate intentions in public goods games? I assume that similar to reciprocity-based theories of social preferences, intentions are inferred based on a comparison of individuals' behavior with the individuals' behavioral options. However, my approach differs in two ways from reciprocity-based theories of social preferences. First, I suggest that individuals do not form higher-order beliefs. This limitation of players' sophistication is in line with Colman's (2003a) findings that players usually do only consider others' goals and do not form or use higher-order beliefs, such as what others' might believe about one's own goals. Second, I suggest that the kindness of an action is evaluated independent of its consequence. Therefore, the behavior of an individual is evaluated cooperatively if the individual intended to behave cooperatively, even if the actual behavior is, for instance, a low contribution.

For public goods games, I suggest that a player view the others' past relative contributions—that is, a player's absolute contribution divided by her maximal possible contribution—as signaling others' intentions. High relative contributions signal positive intentions and low relative contributions signal negative intentions. I consider individuals who react positively (negatively) to others' positive (negative) intentions—independent of absolute contributions and outcomes—as *intentional reciprocators*. Individuals who react positively (negatively) to others' high (low) absolute contributions—independent of relative contributions—are considered *consequential reciprocators*. Consequential and intentional reciprocators imitate different aspects of others' behavior: the former imitate consequences and the latter imitate intentions.

The definition of *intentional reciprocators* based on relative contributions rests on the assumption that people derive intentions from others' past behavior and the external constraints that limit behavioral options. This assumptions seems reasonable, as research on social attribution and theory of mind has shown that people distinguish between internal and external causes of others' behavior (Malle, 1999). The assumption is also justified from a functionalistic point of view, because—in interdependent decision making—if individuals have beliefs about others' actions, they can often make better decisions. Experimental evidence shows that individuals form beliefs about intentions and use them to make decisions. Kelley and Stahelski (1970b), for instance, showed that players can infer intentions from other players' past decisions, and Andreoni and Miller (1993) demonstrated that players use others' past behavior to update beliefs about others' future actions. McClintock and Liebrand (1988) showed that cooperative choices are strongly influenced by expectations about others' behavior.

The assumption that players form beliefs about others' intentions is closely linked to the assumption that people have a theory of mind (e.g. Baron-Cohen, 1994; Leslie, Friedman, &

German, 2004). Accordingly, it was examined how players use information in games to infer other players' goals—represented by beliefs about others' intentions—and how players use information to update the degree of sophistication attributed to other players. Hedden and Zangh (2002) found that players updated their beliefs about the other players' level of reasoning based on observed decisions. More specifically, they found that participants by default assumed the other player to be myopically self-interested; they adjusted their initial belief if the other's choices indicated that he considered their motives (see also Colman, 2003b; Stahl & Wilson, 1995; Zhang & Hedden, 2003). Examining the role of intentions, McCabe, Smith, and LePore (2000) found that players behaved differently when decisions in a game were made sequentially, compared to simultaneously. The authors interpreted this as evidence for the assumption that players use others' earlier decisions to update their beliefs about others' intentions. Supporting this result, McCabe et al. (2000) also found that cooperation increased for repeated interaction with the same partner, because repeated interaction further facilitated the detection of cooperative intentions.

A simple way to disentangle behavioral consequences and intentions in a public goods situation with voluntary contributions is to provide individuals with variable endowments, which can vary between players making simultaneous contribution decisions or vary for a player making repeated decisions. With variable endowments, the same absolute contribution can differ in *relative* terms, allowing players to infer intentions. For instance, player A who is endowed with 10 euros and contributes 3 euros signals a negative intention—that is, he is not willing to cooperate, whereas player B who is endowed with 3 euros and contributes 3 euros signals a positive intention—that is, she is willing to cooperate. In contrast, when evaluating the other's behavior merely on the basis of the behavioral consequences, that is, on *absolute* contributions, the behavior will be evaluated identically. Following the two distinct definitions of reciprocity, an intentional reciprocator contributes less after observing the other's low relative contributions (e.g. 3 out of 10), compared to observing the other's high relative contributions (e.g. 3 out of 3). In contrast, a consequential reciprocator contributes the same amount in both cases, because the absolute contributions are equal.

Public goods games with unequal resources have been examined before (e.g. Marwell & Ames, 1979; Rapoport, Bornstein, & Erev, 1989; van Dijk & Wilke, 2000), but these experiments did not allow us to contrast consequential and intentional reciprocity, because contribution decisions were dichotomous or made simultaneously and the games were not repeated. Nevertheless this research did examine people's preferences for decision rules in one-shot public goods games. Specifically, van Dijk and Grodzka (1992) found that participants informed about asymmetric endowments considered the equal absolute contribution rule less fair than participants not informed about asymmetry. Van Dijk and Wilke (2000) found that participants in a public goods game were generally best described as using the equal proportional contribution rule—which is similar to the intentional reciprocity rule, with the exception that intentional reciprocity is conditional on others' past behavior while the proportional rule is conditional on others' simultaneous contributions—and that 98% of these participants are better described by that rule than by the equal final outcome rule. Note that van Dijk and Wilke did not test for the equal absolute contribution rule and that their classifications were based on fairness ratings and not on actual contributions. Hence, this research extends previous research on decision rules in public goods games with asymmetric endowments in that it examines a rule that prescribes equal absolute contributions, examines rules in repeated games, and classifies participants based on actual contribution decisions.

3.2.1 Models of Reciprocity in Public Goods

I will examine reciprocal behavior in a public goods game where four players *i* have an endowment, *E*, which they can allocate to a private account or to a public good. The contribution to the public good is *c* and the investment in the private account is $E - c$. Investments in the private account lead to a payoff equal to the investment. Contributions to the public good are multiplied by a constant and then equally split among all members of a group. The efficiency gain of a contribution to the public good is expressed as the 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 $\pi_i = E_i - c_i + MPCR \cdot \sum_{i=1}^{N} c_i$. A dilemma occurs when $1/N <$ MPCR < 1, because for $1/N <$ MPCR, contributions lead to efficiency gains, while for MPCR < 1 players have an incentive to defect.

The consequential and intentional reciprocity models I suggest are modified versions of the reciprocity model that was tested in Chapter 2. The key assumption of these models is that people will contribute the same amount as others did in the round before, and that this reciprocity principle is mediated by people's generosity. A probabilistic decision rule captures random deviations from strict reciprocity.

3.2.2 Consequential Reciprocity

In general, the consequential reciprocity model describes the following decision process. In the first round of a game, players have no information about others' cooperativeness and are thus assumed to allocate their endowment equally to the public good and to the private account. In the following rounds players contribute in absolute terms as much to the public good as the others did in the preceding round. If the endowment is not sufficient to reciprocate others' contributions, players contribute their complete endowment in order to reciprocate as much as possible. Reciprocity is assumed to be mediated by individuals' tendency to give more or less than a strict reciprocity principle predicts. For instance, if others in a public goods game contribute 5 units of their endowment, then a generous individual might contribute 6 units. Figure 3.1 depicts a flow chart describing consequential reciprocity.

Formally, the model determines first a *most likely contribution* and then the probability of all possible contributions. The most likely contribution, m, in the first round is $c_m = E / 2$. The probability p_j that one of the possible contributions j is chosen is defined as:

$$p_j = \exp\left(-\left|x_{jm}\right| / 2o_C^2\right) / U \ , \qquad (3.1)$$

where x_{jm} is the difference between any possible contribution j and the most likely contribution m, the free parameter σ_C is a standard deviation, defining to what extent similar contributions to the most likely contribution are also chosen with a substantial probability, and U is a constant that normalizes the likelihoods so that they sum to 1. According to Equation 3.1, the probability of choosing an allocation increases with its similarity to the most likely allocation m. This probabilistic decision rule captures that other non-reciprocal influences lead to small, random deviations from strict reciprocity.

In the second and all following rounds contributions depend on the other players' contributions in the preceding round. To implement the consequential reciprocity principle, the most likely allocation is determined by contributing in absolute terms as much as the other players did in the preceding round,

$$c_m = \begin{cases} E \geq Mdn(o) & \rightarrow \quad Mdn(o) + \gamma \cdot [E - Mdn(o)] \\ E < Mdn(o) & \rightarrow \quad E \end{cases} , \qquad (3.2)$$

with $Mdn(o)$ as the median of other players' contributions in the preceding round and $\gamma \in [-1,1]$ as a generosity parameter determining how much a player contributes more than the others' median contribution in the previous round. The choice probabilities for all possible contributions are again determined according to Equation 3.1.

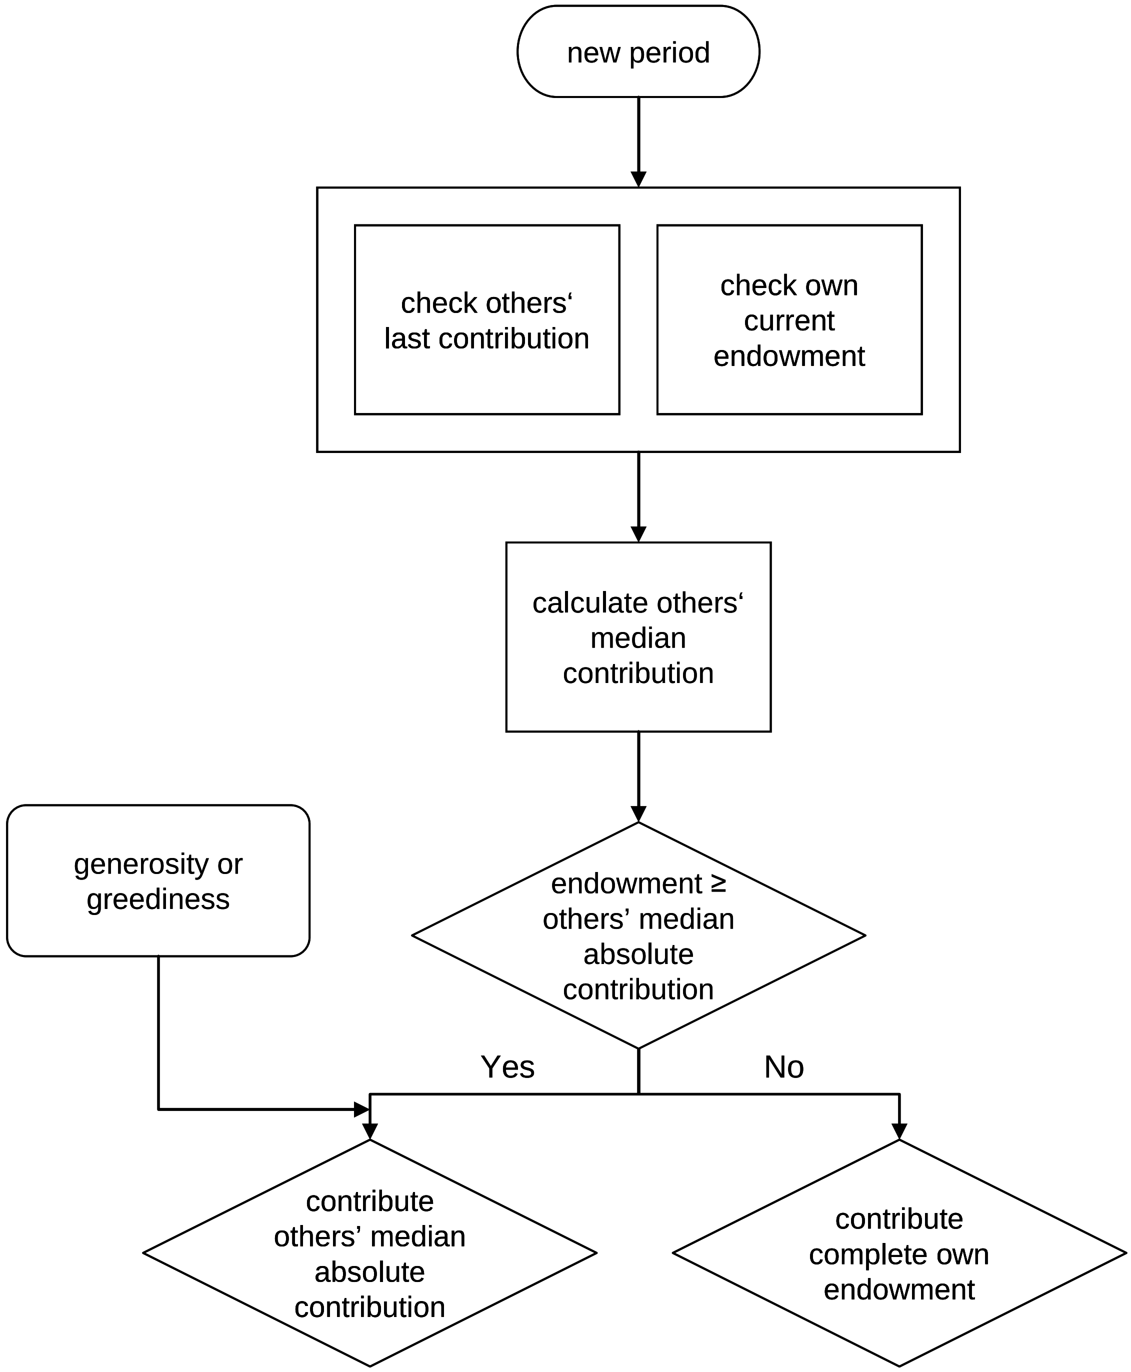

Figure 3.1. Flow chart for consequential reciprocity. Differently than in Chapter 2, the reciprocity heuristic must now consider others' and the own endowment before making a contribution decision.

3.2.3 Intentional Reciprocity

To represent intentional reciprocity I modify the consequential reciprocity model and assume that the players' relative contributions signal their intention. According to this assumption, a player who reciprocates others' intentions should make the same *relative* contribution as the other players. Therefore Equation 3.2 is modified as follows:

$$c_m = E_i \cdot [Mdn(o/E) + \gamma \cdot (1 - Mdn(o/E))] \tag{3.3}$$

with *Mdn*(*o/E*) as the median of the other players' relative contributions. Choice probabilities are determined according to Equation 3.1. Figure 3.2 depicts a flow chart describing the intentional reciprocity heuristic.

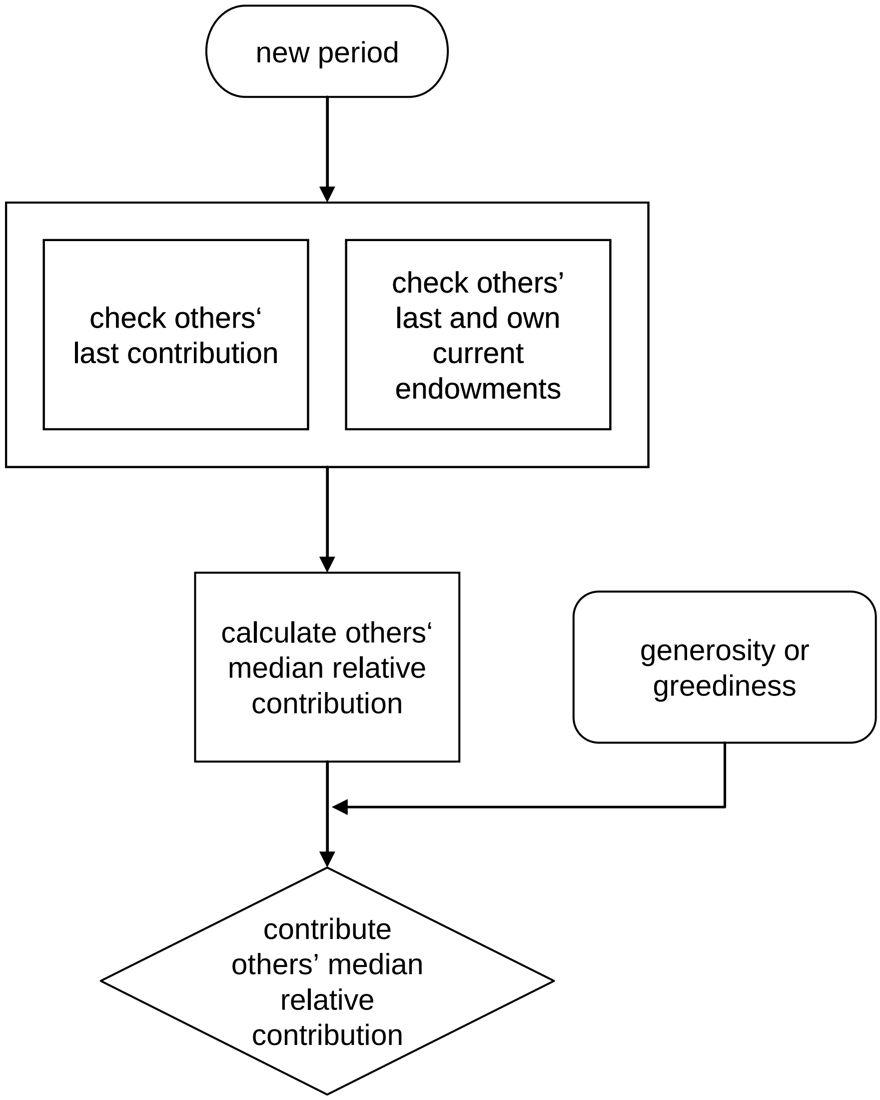

Figure 3.2. Flow chart for intentional reciprocity. The key assumption of the intentional reciprocity heuristic is that relative contributions signal intentions.

3.2.4 <u>The Ecological Rationality of Consequential and Intentional Reciprocity</u>

The investigation of the ecological rationality of decision rules inquires if decision rules are successful in specific decision environments (Gigerenzer et al., 1999; Goldstein &

Gigerenzer, 2002). In my case the decision environment consists of the public goods game defined above and of the distribution of endowments over participants and over time. The analysis of the ecological rationality differs from standard rationality criteria in that it is not concerned with the adherence to norms of rationality (e.g. weakly ordered preferences, consistent beliefs, and correct application of the rules of logic) and it does not look only at the instrumentality of a behavior, but also at how a behavior's instrumentality depends on characteristics of the decision environment.

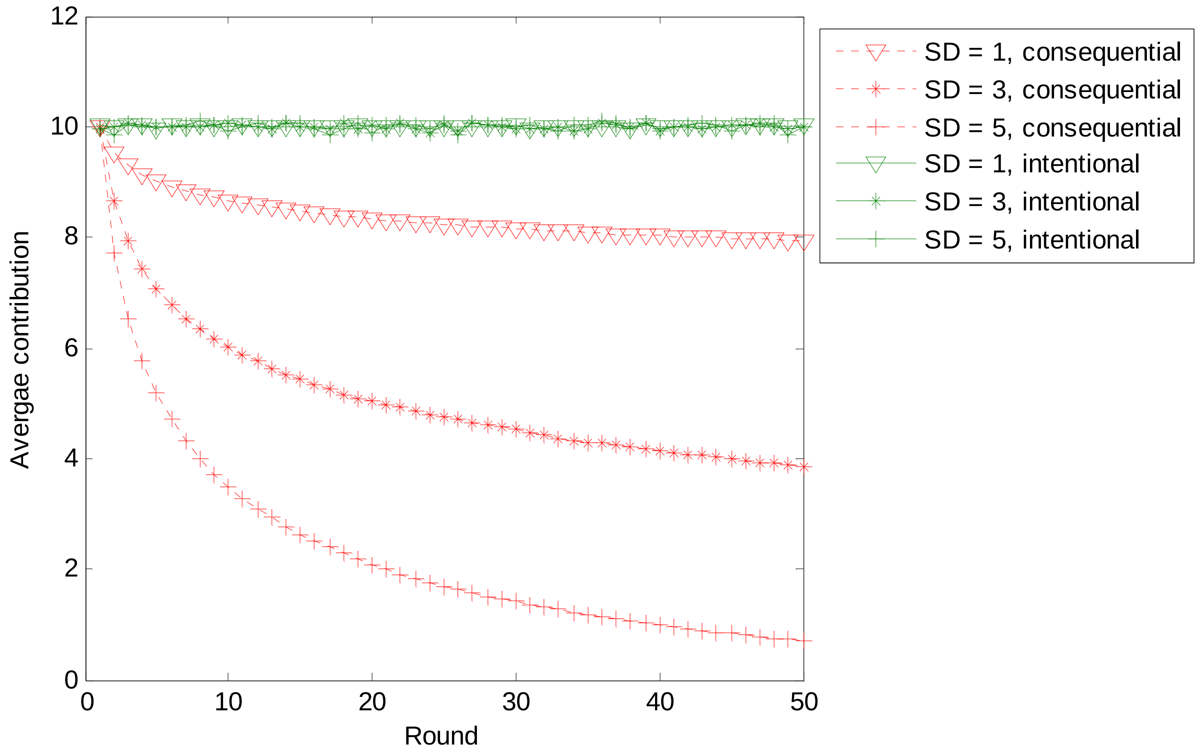

Figure 3.3. Development of contributions for consequential and intentional reciprocators in a public goods game with variable endowments. Results are means over 10,000 simulated four-person public goods games. For every player the mean endowment was 10 points.

An individual who reciprocates consequences only needs to observe the consequences of others' behavior. An individual who reciprocates intentions additionally needs to observe external constraints under which others' made their decisions and then to integrate both pieces of information. Thus, if consequential reciprocity is the easier behavioral rule, why should individuals reciprocate intentions? Intentional reciprocity has the advantage that it can maintain cooperation between individuals whose endowments—that is, their opportunity to cooperate—vary over time. For instance, if player A has few resources available in round *t* and therefore contributes little, this will lead to a lower perceived average absolute contribution of others, so that consequential reciprocity will lead to reduced contributions in round $t + 1$. As a result, player A will perceive a low average absolute contribution in round $t + 2$ and contribute little, even if his endowment has increased again. Figure 3.3 shows how contributions decline when consequential reciprocators interact in public goods games with varying endowments across players. In contrast, as Figure 3.3 also shows, intentional reciprocity maintains mutual cooperation when players have varying endowments.

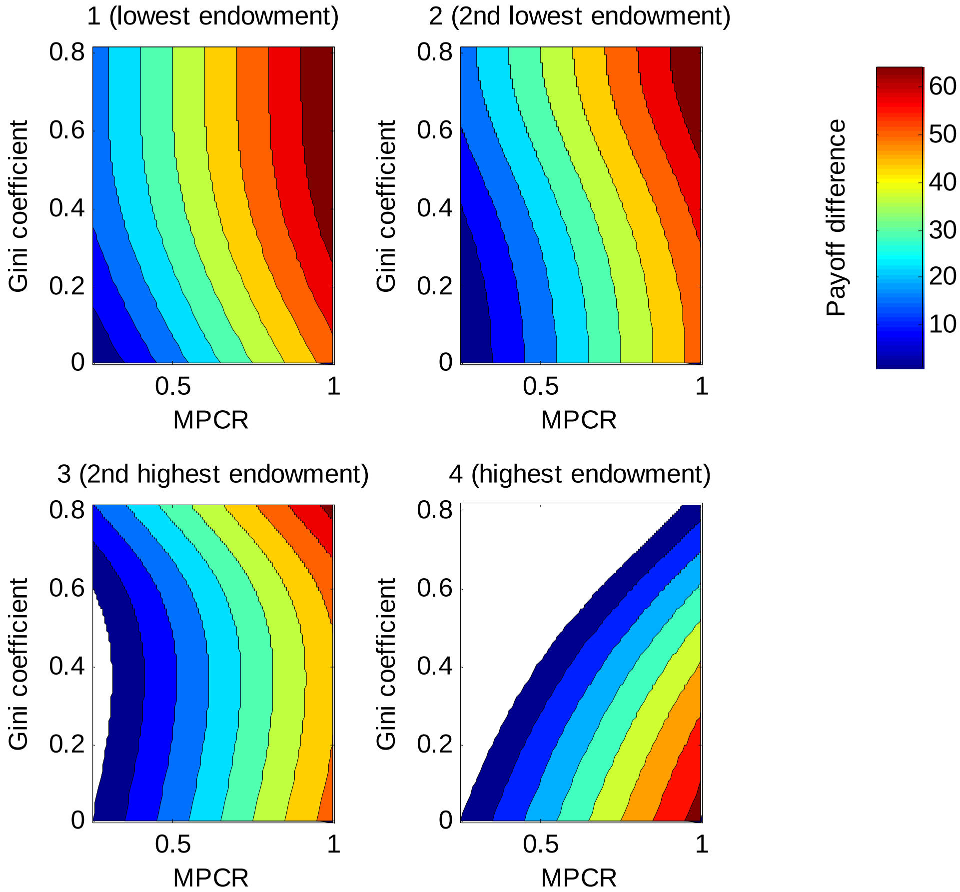

Figure 3.4. The influence of efficiency gain and inequality on players' payoffs. The plots show payoffs for players in a public goods game when all cooperate minus payoffs when nobody cooperates. White areas mark combinations of MPCR and Gini coefficient a player is worse off when he or she contributes his or her complete endowment.

In case some players have on average higher endowments than other players, intentional reciprocity predicts that the wealthier players will on average make larger absolute contributions. Despite this inequality in contributions, the wealthier player will often still be better off compared to when no one contributed. However, in the extreme it can happen that for a wealthier player it might be beneficial to contribute nothing, even if all other players would continuously contribute their whole endowment to the public good if the wealthier player also contributed. To explore this possibility I systematically examined under which combinations of inequality, measured with the Gini coefficient[15] and MPCRs, a player could only lose by contributing to a public good. Figure 3.4 depicts the payoff difference between two possible situations: Full contribution of all players minus zero contribution of all players. The lighter the gray, the higher are payoffs under full contribution compared to zero contributions. The areas with a diagonal pattern mark combinations of inequality and MPCRs, where players can unilaterally guarantee themselves a higher payoff than can be achieved through full cooperation of all players. The figure shows that in a four-person game this can only occur for the two wealthiest players and is mostly restricted to environments with an unusually high inequality (Gini coefficients for the distribution of wealth in developed countries are usually around .35). In most environments, full cooperation of all agents is beneficial to all players. Since only intentional reciprocity maintains cooperation over time and thereby leads (in most environments) to higher payoffs for all players, it can be predicted that intentional reciprocity is more adaptive in social dilemmas than consequential reciprocity; therefore it should represent a better cognitive model of people's behavior in public goods situations.

3.3 Predictions

To test the two models of reciprocity against each other I will determine which model is more suitable to predict participants' contributions round by round in a repeated public goods game with variable endowments. Moreover, the two reciprocity models make several qualitative predictions about participants' contributions that I will test. In addition, since the two reciprocity concepts make different predictions about the information that an individual requires to make a decision, these predictions will be tested by monitoring participants' information search using a computerized information board (Payne et al., 1993; van Dijk & Grodzka, 1992).

According to intentional (consequential) reciprocity, participants will make higher (lower) contributions when others' relative contributions are high in relative terms and at the same time low in absolute terms compared with when others' contributions are high in absolute terms and low in relative terms. Figure 3.7 displays the predictions of both models. To derive predictions, I computed the contributions from the reciprocity models in the conditions with the generosity parameter set to 0. The models make different prediction for cases of asymmetric endowments.

[15] The Gini coefficient varies between zero, indicating equal distribution of resources, and one, indicating maximally unequal distribution (see Atkinson & Bourguignon, 2001; Haitovsky, 2001).

Intentional (consequential) reciprocity predicts that individuals contribute the same relative (absolute) amounts and wealthier players contribute higher absolute (lower relative) amounts.

According to consequential reciprocity individuals will only search for information about others' contributions and will neglect information about others' endowments. According to intentional reciprocity individuals will search for information about others' contributions *and* for information about others' endowments.

It has been shown that not all individuals can be considered as reciprocators; instead, a substantial proportion of individuals defect or cooperate unconditionally (e.g. Kurzban & Houser, 2005). Therefore when testing the two models of reciprocity against each other, the test has to be restricted to individuals who in general behave reciprocally.

3.4 Experiment 3.1

3.4.1 Method

Participants played repeated four-person public goods games. The marginal per-capita return was .375, that is, every point allocated to the public good was multiplied by 1.5 and then equally divided among the group members. Participants' information acquisition was recorded using a computerized information board (Figure 3.5). Information acquisition was costly but relatively cheap, as the aim of imposing costs was to limit search to relevant information. The experiment was conducted with the software CING (Czienskowski, 2004).

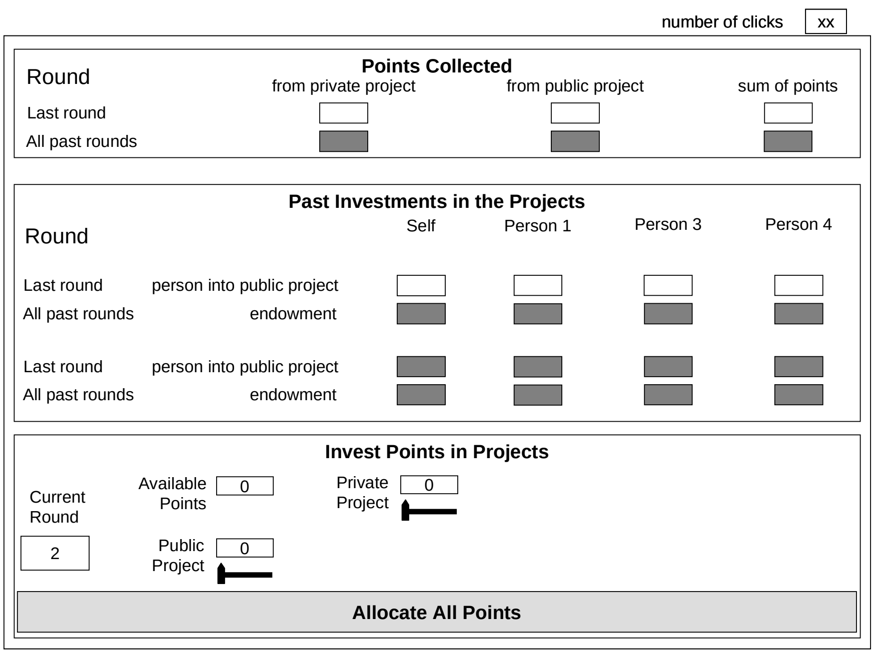

Figure 3.5. Schematic picture of participants' information board (for person #2). Participants had to click and pay to obtain information behind the dark boxes.

3.4.1.1 Design

The first, *random-reduction* condition was designed to test how participants respond to low contributions of other players; I especially wanted to examine if it mattered to players if low contributions were made deliberately or because of low endowments. In most rounds the participants received 12 endowment points; beginning with round 15 in every third round one participant received either 2, 3, or 4 points. Low endowments were randomly distributed across the last 80 rounds of the game. This endowment plan enabled participants—if they intended to cooperate—to first establish high contribution levels regardless of the type of reciprocity principle they followed and forced them to make low contributions from round 15 on.

Table 3.1 Distributions underlying the endowment schedules (frequency of endowments).

Condition	Endowment points																					
	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
Random-reduction	4	3	3								85											
Random-equal						12	12	12	6	12	6	12	12	12								
Random-unequal (wealthy)										6	6	12	6	6	12	6	12	6	6	6	6	6
Random-unequal (poor)	12	12	12	12	18	11	18															

The distribution of endowments in the random-reduction condition was still very transparent for the participants, so that after some experience with the game, they could infer how the endowments were distributed across the participants, providing low incentives to search for information about endowments as predicted by the intentional reciprocity model. Therefore in the second, *random-equal* condition, endowments were randomly drawn from a discrete distribution with a minimum of 7 and a maximum of 15. The rows in Table 3.1 show the frequencies of the possible contributions in the three conditions. The random-equal condition also allowed us to test if contributions were made conditional on others' relative or absolute contributions. The third, *random-unequal* condition tested if relative or absolute contributions differed between wealthy and poor players. Endowments of the two wealthy (poor) players were drawn from a distribution with a mean of 16.75 (5.25), a maximum of 23 (8), and a minimum of

11 (2); see also Table 3.1. Endowments in the last two conditions were created so that within each block of sixteen rounds all players received their mean endowment. The sum of all players' endowments for a game was the same in all conditions.

3.4.1.2 Participants and Procedure

The 72 participants (34 women) with an average age of 24 years were randomly assigned to one of the three conditions. Participants were mainly students from different departments of the Free University Berlin. Each condition was conducted in four sessions with groups of 8 participants. When participants arrived at the laboratory they were seated in four rooms, each with two separate cubicles, preventing any communication. Participants were not informed who the other participants in their group were and participants in the same room always belonged to different groups.

The instructions explained to the participants that they would take part in a repeated group decision-making task together with three other persons. Then the instructions explained the public goods game, first in text form and then with some numerical examples. It was further explained that for every collected point, 0.01 euros would later be paid to the participants, that every click for information cost 0.1 cents, and that the game would be repeated but that no participant knew how often. Participants made 95 decisions in the random-reduction and 96 in the random-equal and random-unequal conditions. 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 a particular goal. To inform participants about the variability of endowments, participants in the random-reduction condition were instructed that everyone would generally receive 12 points per round, but that occasionally one participant would receive a lower endowment. Participants in the random-equal and random-unequal condition were instructed that endowments could vary for the same person across rounds and between persons within each round.

A quiz tested participants' understanding of the game. If anybody failed to answer a question, the experimenter clarified any misunderstanding, so that all questions could be answered correctly. Finally participants received instructions for the computerized information board used to conduct the game. The game started after all participants had correctly completed a second test quiz about the interface and had received a summary of the instructions.

In the first round of a game, participants received their endowments and made their contribution decisions. After all participants had made contributions in a round, the next round started automatically. From the second round on participants could access information on the computerized information board (see Figure 3.5), by clicking on information boxes. Participants could access information about their collected points, about others' contributions to the public projects, about others' endowments, and about their own contributions; all information was available for the last round and cumulated for all past rounds. Opening a box automatically closed the previously opened box. Participants were allowed to click for information as often as they wanted. After a decision was made, no information could be acquired until the next round started. Participants could allocate their endowment at any time in a round.

Participants received their payments individually. Sessions lasted on average 120 minutes and participants earned on average about 24 euros.

3.4.2 Results

3.4.2.1 Contributions

Figure 3.6 depicts participants' contributions in the three conditions. In the random-reduction condition participants contributed on average 4.71 points (*SD* = 4.81), or 43% (*SD* = 41%) of their endowment. In the random-equal condition participants on average contributed 3.51 points (*SD* = 4.01), or 33% (*SD* = 36%). In the random-unequal condition participants on average contributed 3.15 points (*SD* = 3.81), or 32% (*SD* = 33%). Mean contributions in the first 10 and last 10 rounds were 56% (*SD* = 23%) and 37% (*SD* = 24%) in random-reduction, 41% (*SD* = 12%) and 32% (*SD* = 27%) in random-equal, and 40% (*SD* = 10%) and 24% (*SD* = 17%) in random-unequal conditions, respectively. A repeated measurement analysis of variance (ANOVA) with conditions as a between-subjects factor and block—the average contribution in the first 10 rounds versus the last 10 rounds—as a within-subject factor tested for differences in average contributions. While conditions did not influence contributions, $F(2, 15) = 2.27$, $p = .137$, the factor time had an influence on participants' contributions; $F(1, 15) = 8.46$, $p = .011$, $\eta 2 = .36$, indicating that contributions decreased with time.

3.4.2.2 Classification into Cooperative Types

Similarly to Kurzban and Houser (2005), I classified participants based on a linear regression analysis that predicts participants' contributions based on other players' contributions in the preceding round. For each participant I computed one regression analysis with absolute and one with relative contributions. Specifically, a player's contributions in round t, c_t, was regressed on o_t-1, the median of the other players' contribution in round t-1, according to the function $c_t = g + \beta \cdot o_t\text{-}1$; with g as the intercept representing generosity, and β as the slope representing reciprocity. Each participant with a standardized $\beta > .25$ in at least one of the two regression models was classified as a reciprocator. Of the non-reciprocators, those with an intercept < 1/3 were classified as free riders, with stable small contributions, and those with an intercept > 1/3 were classified as unconditional cooperators with stable high contributions. Altogether 35 (49%) participants (13, 14, and 8 participants in the random-reduction, equal, and unequal conditions, respectively) were classified as reciprocators, 24 (33%) as free riders, and 8 (11%) as unconditional cooperators; five participants (7%) could not be classified.

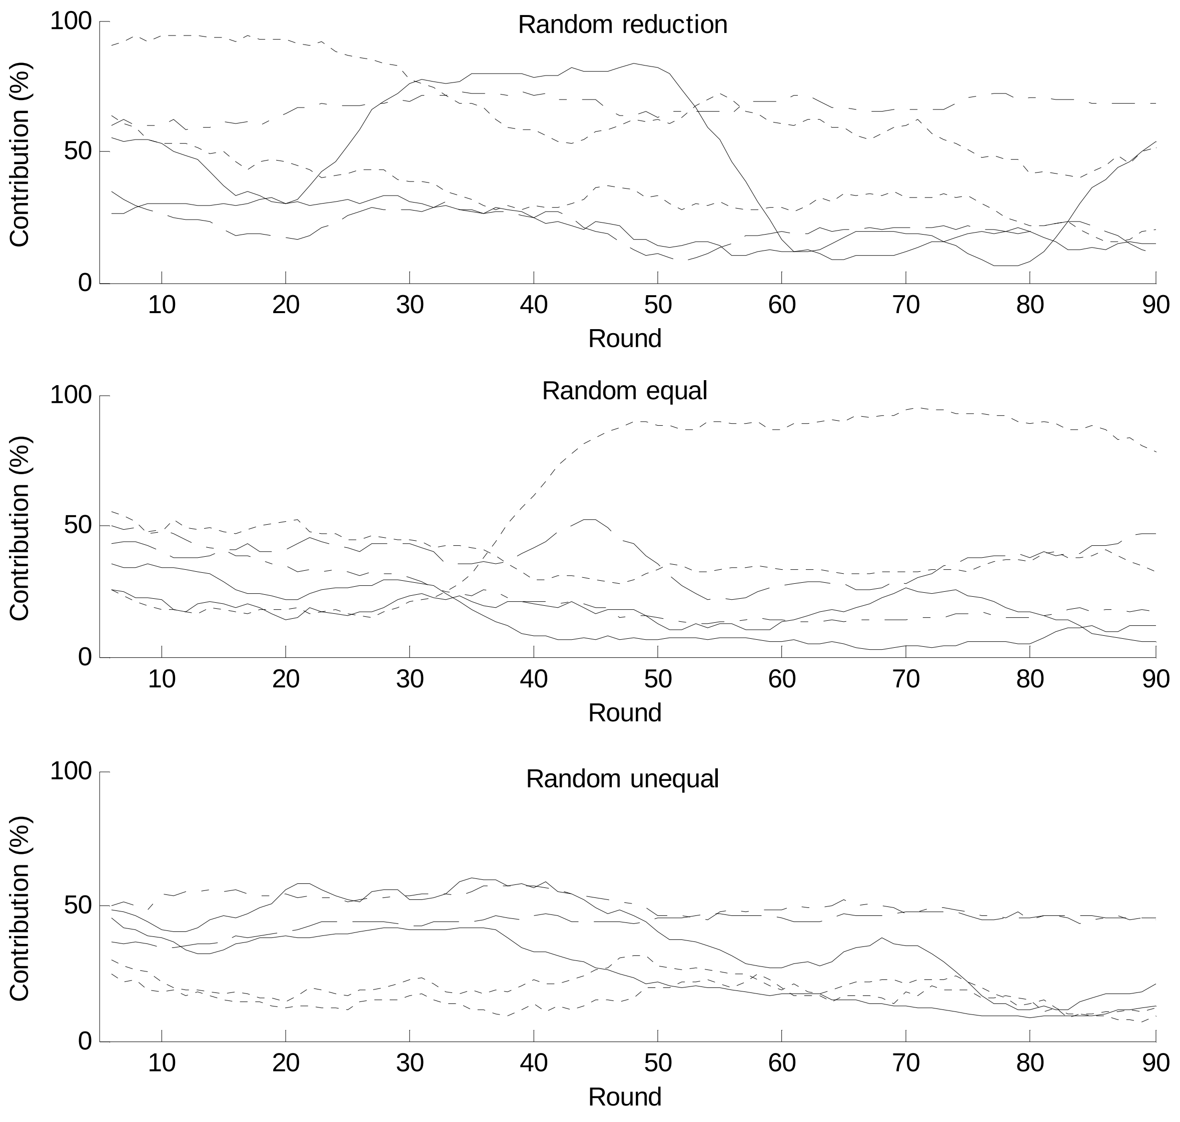

Figure 3.6. Average contributions of groups. Panels show contributions in the different conditions. Contributions are running averages with a window size of 11 rounds. Different lines represent different groups in a condition.

3.4.2.3 Model Fits

The comparison of model performance was restricted to the participants classified as reciprocators. Both models predict for each round the probability with which any possible contribution is made, 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 (Pitt & Myung, 2002), parameters were optimized by using only participants' decisions in the first 47 rounds of the games and the models' predictions were then cross-validated for the remaining rounds. As a goodness-of-fit measure the sum of squared errors (see e.g. Selten, 1998) was employed, by calculating for every possible contribution in a round the sum of the squared differences of the predicted probability and the observed decisions (assigning the chosen contribution a value of 1 and all other contributions a value of 0). The sum of squared errors lies between 0 and 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 participant, first a grid search was performed to identify good starting parameter values. These values were then optimized with the sequential quadratic programming method (Fletcher & Powell, 1963), as implemented in MATLAB®. The parameters of the models were restricted to $-1 \leq \gamma \leq 1$ for the generosity parameter, and to $.1 \leq \sigma c \leq 16$ for the standard deviation.

Table 3.2 *Means and standard deviations of model fits (from cross-validation) and parameter values for participants classified as reciprocators.*

Model	Parameter/Fit	Condition		
		Random reduction *n* = 11	Random-equal *n* = 12	Random-unequal *n* = 14
Consequential reciprocity	Fit	.95 (.16)	.81 (.09)	.81 (.05)
	Generosity	.12 (.83)	-.29 (.47)	.04 (.5)
	SD	3.43 (3.43)	3.91 (1.19)	3.68 (2.05)
Intentional reciprocity	Fit	.93 (.20)	.8 (.1)	.89 (.28)
	Generosity	.1 (.87)	-.3 (.45)	-.35 (.64)
	SD	2.43 (1.98)	3.36 (1.16)	6.74 (6.19)

Note. Numbers is parentheses are standard deviations.

The consequential reciprocity model predicted participants' contributions in the second half of the experiment with an average fit of .86 (*SD* = .14) compared with the intentional reciprocity model with an average fit of .87 (*SD* = .19). Table 3.2 shows the average fits separately for the different conditions. To test whether one of the two models was more suitable to predict participants' contributions, I performed a repeated measurement ANOVA with the type of model (intentional vs. consequential) as a within-subject factor and the three conditions as a between-subjects factor. The average fit of the two models did not differ, $F(1, 32) = .75$, $p = .39$, nor did the fit differ between the three conditions, $F(2, 32) = 2.45$, $p = .102$, indicating that intentional and consequential reciprocity predict reciprocators' behavior across all conditions equally well.

3.4.2.4 Contribution Predictions

The main contribution prediction is that according to intentional reciprocity participants' contributions depend only on others' relative contributions, whereas according to consequential reciprocity they depend only on others' absolute contributions.

To test this prediction I first examined whether participants cooperate more when others' low contributions were made from low endowments (i.e. contributions were high in relative terms), compared to when they had high endowments (i.e. contributions were low in relative terms). Specifically, I split participants' contributions into those (a) following a round in which one of the other players received a low endowment and (b) following a round in which no participant had a low endowment. Then I computed for each participant two contribution vectors, with their contributions conditional on others' median contribution. The difference score Δ_C is the mean of the difference vector between contributions after (a) forced versus (b) deliberate low contribution. The difference score Δ_C is larger (smaller) than zero if participants contribute more (less) after low contributions were forced. For the 14 reciprocators in the random reduction, the mean Δ_C was .02 ($SD = 1.21$). A *t*-test for paired comparisons showed no difference in contributions after forced compared to after deliberate low contributions, $t(12) = -.06$ (one tailed), $p = .52$. Hence, contrary to the prediction from intentional reciprocity, participants did not contribute more after others were forced to make low contributions compared to when they did so deliberately.

Table 3.3. *Players' relative and absolute contributions given others' relative and absolute contributions in the previous trial.*

Others' absolute contributions	Others' relative contributions	
	Low	High
Low	17% (12%) 1.72 (1.19)	30% (21%) 3 (2.03)
High	32% (35%) 3.6 (4.35)	40% (26%) 4.44 (3.08)

Note. The first row in a cell reports means of players' relative contributions; the second row reports absolute contributions. Standard deviations are in parentheses.

As a second test of the contribution prediction I examined whether participants' contributions varied with others' relative contributions or with others' absolute contributions. Only the random-equal and random-unequal conditions provide the necessary combinations of relative and absolute contributions for this test. I classified others' relative and absolute median contributions in the preceding round as either low or high, using the mean of others' median contributions (relative or absolute) a player experienced in a game as a cutoff point. Table 3.3 and Figure 3.7 depict participants' contributions conditional on the four past contribution conditions created by the classification.

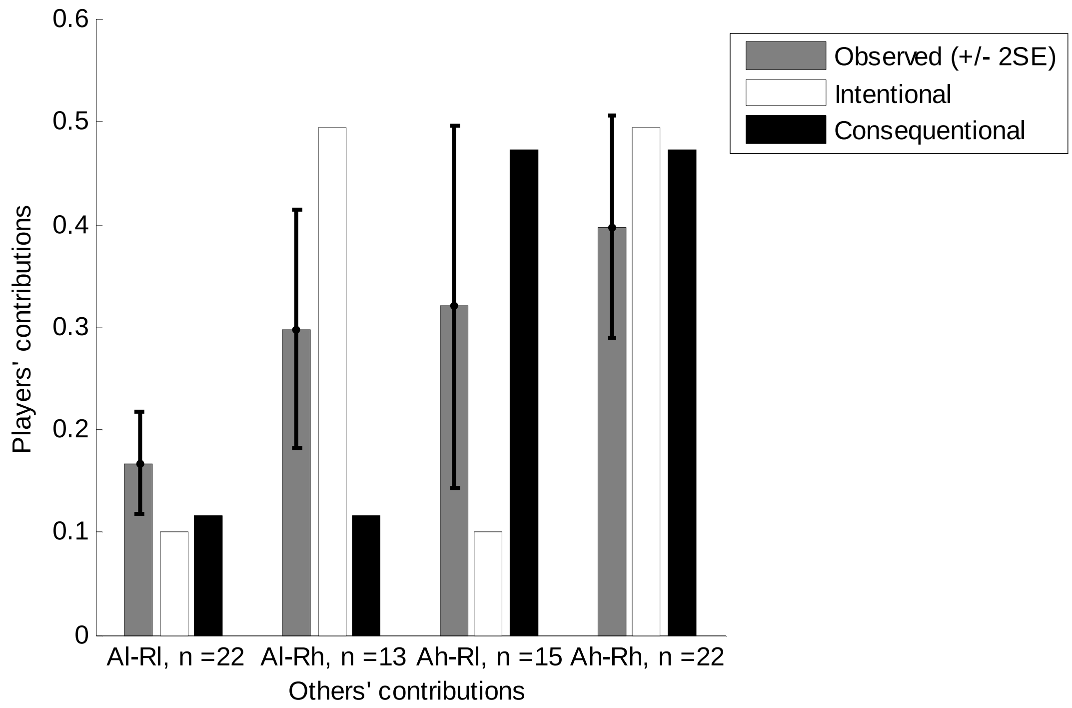

Figure 3.7. Participants' contributions conditional on others' absolute low and high (Al and Ah, respectively) and others' relative low and high (Rl and Rh, respectively) contributions in the preceding round, together with contributions predicted from intentional and consequential reciprocity. Predictions were computed according to the models and given other participants' contributions in the preceding round and averaged over participants.

The second test of the contribution prediction examined if, as predicted by intentional (consequential) reciprocity, contributions after others' absolute low and relative high contributions are higher (lower) than after others' absolute high and relative low contributions. Reciprocators' average relative (absolute) contributions of 30% (3 points) after others' high relative and low absolute contributions were similar to their average relative (absolute) contributions of 32% (3.6 points) after others' low relative and high absolute contributions (see also Table 3.3 and Figure 3.7). Hence, instead of being different, as predicted—with varying direction—by the two reciprocity models, participants contributed the same amount when others' preceding contributions were not high or low in absolute *and* relative terms, $t(7) = .86$, $p = .41$.

The third test of the contribution predictions examined if wealthy and poor reciprocators made the same relative or absolute contributions. Only eight participants were identified as reciprocators in the random-unequal condition. The five wealthy reciprocators contributed on average 2.6 points or 16% of their endowment per round (individual values were .8, 4.2, 5, 1.5 and 1.5 points and 5, 27, 31, 9, and 11%) and the three poor reciprocators contributed on average 2.2 points or 42% (individual values 3.1, 1.8, and 1.6 points and 61, 36, and 30%). Although the sample size is too small to conduct any inferential statistic, the raw data indicate that contrary to the prediction from intentional reciprocity poor participants contributed a higher proportion of their endowment. The similar absolute contribution of wealthy and poor reciprocators speaks rather in favor of consequential reciprocity. [16]

3.4.2.5 Information Search Predictions

According to consequential reciprocity participants should only search for information regarding others' contributions, whereas according to intentional reciprocity they should also search for information about others' endowments. Figure 3.8 depicts the frequency with which reciprocators' searched on average for different types of information. Over all the frequency of clicks for information was rather low, which might be due to the costs (.1 cent per click) associated with information search. Across all conditions the median of reciprocators' mean click frequency across all rounds for others' contributions was .74, and .13 for clicks for others' endowments (see Table 3.4 for details on standard deviations and conditions).

In rounds in which reciprocators did not look up others contributions, their median click frequency for others' endowments was 0 whereas it was .59 if participants looked up others' contributions in the same round. Participants not classified as reciprocators clicked .14 ($SD = .4$) times per round for others' contributions and .02 ($SD = .27$) times per round for others' endowments. Because the click frequencies differed between reciprocators and non-reciprocators for others' contributions, $t(72) = -5$, $p < .001$; $d = 1.01$, and for others' endowments, $t(72) = -2.5$, $p = 015$; $d = .58$, I infer that participants' information search was in accordance with their

[16] Adjusting for group differences by subtracting the group means from individual contributions, wealthy contributed on average .53 points or -8% (-.04, .67, 1.84, .08, .1; -12, -12, 4, -11, -9%) and poor contributed -1.2 points or -8% (-.41, -1.72, -1.55; 22, -3, 4%). Different from the non-adjusted values this speaks rather for intentional reciprocity.

general contribution rule. That is, reciprocators who needed information about others' contributions (and endowments) looked up this information, whereas non-reciprocators—whose contributions were independent of others' contributions—did not look up others' contributions or endowment.

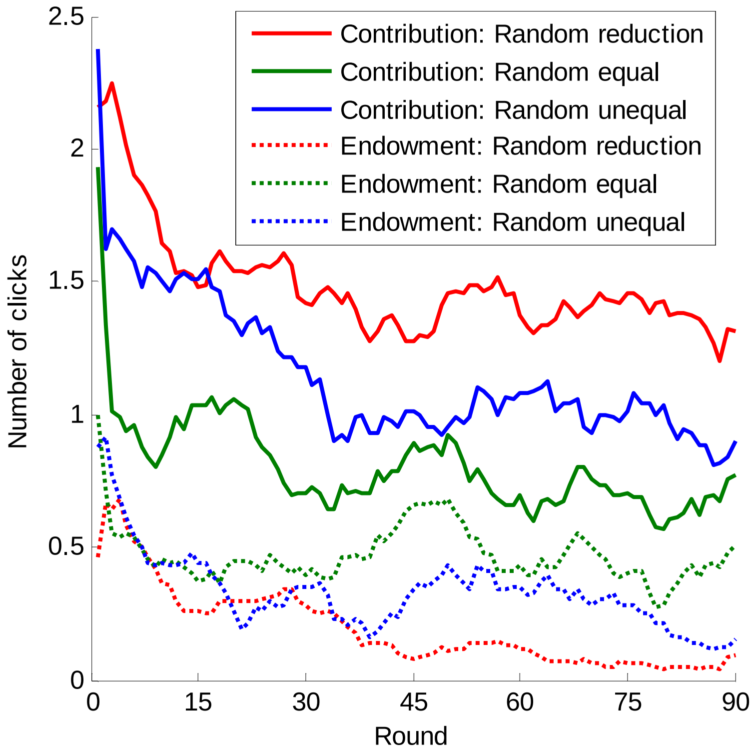

Figure 3.8. Participants' information search. Lines are running averages (window size 11) of clicks per round. Solid lines are for clicks about other's contributions, dotted lines are for clicks about others' endowment.

Table 3.4. *Information search of reciprocators.*

Condition	Contribution	Endowment	Endowment given high contribution	Endowment given low contribution
Random-reduction	1.11 (1.06)	0.07 (.3)	0.18 (.49)	0.24 (.34)
Random-equal	0.61 (.62)	0.26 (.52)	1.27 (.99)	1.13 (.93)
Random-unequal	.92 (1.15)	0.07 (.05)	1 (.76)	0.84 (.62)
All	.74 (.95)	0.13 (.45)	0.57 (.89)	0.54 (.77)

Note. Medians and standard deviations of clicks per round

In sum, the analysis of participants' information acquisition shows that information search was congruent with contribution decisions and supports rather intentional reciprocity because participants also looked up others' endowments, which they only needed to reciprocate intentions. However, information was acquired to a lesser extent than expected.

3.4.3 Discussion of Experiment 3.1

Experiment 3.1 examined which aspect of others' behavior players reciprocate in public goods games. For this test, I compared the fit of the two models and I examined qualitative predictions from the two models. I observed large heterogeneity of individuals' behavior, so that some individuals cooperated or defected unconditionally. However, the largest group of 49% of participants reciprocated. For these reciprocators the two models did not differ substantially in predicting their contributions. Testing the contribution predictions derived from the models, I found more evidence for consequential reciprocity. That the difference in contributions after forced or deliberate low contributions was on average zero supports more consequential reciprocity. However, as reflected by the large standard deviation, the mean difference results partly from large deviations in both directions, making it unreasonable to assume that the mean adequately represents individuals. I further found that (relative and absolute) contributions were similar after others contributed high in relative terms and low in absolute terms, thus not supporting the predictions derived from both reciprocity models. In sum, reciprocators' contributions did not clearly speak in favor of one of the two reciprocity concepts.

Participants' information search supports this interpretation. In general, probably due to information search costs, participants did not search for a lot of information regarding others' contributions and endowments. Nevertheless, the important finding was that participants classified as reciprocators searched for much more information than participants classified as non-reciprocators. In general, participants not only searched for contribution information but also for information about others' endowments, which is in line with the intentional reciprocity concept. In sum, participants' information search corresponded to their contribution decisions, and supported more intentional than consequential reciprocity.

When comparing the fit of the two reciprocity models, participants' contributions, and their information search, I conclude that reciprocators cannot adequately be described by assuming that they reciprocate one specific aspect of others' behavior, consequences, or intentions. Therefore it can be asked whether there might be an alternative reciprocal rule explaining cooperative behavior. One plausible alternative explanation for reciprocal behavior consists of including both aspects, by assuming that individuals consider the consequences *and* the intentions guiding others' behavior. A decision strategy that follows this idea, and is also in line with the main results of Experiment 3.1, is *opportunistic reciprocity*, which predicts that a player contributes his/her complete endowment if others' median contribution is higher than the player's endowment and otherwise reciprocates intentions. The concept can be called opportunistic because it dictates intentional reciprocity only as long as it does not threaten to reduce others' cooperativeness. To describe opportunistic reciprocity formally, Equation 3.2 of the consequential reciprocity model can be replaced by

$$c_l = \begin{cases} E \leq Mdn(o) & \rightarrow \quad E \\ E > Mdn(o) & \rightarrow \quad f(o) \end{cases}, \tag{3.4}$$

where *f*(*o*) equals Equation 3.3, which represents intentional reciprocity. Figure 3.9 depicts a flow chart of consequential reciprocity.

Opportunistic reciprocity predicts that players consider others' relative and/or absolute contributions, depending on how high the current endowment is compared to others' contributions in the last round. The concept provides a possible explanation for why the two reciprocity models predicted participants' behavior in the first experiment equally well. The model can also account for the observation that intentions only mattered when others contributed relatively little (see Figure 3.7), that is, when players had at least as much resources available as the median of others' contributions. Regarding behavior of poor and wealthy players, the model can explain the higher relative contribution rates of poorer participants, because according to opportunistic reciprocity they try to match the wealthier players' contribution.

However, it also predicts—in contrast to my findings—that wealthy players contribute the same relative amount as poor players. Finally, the opportunistic reciprocity model only requires information about others' endowments when others' contributions are low compared to one's own endowment, so that it can explain why the participants in Experiment 3.1 searched only to a moderate extent for endowment information. In sum, the opportunistic reciprocity model can explain several findings of Experiment 3.1, which the two other reciprocity models cannot explain. But it is difficult to test the model with data from Experiment 3.1: First, it predicts different behavior dependent on the relation of a players' endowment and others' contributions, but the level of wealth was only manipulated between and not within players in Experiment 3.1. Second, it predicts intentional reciprocity if the players' endowment is high, but intentional reciprocity was difficult to test in Experiment 3.1 because cases in which others' endowments were high (low) in absolute terms while being low (high) in relative terms were infrequent.

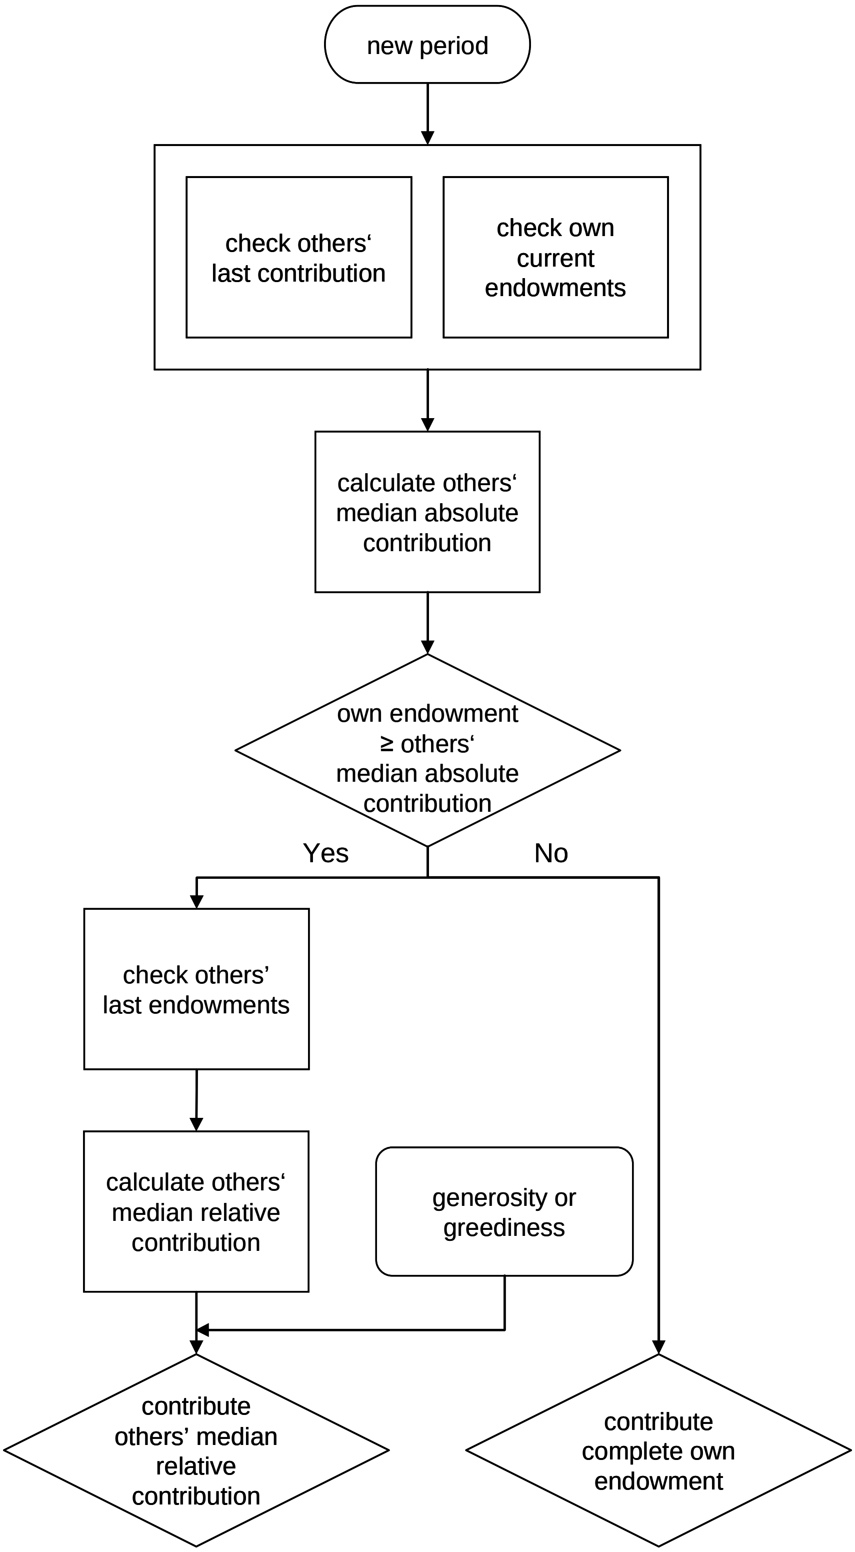

Figure 3.9. Flow chart for opportunistic reciprocity. Others intentions are only considered, when the own endowment is higher than others median contribution in the previous round.

Using opportunistic reciprocity to explain the data of Experiment 3.1 maintains the assumption that all reciprocators reciprocate the same aspect of others' behavior. An alternative possibility is that reciprocators differ in their regard for the consequences or intentions inferred from others' behavior; some reciprocate intentions, some consequences, and some act opportunistically. As described above, participants did not decide frequently enough in situations for which the models make different predictions to test among the models for participants separately. Hence, Experiment 3.2 will investigate if opportunistic reciprocity is the best model for all reciprocators and if all reciprocators can be described best with the same reciprocity model.

3.5 Experiment 3.2

The opportunistic reciprocity model predicts that if an individual's own endowment is lower than the other players' absolute contributions, the individual should contribute his complete endowment, even if the other players' contributions are low in relative terms. If the individual's own endowment is equal to or higher than others' past contributions in absolute terms, the player should contribute the same relative amount as the other players. Consequential reciprocity differs form opportunistic reciprocity in that it reciprocates intentions, and it differs from intentional reciprocity in that it reciprocates only when in absolute terms the individual's own endowment is larger than others' contributions.

3.5.1 Method

Experiment 3.2 was designed so that participants would frequently encounter situations discriminating between the different reciprocity models. Participants played two rounds of a public goods games by using the "strategy method" (Selten & Stöcker, 1986). When using this method participants indicate their decisions for all possible stages of a game, which are derived from hypothetical decisions—that is, all possible combinations of contributions all players can make. Specifically, participants had to decide about their contribution in the present round conditional on all hypothetical contributions in the preceding round (henceforth: vector of past contributions). As in Experiment 3.1, participants' payment depended on their decisions.

3.5.1.1 Design

The Experiment varied individuals' *endowment level*, which was either high (twice the median of others' absolute contributions) or low (half the median of others' contributions), and the median of others' *last relative contribution*, which was either 50 or 100% of their endowment[17], resulting in four conditions, referred to as high-50%, high-100%, low-50%, and low-100%. Participants' relative contributions were the main focus of examination. To limit the

[17] Actually, participants also made contribution decisions conditional on others' contributions between 50 and 100%. However, these decisions were part of the experiment only to enable the calculation of payoffs.

number of possible stages of the games, the feasible contributions in the first round were limited, so that participants had to choose one of two contributions for every endowment they had. Table 3.5 depicts endowments and corresponding possible contributions in round 1 as well as players' endowments for their decisions in round 2. The feasible contributions in round 1 were chosen so that most decisions in round 2 combined the players' own endowments (high or low) and others' past contributions (50 or 100%) according to the four conditions.

Table 3.5. *Endowments and possible contributions in Experiment 3.2.*

Game	Endowment			
	Participant 1	Participant 2	Participant 3	Participant 4
A	12 (12 / 6)	12 (12 / 6)	12 (12 / 6)	12 (12 / 6)
B	12 (10 / 6)	12 (10 / 6)	12 (10 / 6)	12 (10 / 6)
C	12	6 (6 / 3)	6 (6 / 3)	6 (6 / 3)
D	8 (4 / 8)	8 (4 / 8)	12 (6 / 12)	16 (8 / 16)
E	20 (10 / 20)	12 (6 / 12)	16 (8 / 16)	8 (4 / 8)
F	4 (3 / 4)	8 (4 / 6)	12 (6 / 9)	16 (8 / 12)

Note. Endowments of players for conditional contribution decisions are underlined. Feasible contributions for players in round 1 are in parentheses. For distributions where endowments differed between players, participants made one decisions per possible personal endowment. Every endowment point was worth 0.04 euros.

In the first round, participants made decisions for six public goods, differing only in the distribution of endowments across players (see Table 3.5). In the second round, participants received information about others' hypothetical contributions and their real endowments in the preceding round and made unrestricted contribution decisions. For instance, in game A presented in Table 3.5, all players were endowed with 12 points in the first round and could contribute 6 or 12 points. In the second round of game A, participants decided how much of their 12-point endowment to contribute for the hypothetical cases of the following past-contribution vectors: [6 6 6], [6 6 12], [6 12 12], and [12 12 12]. In the second round participants made 36 decisions, one for every vector of past contribution that could arise from the 6 games of the first round. To motivate participants, in each round one decision per game was given as payment for participation in the Experiment.[18]

3.5.1.2 Participants and Procedure

The 40 participants[19] (25 women and 15 men) with an average age of 25 years were mainly students from different departments of the Free University Berlin. Sessions were conducted in groups of 4 participants. When participants arrived at the laboratory they were seated in two rooms, each with two separate cubicles, preventing any communication. After the public goods game was explained with the same instructions as in Experiment 3.1, participants were instructed that they would make decisions in two modes: mode 1 providing information about others' endowments in the same round, and mode 2—starting in the second round—providing information about others' hypothetical contributions and endowments in the preceding round. Participants were further informed that within each round they would make decisions for six different games, which differed in the endowment distributions. Participants were not told how many rounds the experiment would last. Finally it was pointed out that at the end of the experiment for every round one payoff per game would be determined, based on participants' randomly matched contribution decisions. After reading the instructions, participants completed the same test quiz as in Experiment 3.1 and received a summary of the instructions.

Participants entered decisions on a computer interface, which depicted information about others' endowments and hypothetical contributions in table form. All participants made decisions in the same order. For the conditional contributions of round 2, others' hypothetical contributions were ordered so that median contributions increased within games. Finally participants made unconditional contributions in round 2 that were used to calculate payoffs. After finishing the second round participants drew a lot determining which of the six games were used to calculate payoffs and were paid individually. Sessions lasted on average 70 minutes and participants earned on average about 13 euros, including 5 euros as a show-up payment.

3.5.2 Results

Before I tested the predictions, participants were first classified into reciprocators and non-reciprocators. The classification procedure of Experiment 3.2 differed from that of Experiment

---

[18] For the first round it was randomly determined which player assumed which role (i.e. had which endowment) for each game, before decisions were combined to determine the payoff. For every game the conditional contribution in round 2 of one participant was combined with the unconditional contributions of the other three participants in round 2. Which role (conditional, and which of the unconditional decisions) a player assumed was again determined randomly. In the experiments every participant would draw one of four profiles from an envelope and then payoffs were determined according to participants' roles and their contributions.

[19] Due to technical problems I lost data of one participant.

3.1, because I wanted to classify people on a set of decisions for which all models made the same contribution prediction. The classification was based on the two conditional contributions in game A (see Table 3.5), for which all models predicted a contribution of the same absolute amount as others in the previous round, as participants had the same endowment as the other players in the preceding round. Twenty-three of 39 participants (59%) who increased their contributions by at least 3 points with the increase of others' median contributions from 6 to 12 points (corresponding to a slope of .5) were classified as reciprocators.

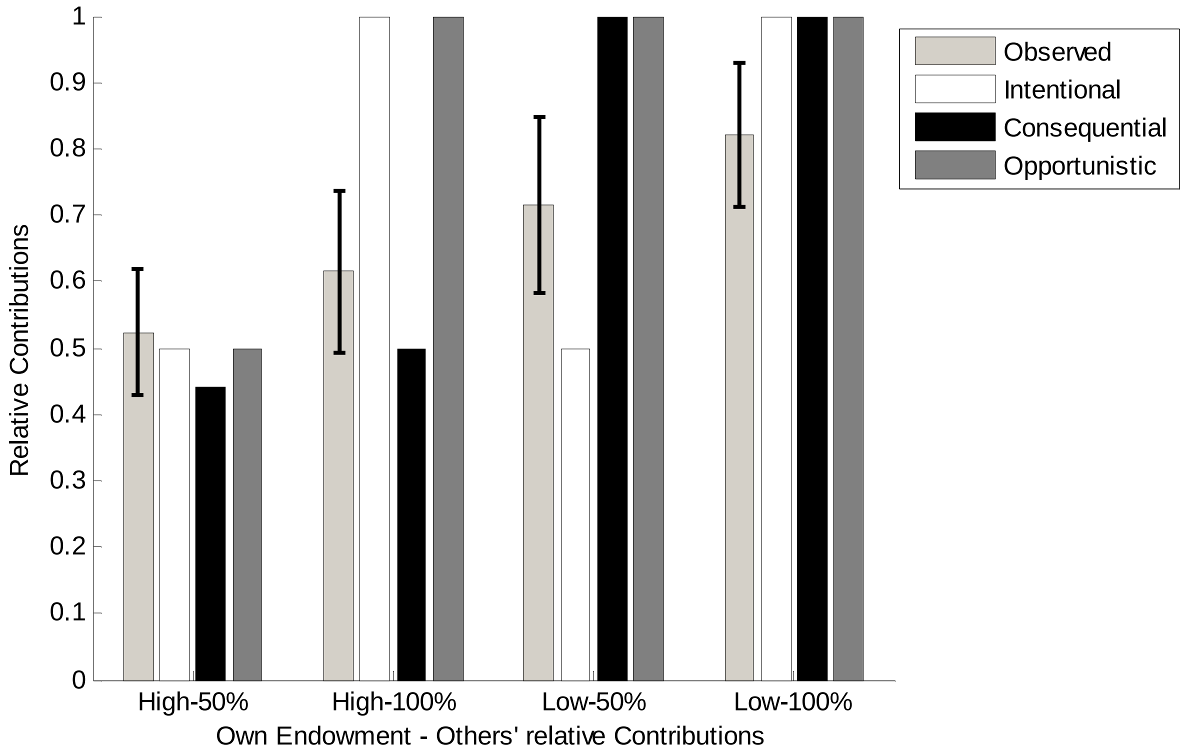

Figure 3.10. Mean and standard errors of participants' conditional contributions in the four conditions in Experiment 3.2.

Mean contributions of reciprocators were 52% (*Mdn* = .50%, *SD* = 13%) in high-50% (based on 10 decisions per participant in round 2), 61% (*Mdn* = 67%, *SD* = 26%) in high-100% (6 decisions), 71% (*Mdn* = 75%, *SD* = 21%) in low-100% (8 decisions), and 82% (*Mdn* = 83%, *SD* = 16%) in high-100% (6 decisions), compare also Figure 3.7. To test the opportunistic reciprocity model I conducted an ANOVA with others' relative contributions (50 vs. 100%) and the player's own endowment level (low vs. high) as two repeated measures factors. The intentional reciprocity model predicts a main effect for others' relative contributions, whereas consequential reciprocity predicts a main effect for the endowment level. Opportunistic reciprocity predicts both main effects and also the interaction, as relative contributions should not matter when the endowment is low. The results show a main effect for the endowment level, $F(1,22) = 38.63$, $p < .001$, $\eta 2 = .64$, and for others' relative contributions, $F(1,22) = 9.56$, $p = .005$, $\eta 2 = .3$, but no interaction effect, $F(1,22) = .14$, $p = .714$, $\eta 2 = .006$. While the finding of two main effects is only in line with the reciprocity model, the model is not fully supported by the data, because the interaction that was also predicted did not appear.[20]

To explore the reason for this ambiguous result I examined if the mean results found are representative for all reciprocators or if reciprocators differed in the aspects of others' behavior they imitate. In a first step I classified reciprocators as being intentional, consequential, or opportunistic. The classification was based on the Euclidean distance between reciprocators' observed contribution vectors in the four conditions and the vectors predicted by the models. Reciprocators were classified according to the reciprocity model with the smallest distance to their observed vector. I identified 6 intentional, 6 opportunistic, and 11 consequential reciprocators. To validate this result I conducted a cluster analysis. The rationale was to examine if contribution vectors describing the three concepts of reciprocity would emerge without being pre-specified and if the same participants would be grouped together. Specifically, I employed a k-means cluster analysis with random start clusters and with reciprocators' normalized mean contributions in the four conditions as the dependent variable. The maximal number of clusters was set to four, as most research has identified three to four types of players in public goods games (e.g. Fischbacher et al., 2001). Five thousand iterations of the cluster analysis were performed, and the mean silhouette value (the silhouette value can vary between -1 and 1, values larger values indicating a better solution, Martinez & Martinez, 2004) of participants was used as selection criterion. The best solution with a silhouette value of .43 had four clusters. Table 3.6 depicts the centers of the identified clusters. The first cluster with eight participants (which included six opportunistic and two consequential reciprocators as classified above) shows a contribution pattern where contributions are low in high-50% and higher in all other conditions, indicating that participants in this cluster behaved according to opportunistic reciprocity. The second cluster (six participants who are identical with intentional reciprocators as classified above) shows higher contributions when others' relative contributions were high and lower contributions when others' relative contributions were low, suggesting that participants in this cluster were intentional reciprocators. The third cluster shows a contribution pattern in

[20] I also compared the models with the likelihood ratio test described in Dixon (2003), which is the ratio of the variance unexplained by two models to the power of half the number of observations in the experiment. Comparing the two models with the most unexplained variance I found that opportunistic reciprocity fits the data better than intentional reciprocity, $\lambda = 122.87$. Comparing the two models with the least unexplained variance, I found that consequential reciprocity fits the data better than opportunistic reciprocity, $\lambda = 28.1$.

accordance with consequential reciprocity and includes five participants who had also been classified as consequential reciprocators. The fourth cluster includes four participants (all classified above as consequential reciprocators) who contributed more when their endowment was low and decreased their contributions when others' relative contributions increased.

Table 3.6. *Cluster centers of cluster analysis for Experiment 3.2.*

Cluster	Others' contributions				*N*	Silhouette value
	High-50%	High-100%	Low-50%	High-50%		
Opportunistic	.63	.81	.86	.91	8	.48
Intentional	.57	.75	.47	.84	6	.69
Consequential	.37	.38	.79	.79	5	.08
Non-classified	.44	.3	.69	.38	4	.40

Note. Cells show relative contributions.

3.5.3 Discussion of Experiment 3.2

Experiment 3.2 examined if participants' behavior is best described by the consequential, opportunistic, or intentional reciprocity model. On the group level I found that participants reacted to both aspects of others' behavior—intentions and consequences—and not only to one of the two aspects as predicted by consequential or intentional reciprocity. However, because the interaction between the two factors was not there as predicted by opportunistic reciprocity, I cannot conclude that reciprocators are accurately described as opportunistic reciprocators on the group level. The results of the two classification procedures suggest that the mixed result on the group level is due to differences between players, who reciprocate different aspects of others' behavior. The two classification procedures identified all types of reciprocators I examined in this paper: intentional, consequential, and opportunistic.

3.6 General Discussion

The research reported here was founded in the assumption that people can reciprocate different aspects of others' behavior. Specifically, I first argued that at least two interpretations of reciprocity can be distinguished, namely, consequential and intentional reciprocity. The concept of intentional reciprocity assumes that people reciprocate the intentions they inferred from others' behavior. Consequential reciprocity assumes that people reciprocate the observed consequences of others' behavior. To contribute to the understanding of reciprocity for cooperation in groups, I examined the two concepts in public goods games. For public goods games, I defined intentional reciprocity as contributing the same relative amount as others and consequential reciprocity as contributing the same absolute amount as others.

Consequential and intentional reciprocity were tested against each other in Experiment 3.1 by examining people's behavior in repeated public goods games, with endowments varying within and between participants. Approximately half of the participants (49%) in Experiment 3.1 were classified as reciprocators. Comparing model fits, I found that the intentional and the consequential models predict reciprocators' behavior equally well. In contrast to the prediction of the intentional reciprocity model, I found that participants contributed the same amount when others had no other choice than making a low contribution, compared with when other players defected deliberately. Participants also contributed similar amounts subsequent to rounds in which others' contributions were high in relative terms and low in absolute terms, compared to rounds with high absolute and low relative contributions. This indicates that participants' behavior cannot exclusively be accounted for by intentional or consequential reciprocity. Comparing contributions of wealthy and poor participants, I found that contributions were more in line with consequential reciprocity. Participants searched—in accordance with intentional reciprocity—for information about others' contributions and endowments. Altogether, neither intentional nor consequential reciprocity alone could account for the findings in Experiment 3.1. Hence I argued that even the distinction of intentional and consequential reciprocity does not provide a sufficient description of reciprocal behavior. Instead I suggested the alternative concept of opportunistic reciprocity, which can explain several findings of Experiment 3.1.

In Experiment 3.2 the opportunistic reciprocity model was tested against the two alternative reciprocity concepts. The opportunistic reciprocity model predicts that individuals contribute their complete endowment when it is lower than the median of others' absolute contributions in the preceding round, and otherwise reciprocate intentions. Even though Experiment 3.2 showed that reciprocators react to variations in others' absolute and relative contributions as predicted by opportunistic reciprocity, this model cannot explain contributions of all participants. Instead I found that similar numbers of participants were best described with consequential, intentional, or opportunistic reciprocity. Only when taking all three reciprocity concepts presented in this article together could reciprocal behavior be explained sufficiently.

While my experiments add to the understanding of cooperation in groups, the conclusions also need to be limited. In Experiment 3.1 participants played a repeated public goods game in which each participant had a free choice of how much to contribute to the public good in every single round. The experimental setting is comparable to other public goods experiments and captures the essence of public good situations outside of the laboratory. Nevertheless, without any restrictions on participants' behavior, participants frequently made decisions for which the different reciprocity models I considered make the same predictions, thus decreasing the experimental power of Experiment 3.1. Another limitation of my conclusions arises from participants' information search. Participants in Experiment 3.1 had to pay for each piece of information. The search costs were introduced to examine which information was regarded as essential by the participants, but it led to relatively limited information search so that even reciprocators were not always fully informed about others' contributions in the preceding round. However, participants in Experiment 3.2, in which no monetary search costs were imposed, made similar contribution decisions. Finally, the classification of reciprocators into different types needs to be qualified. While the cluster analysis identified types that are readily interpretable in terms of the reciprocity models and the classification of participants showed a high overlap with the first classification based on the three models of reciprocity, the cluster analysis was based on a relatively small number of participants. Hence I could not use a cross-validation procedure that tests if different sub-samples would identify the same clusters and a similar frequency of participants in the clusters.

A research program related to the research presented in this chapter is that of van Dijk and colleagues (van Dijk & Wilke, 2000; van Dijk & Wilke, 1995; van Dijk, Wilke, Wilke, & Metman, 1999), who examined decision rules individuals use in asymmetric social dilemmas. My work supplements their work in two ways. First, while van Dijk and Wilke's proportional rule, which models—similar to my intentional reciprocity model—people's contributions as contributing the same proportion as others' to public goods, is motivated by fairness considerations, my analysis of the ecological rationality of intentional reciprocity shows that self-interested or efficiency motives can also motivate proportional contributions. The self-interest hypothesis is especially appropriate to describe the behavior of poor players in Experiment 3.1, who contributed more in relative terms, and with the opportunistic reciprocators in Experiment 3.2, who contributed higher relative amounts than others when they were relatively poor themselves. While this might seem "unfair," it helps to maintain beneficial cooperative interaction in the long run. This "selfish" interpretation of reciprocity is in agreement with findings that reciprocity is used by self-interested players (Andreoni & Miller, 1993; Kreps & et al., 1982), who reciprocate because it leads to a high payoff in the long run. On the other hand, Marwell and Ames (1979) reported that 75% of their participants were concerned with being fair when contributing to a one-shot public goods game.

Second, while van Dijk and Wilke (2000) found that nearly all participants in a game can adequately be described with the proportional contribution rule, I found that only some participants contribute the same relative amount as others. Different reasons can explain my divergent finding. First, van Dick and Wilke's classification was based on fairness ratings about contributions others could make, whereas my classification is based on contributions actually made by participants. Second, van Dijk and Wilke explicitly defined two classes of decision rules, equal final outcome and proportional contribution. By contrast, I compared three decision rules, of which only one corresponds to the proportional decision rules.

The argument that different individuals apply different decision strategies has been made before. For instance, Zwick and Rapoport (2002) argued that players in a market entry game are heterogeneous in that they are best described with different decision mechanisms, and Fischbacher et al. (2001) and Kurzban and Houser (2005) have shown in the context of public goods situations that some players cooperate or defect unconditionally, whereas others reciprocate. Likewise, in the domain of individual decision making it has been argued that people's decision strategies differ (e.g. Bröder, 2003; Howe & Loftus, 1992). Consistent with previous research on public goods games, in both experiments the largest group of individuals could be classified as reciprocators. However, I also showed that reciprocators are not a homogenous group. While reciprocators all cooperate conditional on others' behavior (and their own wealth position), they differ in the aspects of others' behavior they imitate. Previous experiments could not identify different types of reciprocators, because they did not use asymmetric or variable endowments in repeated or sequential public goods games. As most daily interactions takes place in an environment with unequal and variable access to resources, consideration of different types of reciprocators will foster the understanding of cooperative behavior in groups.

While the models describe players' contributions, they do not explain why individuals choose certain strategies or predict which person adheres to which strategy. One explanation for the different types of reciprocators is that they might have different social motives (van Lange, 1999). An examination of effects of the different models of reciprocity on outcomes in repeated games can help to identify potential motives of different types of reciprocators. Consequential reciprocity will generally lead to unequal outcomes in favor of the wealthier person and might thus be favored by people with individualistic social motives who try to maximize their own outcome. However, because consequential reciprocity will also lead to decreasing contributions over time, which should not be in the interest of individualistic people, a straightforward interpretation of consequential reciprocity in terms of social motives seems difficult. Intentional reciprocity leads to outcomes that are generally more equal than when applying consequential reciprocity and also insures generally higher contribution rates and might thus be favored by people with prosocial motives who maximize joint outcomes and/or equality in outcomes. Opportunistic reciprocity will lead to outcomes that are less equal but more efficient (in terms of joint outcomes) and might thus also be favored by people with prosocial motives. Support for the hypothesis of different motivations underlying the same behavior comes, for instance, from Kuhlman and Marshello (1975), who found that not only prosocials but also individualists reciprocate against a Tit-For-Tat strategy in repeated prisoner's dilemma games (see also van Lange, 1999).

Assuming a direct link from motives to strategies is only one way to explain why people choose certain strategies. An alternative explanation is that strategies are acquired by individual learning (e.g. Rieskamp & Otto, submitted for publication) or social learning (e.g. Joseph Henrich & McElreath, 2003) and are applied contingent on signals connecting decision strategies to appropriate environments (March, 1996; J. M. Weber et al., 2004). For instance, a player in interaction with other wealthier players might follow the rule "Be nice to the people you need," whereas a player interacting with equally wealthy players might use the Tit-For-Tat rule. The rule-based approach seems especially relevant for repeated interaction, as there is little evidence that social values are predictive for cooperative behavior in repeated interactions, as, for instance, Chapter 2 showed. Future research will have to examine how individuals learn and choose among reciprocal strategies.

Previous research examining cooperation in groups neglected the different types of reciprocity that can explain cooperative behavior. The role of intentions for cooperative behavior had primarily been studied in dyadic interactions. I extended this research to cooperation in groups and inquired if people reciprocate intentions or consequences of others' behavior by proposing and testing three realizations of reciprocal cooperation under variable endowments. As it turned out, no single concept of reciprocity can explain cooperation; instead I could show that most—opportunistic—reciprocators considered others' intentions only if their own endowment was relatively high, whereas only a minority of reciprocators considered exclusively consequences or intentions. In general, the different types of reciprocity illustrate that reciprocity is not a one-dimensional concept but has to be understood in its multiple facets.