

CHAPTER 5

Summary and Conclusion

5.1 Summary

The aim of this dissertation was to explore how people maintain cooperation in groups and how social learning helps them to make good decisions. Chapter 1 of introduced the public goods game as a general model of cooperation in groups. A game theoretic analysis of the game showed that cooperation is difficult to maintain because defection is individually rational in finitely repeated public goods games. However, a review of experimental results highlighted that people cooperate more than predicted by classical game theory, and that cooperation declines with time, affording alternative theories of cooperation. To set the stage for the experimental chapters, Chapter 1 briefly described models of cooperation in groups, the decision rule approach, the learning approach, and the social values approach, which have all been suggested to explain cooperation in groups. Finally I sketched approaches describing how individuals learn to make good decisions. These approaches were evolutionary adaptation, individual learning, and social learning.

The aim of Chapter 2 was examine how decision rules and the interaction structure work together to determine cooperative behavior. Specifically, I observed participants' behavior either the repeated public goods games, or the in the repeated Social Dilemma Network. In the Social Dilemma Network, each of the four players could simultaneously contribute to three different two-person public goods games, each participant with every other player in the game. By using the information board technique, I also examined players' information search behavior. The results showed that while contributions were stable on an intermediate level in the public goods game, contributions in the Social Dilemma Network started on a higher level and increased further. I thus concluded that partner selection, as possible in the SDN, increases cooperation in cooperative groups. I compared three approaches explaining cooperation in groups, the reciprocity heuristic, simple reinforcement and local adaptation learning, and social value orientation. Social value orientation did not predict contributions. Among the learning models and the reciprocity heuristic, the reciprocity heuristic was the best model to describe participants' behavior because it predicted participants' contributions and their information search. Finally, I also found that players adapt their decision rule to the environment because players in the Social Dilemma Network behaved more according to the reciprocity heuristic than players in the standard public goods game. In sum, Chapter 2 showed that cooperation in groups is best described by the reciprocity heuristic, that players adapt their decision rules to the interaction

environment, and that partner selection increases cooperation in groups, even when efficiency gains of cooperation are higher in groups without partner selection.

Most models of reciprocity, including the model tested in Chapter 2, implicitly assume that people reciprocate the consequences of others' behavior. The aim of Chapter 3 was to investigate to which aspect of the others' behavior, consequences or intentions, reciprocators respond. I defined intentional reciprocity as contributing the same relative amount as others, and consequential reciprocity as contributing the same absolute amount. Experiment 3.1 examined cooperation in repeated public goods games, where endowments varied within and between participants. Approximately half of the participants were classified as reciprocators. Reciprocators looked up information on the others' contributions and endowments, as was predicted by the intentional reciprocity model, but reciprocators' contributions were equally well-described by intentional and consequential reciprocity. Hence, I suggested an alternative rule, the opportunistic reciprocity, which could account for most of the results in Experiment 3.1. Opportunistic reciprocity predicts that individuals contribute their complete endowment when it is smaller than the median of the others' absolute contributions in the preceding round, but otherwise reciprocate intentions. Experiment 3.2 showed that this model well-described reciprocators' behavior on the group level. Beyond this, however, the analysis of individual contributions revealed that one third of the reciprocators respectively reciprocated intentions, consequences, or opportunistically. In sum, the results of Chapter 3 suggest that a large part of reciprocators consider others' intentions. Intentional reciprocators consistently consider others' intentions, opportunistic reciprocators consider others' intentions when they are relatively wealthy, and consequential reciprocators disregard others' intentions and instead reciprocate the outcome of others' behavior.

While the different models of reciprocity describe players' contributions, they do not explain why individuals choose certain strategies or predict which person adheres to which strategy. Possible explanations are that decision rules are acquired by individual learning (e.g. Rieskamp & Otto, submitted for publication) or social learning (e.g. Bandura, 1977). The rule-based approach seems especially relevant for repeated interactions, as Chapter 2 showed that social values are do not predict cooperative behavior in repeated interactions (see also Parks, 1994). Hence the next chapter examined social learning as one way of acquiring decision rules. When examining the learning of decision rules, two processes, the acquisition and the selection of decision rules, determine individuals' behavior.

The aim of Chapter 4 was to examine how people learn to make good decisions by combining advice, as one form of social learning, with their own experience. To investigate

social learning, I observed choices in the multi-armed bandit paradigm, specifically the Iowa Gambling Task (IGT), where advice receivers first received a decision strategy as advice, which prescribed to choose always one of four options, and then made 100 choices. Receivers in both experiments in Chapter 4 chose, on average, the advised choice option more frequently than the corresponding option with the same expected value, indicating that choices were influenced by advice. Moreover, in Experiment 4.2, people chose a bad option with a lower expected value that was advised more frequently, than better options with high expected values that were not advised. A comparison of the earnings of decision makers with and without advice showed that good advice improved performance. To identify how participants combine advice and their own experience with choice options, I proposed one model of individual and four models of social learning (models of advice reinforcement combination, ARC). A first comparison of the social learning models identified ARC-Reinforcement (assuming comparatively higher reinforcements from payoffs from the advised options), ARC-Decay (propensities of advised options decay less), and ARC-Choice (receivers choose the advised option when they are uncertain which is the best) as the best social learning models because only these models could predict the choice frequencies of good decks as well as the level of adherence to advice. Experiment 4.2 further tested between ARC-Reinforcement, ARC-Decay, and ARC-Choice by implementing a multi-armed bandit task, for which the three models predicted different behavior. Experiment 4.2 showed—contrary to the prediction from ARC-Decay—that bad advice led to more choices of bad options and that ARC-Choice overestimated adherence to advice to a large extent. Only ARC-Reinforcement was able to predict the probability with which participants chose good decks, and also predicted that adherence to advice was higher after receiving good advice. In sum, results in Chapter 4 showed that social learning can help people to make good decisions, and that social learning in repeated choice tasks is best modeled by assuming that payoffs from advised options lead to stronger reinforcements.

5.2 Conclusion

What can the reader of this dissertation learn about cooperation in groups and social learning? This conclusion derives four theses from the three empirical chapters. Briefly, these theses are: (1) people reciprocate in groups, (2) intentions matter—but not for all in the same manner, (3) people choose decision rules dependent on the decision environment, and (4) social learning influences decisions persistently. To examine the relevance of these theses, I relate them to previous research, show what they add to past and current research.

People reciprocate in groups. While this statement is certainly not provocative, it nevertheless adds to what was known so far about cooperation in groups. Earlier experiments in social psychology could not identify reciprocity in groups at all, if it was not seeded by a simulated player (see Bornstein et al., 1994; Komorita et al., 1992). Recent research on reciprocity in groups identified reciprocity in rather specific institutions. For instance, the others' contribution was the only information available (Fischbacher et al., 2001), or people could make binding commitments (Kurzban et al., 2001). Hence, the results of my dissertation extends knowledge on cooperation in groups by showing that reciprocal behavior can emerge without a specific institutional design—such as the possibility of binding commitments—and without external sources providing role models that can be imitated (Komorita et al., 1992). A further novel contributions is the direct comparison of reciprocity with alternative explanations. The proposed reciprocity heuristic explained participants' contribution more successful than alternative explanations, such as reinforcement learning, directional learning, or social value orientation. A question that remains unanswered from this dissertation is, to what aim do people reciprocate. Some researchers suggest that people reciprocate for egoistic reasons (Andreoni & Miller, 1993; van Lange, 1999), but others argue that reciprocate intentions to maintain equity (e.g. Falk & Fischbacher, 2000).

Intentions matter—but not for all in the same manner. Building on the result that people reciprocate, Chapter 3 investigated the role of intentions for reciprocal cooperation in groups. The results suggest that players care about why others cooperated or defected. While psychological research rarely considered the role of intentions for reciprocal cooperation (but see Kelley & Stahelski, 1970b), economists' work considered, so far, only situations in which positive intentions could actually be realized (e.g. Falk & Fischbacher, 2000). I contributed to this research by independently varying outcomes and intentions behind others behavior. To incorporate intentions in reciprocal cooperation, I proposed a simple reciprocal heuristic, which uses a relative contribution as an indicator for the others' intentions. The examination of players' contributions revealed that different individuals reciprocate different aspects of the others' behavior. Some reciprocate intentions, some reciprocate consequences, and some reciprocate opportunistically. Similar results were found by Howe and Loftus (1992), who reported that people differ in their regard for plaintiffs' intentions and the consequences of plaintiffs' behavior when choosing levels of punishment. As an explanation for opportunistic reciprocity, I suggested, in line with previous self-interested interpretations of reciprocal behavior (Andreoni & Miller, 1993), that poor players contribute more than the same relative amount of others because they want to maintain the others' cooperation. The opportunistic reciprocity heuristic

highlights another theme of this dissertation, which is that people select decision strategies dependent on characteristics of the social environment.

People choose decision rules dependent on the decision environment. Opportunistic reciprocity is sensitive to the environment, in that the rule prescribes to cooperate unconditionally only if one is a poor person among wealthy people. Participants in the Experiment 2.1 also adapted their behavior to the environment. This insight resulted from the finding that participants in the Social Dilemma Network were better described by the reciprocity heuristic than players in the ordinary public goods game. Others have examined the effect of partner selection on cooperation in public goods games (Coricelli et al., 2004), but they did not test if different institutional settings would influence participants' decision rules. The finding that people choose cooperative decision rules dependent on the social environment generalizes results on individual decision making, which showed that people adapt decision rules to the environmental structure (e.g. Rieskamp & Hoffrage, 1999). The reason for my finding of more reciprocity in the Social Dilemma Network might be that reciprocation in a two-person interaction always directs reward and punishment to the desired target, whereas this is not always possible in public goods games. If a player observes two cooperators and defectors in a public goods game, he or she cannot reward cooperators and punish the defector simultaneously as he or she could in a Social Dilemma Network. The question that poses itself when reporting on the influence of different institutional features (i.e., partner selection, the available information, the possibility of binding commitments) on the level of reciprocation is how these can be systematized, and which of them is more efficient in supporting reciprocal cooperation. Pursuing which institutions support decision rules that people already have seems especially worthwhile because my results showed that allowing reciprocal heuristics to work effectively increased cooperation more than doubling the efficiency gain through cooperation.

Social learning influences decisions persistently. While the first part of the dissertation was concerned with how people make cooperative decisions, the last part did first steps to explain how people learn to make decision. Others have argued before that cooperative behavior has its roots in people's docility, their propensity to learn from others (Dewey, 1922; Simon, 1990). A first finding was that participants' learning process followed an adherence-exploration-adherence pattern. Social learners first adhered to the decision strategy they had received as advice, then explored alternative options, and finally returned to the advised decision strategy. Interestingly, this result was not only true for participants receiving good advice but also, to a lesser degree, for participants receiving bad advice. Previous social learning research examined situations where decision makers received new advice before every single situation (e.g. Yaniv, 2004b). The key

contribution of the chapter on social learning is that it shows how people use a single piece of advice for repeated choices. This is a relevant contribution, because people often get social information—like advice—only once, or a few times, and then have to go on themselves. Further, learning theories mostly describe either individual learning (e.g. Yechiam & Busemeyer, 2005) or social learning (e.g. Bandura, 1977). The formulation of ARC-Reinforcement supplements these theories by describing the mechanism with which individual experience and social information are combined: people perceive outcomes from options preferred from others as more positive than the same outcome from other options.

The general aim of this dissertation was to examine how people maintain cooperation in groups, without being supported by specific institutions. The general conclusion is that most people maintain cooperation by using a reciprocal decision rule, that is, they cooperate as much as others' did in the previous interaction. Accordingly, cooperation can be improved if one makes use of people's tendency to reciprocate. When people reciprocate, most of them consider others' intentions, allowing them to maintain cooperation also in social environments where the ability to cooperate varies. Finally, I could show that social learning is a possible explanation of how people learn to make successful decisions. Indeed, it might be people's propensity to social learning, their docility, which sustains cooperative behavior.

“...society can “tax” docile individuals by sometimes persuading them to take altruistic actions that decrease individual fitness. So long as the cost of the altruism to them is less than the gain obtained from improving their choices [through docility], docile individuals will remain fitter than nondocile ones.”

Herbert Simon (1993), p. 157