University of Groningen

Learning and framing in social exchange

Flache, A.; Macy, M.W.

Published in:
Solidarity and Prosocial Behavior

DOI:
[10.1007/0-387-28032-4_3](https://doi.org/10.1007/0-387-28032-4_3)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2006

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Flache, A., & Macy, M. W. (2006). Learning and framing in social exchange. In D. Fechtenhauer, A. Flache, A. Buunk, & S. M. Lindenberg (Eds.), *Solidarity and Prosocial Behavior: An Integration of Sociological and Psychological Perspectives* Springer. https://doi.org/10.1007/0-387-28032-4_3

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

CHAPTER 3

Learning and Framing in Social Exchange

Andreas Flache and Michael W. Macy

Reciprocity is one of the most widespread and persistent norms for regulating behavior in long-term social relationships (Gouldner, 1960). Successful relationships are typically characterized by a pattern of reciprocal solidarity, while “sour” relationships display the mirror image—a pattern of mutual recrimination (e.g., Buunk and Dijkstra, this volume). We define solidarity as behavior that benefits others at some cost that is not immediately compensated (cf. Lindenberg, Fetschenhauer, Flache, and Buunk, this volume). Reciprocal solidarity is widely observed in exchanges between family members, firms, and nation-states.

The Puzzle of Reciprocity

The robustness and prevalence of reciprocal solidarity presents a puzzle. Although both sides benefit from mutual cooperation (the exchange of valued resources), each is also tempted by incentives and opportunities to unilaterally defect (to fail to reciprocate the partner’s solidary behavior). For example, in business relations such temptation may occur when a firm has already received payment for a shipment from a long-standing client, but then suddenly receives an order from a new client that may be willing to engage in future contracts. The supplier is tempted to give priority to the potential new business relation and to delay the shipment to its old client, violating in the process the norm of due delivery for due payment in relations with long-standing partners. The result may be something that neither of the two old business partners actually wants: deterioration of the relationship into mutual distrust or even disruption of the business contact. Yet numerous experimental studies of exchange behavior have demonstrated the robustness of reciprocal solidarity despite opportunities to cheat the exchange partner without danger of being detected or punished (e.g., Buunk and Schaeffeli, 1999; Ligthart, 1995). Why is reciprocal

solidarity in ongoing exchange relations so widespread and why is the norm of reciprocity so robust across a variety of exchange situations?

In this chapter, we compare two proposed approaches—learning and framing—that address why and when reciprocity may prevail despite the opportunistic temptation to cheat. In framing theory (Lindenberg, 1998 and this volume), situational characteristics and the history of the relationship determine whether actors frame their decisions in the exchange primarily in terms of gain or whether gain as a dominant goal is tempered or even replaced by the goal to follow norms of solidarity. These norms can vary in strength. Solidarity norms are strong when solidary behavior is the actors' dominant goal (i.e., their frame) of the exchange situation. When solidarity norms are present but weak, then the gain motive still dominates, but it is tempered by solidarity as a goal in the background. With strong or weak solidarity, participants in an exchange tend to ignore incidental temptation to defect and they also tend to forgive each other if things occasionally go wrong. Only strong temptations or consistent violations of reciprocity by the exchange partner may bring about such a salient gain frame that the relationship deteriorates into mutual defection.

We compare framing theory to a learning-theoretic alternative. Our learning model draws upon Thorndike's (1911) Law of Effect, which is based on the principle that "pleasure stamps in, pain stamps out." We show that the Law of Effect suffices to model reciprocity in long-term relationships, without the need for elaborations to the theory such as role modeling or social learning (Bandura, 1977) in which actors imitate rules observed in significant or successful others. Using Rapoport and Chammah's (1965) application of the Bush-Mosteller stochastic learning model, Macy (1991) showed how penalty-aversive, reward-seeking agents can elude the trap of mutual defection and establish a successful ongoing exchange relationship. A random sequence of bilateral outcomes (either mutual cooperation or mutual defection) can lead adaptive agents out of the "social trap" of mutual defection into stable mutual cooperation, a process he characterized as "stochastic collusion."

Growing interest in framing and learning reflect widespread criticisms of two alternative theories of solidarity in social exchange—rational calculation and natural selection. In the next section, we outline these criticisms and show how framing and learning approaches avoid these limitations. Then we compare the behavioral assumptions underlying framing and learning theory and the implications of framing and learning for the conditions and dynamics of reciprocal solidarity. We conclude with a discussion of the complementarities of the two theories.

Competing Approaches: Rational Choice and Evolution

Rational choice theory models actors as cognitively sophisticated and self-interested decision makers who evaluate all possible future consequences of alternative actions and select the action that maximizes their self-interest (Coleman, 1990, pp. 13–19). In this view, reciprocal solidarity is a rational response to enlightened self-interest. In repeated interactions, egoistic actors who value future outcomes may be better off if they resist the temptation to take advantage of a cooperative partner and instead reciprocate. So long as there is a sufficiently long “shadow of the future” (Axelrod, 1984; cf. Friedman, 1971), the expectation of long-term gain through a cooperative relationship will deter rational actors from adopting a strategy of “hit and run” that is likely to bring the relationship to an end. This expectation, in turn, rests on the rational expectation that one’s partner will retaliate if cheated and will likewise resist the temptation to cheat. Accordingly, orthodox rational choice explanations rest on the assumption of a “reflexive rationality of actors anticipating each others’ choices” (Scharpf, 1990, p. 471).

Numerous critics have argued that this model of action is psychologically unrealistic because it overestimates the capacity and willingness of actors to calculate the long-term cumulative benefits against the short-run advantage of “hit and run.” For example, Simon (1992, p. 36) regarded strategic rationality as at best a prescriptive model of how choices should be made, but one that bears little resemblance to actual decision making. These criticisms have led researchers to include “bounded” rationality explicitly in rational choice explanations. Examples are approaches that take into account imperfect information processing or models that maximize utility only in the short term and fail to anticipate long-term future consequences (e.g., Fudenberg and Levine, 1998).

Theories of learning and framing assume at most a bounded rationality and thus provide microfoundations for cooperative reciprocity that do not rely on heroic assumptions about perfect rationality and full information. Instead, learning and framing assume adaptive heuristics or “rules of thumb” that impose relatively small cognitive demands compared to the assumptions in analytical game theory (Orbell and Dawes, 1991). These heuristics have been identified through experimental analyses of human decision making. With these heuristics, Orbell and Dawes argue, real decision makers “economize on cognitive effort” (1991, p. 517) rather than pursue perfectly rational solutions.

Parallel with the development of models of bounded rationality, criticism of rational choice explanations led game theorists to explore evolutionary alternatives. Studies in evolutionary game theory (e.g., Axelrod, 1984; Maynard-Smith, 1982) avoid the need to assume that individual actors have highly sophisticated cognitive abilities. The

optimizing mechanism operates not at the level of individual cognition, but at the population level, through competitive pressures that favor the survival and replication of behavioral strategies that are successful relative to the population average. In a celebrated computer tournament, Axelrod (1984) showed that strategies based on “tit for tat” (a rule to reciprocate cooperation with cooperation and cheating with cheating) were far more successful than more aggressive or predatory mutants. Tit for tat succeeds because it never cheats and it never tolerates cheating by others. It therefore receives the long-term benefits of ongoing mutual cooperation while minimizing its vulnerability to cheaters. Cheaters would do well in a population of naive cooperators, but by driving the latter to extinction, they dig their own graves. In a population of reciprocators, cheating triggers retaliation. Thus, the short-term benefit of cheating cannot keep up with the long-term benefits of mutual cooperation enjoyed by reciprocators.

Critics of evolutionary explanations based on natural selection have pointed out that genetic replication and selection may be a misleading template for models of adaptation at the cognitive level (Aunger, 2001). A central problem is that behavioral strategies for exchange in long-term relationships are not simply “hardwired” programs that successful actors automatically pass on to their biological offspring. Adherents of evolutionary psychology argue that the effects of natural selection on contemporary human behavior may be much more indirect (cf. Cosmides and Tooby, 1992). These authors emphasize that natural selection requires a long time span with stable environmental conditions to effectively shape the genetic basis of human behavior. Moreover, genetic predispositions do not carry detailed information about which behavioral response the organism should choose in a given situation. Instead, genetic programs may elicit certain emotional responses (e.g., anger) to certain situational cues (e.g., failure of partner to reciprocate). Within these constraints, genetic dispositions still leave room for extragenetic behavioral change, for example, based on learning or conscious deliberation.

Theories of learning and framing address this gap in evolutionary models of adaptive behavior. In evolution, strategies compete *between* the individuals that carry them, not *within* them. That is, evolutionary models explore changes in the global frequency distribution of strategies across a population. By contrast, models of learning and framing operate on the local probability distribution of strategies within the repertoire of each individual member. Put differently, these models provide a microfoundation for the extragenetic behavioral change that is missing in evolutionary approaches. While evolutionary theory explains long-term cognitive developments such as the human capacities for learning and framing, these capacities in turn complement evolutionary theories by explaining more fine-grained adaptive responses to short-term changes in the environment.

To sum up, both orthodox game theory based on assumptions of perfect rationality as well as an evolutionary alternative have been widely used to explain strategies of reciprocal solidarity in social exchange. Relentless criticisms of these game theoretic approaches have motivated interest in both framing and learning as alternative explanations of reciprocal solidarity (cf. Lindenberg, this volume; Lindenberg et al., this volume; Macy and Flache, 2002). However, these two approaches have developed in parallel, with no attention to their theoretical differences and similarities. It is to that question we now turn.

Framing and Learning: Assumptions Compared

The behavioral assumptions in framing and learning theories occupy a similar niche between evolution and full rationality but are otherwise very different from one another. An extensive account of framing explanations of solidary behavior has been given in Chapters 1 and 2 of this volume. Reciprocity in exchanges can be seen as an instance of weak solidarity. Under conditions of weak solidarity, participants in an exchange feel legitimated to pursue the goal of improving their personal resources through the exchange (Lindenberg, 2001 and this volume), but compliance with social norms remains a salient secondary goal. That is, solidarity is not so strong a goal that individuals are willing to sacrifice resources without the expectation of reciprocity. But the goal is salient enough that moderate temptations to cheat are ignored and occasional failures to reciprocate (perhaps due to mishaps) are forgiven when excuses have been made.

Framing theory specifies relational signals as a crucial mechanism that stabilizes weak solidarity against the continuous “pull” from opportunistic temptation (Lindenberg, 1998). Relational signaling requires that an occasional failure to reciprocate be accompanied by an unambiguous signal that no cheating was intended. In a long-term exchange, such a signal might, for instance, be given through temporary unconditional cooperation by the party that violated the norm. In addition to relational signals, framing theory posits loss avoidance as a mechanism that safeguards against opportunism. As Lindenberg (2001) argues, when an actor feels threatened by a severe loss (for example, when cheating by a long-term exchange partner may be particularly costly), this may trigger a loss frame in which avoidance of the loss dominates normative or gain-oriented motives. Actors’ anticipation of such a frame switch, in turn, may stabilize cooperation, particularly in situations where unsolidary behavior may elicit retaliatory responses that can inflict severe losses on the transgressor. For example, Mühlau (2000, p. 211) points out in an analysis of framing effects in organizational governance that “the higher the damage potential the

other party controls, the more willing an actor will be to bear the costs associated with relational obligations” (cf. Lindenberg, 1988). Translated into social exchanges, this loss avoidance implies that cooperative reciprocity will be particularly stable when the loss of resources obtained from the exchange partner could be highly damaging for a participant.

In sum, framing theory does not neglect self-interested motives and leaves ample room for goal-directed individual choice. However, the model also posits a “cognitive miser” (Orbell and Dawes, 1991) in that it assumes that the complexity of individuals’ decision making in most decision-making situations is greatly reduced by a focus on one foreground goal at a time.

Like framing theory, learning theory also relaxes key behavioral assumptions of the orthodox rational choice approach without disregarding self-interest and goal-driven decision making. There are three key differences with analytical game theory based on standard rational choice:

- Propinquity replaces causality as the link between choices and payoffs.
- Reward and punishment replace utility as the motivation for choice.
- Melioration replaces optimization as the basis for the distribution of choices over time.

Propinquity, not Causality

Compared to analytical game theory, the Law of Effect imposes a lighter cognitive load on decision makers. It assumes experiential induction of the future consequences of actions that were previously encountered. By contrast, rational behavior assumes logical deduction of actions that may never have been experienced. In learning theory, players develop preferences for those actions associated with better outcomes in the past, even though the association may be coincidental, “superstitious,” or causally spurious.

Reward and Punishment, not Utility

Learning theory differs from game-theoretic utility theory in that it posits two distinct cognitive mechanisms that guide decision makers toward better outcomes: *approach* (driven by reward) and *avoidance* (driven by punishment). The distinction means that aspiration levels are very important for learning theory. The effect of an outcome depends on whether it is coded as gain or loss, satisfactory or unsatisfactory, pleasant or aversive.

Melioration, not Optimization

Melioration implies a tendency to repeat choices with satisfactory outcomes even if other choices have higher utility, a behavioral tendency March and Simon (1958) call “satisficing.” A good example is the decision whether to cooperate in an ongoing exchange. Melioration can imply that each side is satisfied with its current choice when the partner cooperates and dissatisfied when the partner defects. Unsatisfactory outcomes increase the probability that alternative actions will be taken, including a tendency to revisit alternative choices whose outcomes are even worse, a pattern we call “dissatisficing.”

While the three learning principles may describe decision making that is suboptimal by conventional game-theoretic criteria, they may be more effective in leading actors out of social traps than more sophisticated decision-making rules. The outcomes of the exchange that are regarded as rewards, such as a successful and mutually profitable business transaction, induce approach behavior, the tendency to repeat the associated choices even if other choices have higher utility. In contrast, outcomes that are coded as punishments, such as being cheated by an exchange partner, induce avoidance. Taken together, approach and avoidance imply the possibility that reciprocal solidarity may become self-reinforcing in exchange relations. As long as participants are sufficiently satisfied with mutual cooperation (approach) and they are sufficiently dissatisfied with failure to exchange (avoidance), learning dynamics may lead actors to engage in reciprocal cooperation.

Comparison of Theory Implications

Both framing and reinforcement learning theories predict behavior that corresponds with reciprocal solidarity. But do the models also make similar predictions about the conditions and dynamics of reciprocal solidarity? We found remarkable overlap between the two sets of predictions, but with one interesting exception. Framing suggests that frequent norm violations lead inevitably to irreversible deterioration of the exchange. In contrast, learning theory implies that recovery is possible, even when actors adapt their aspirations to recent experience.

To compare the implications of framing and learning theories, we used formal games as stylized representations of strategic interdependence in ongoing exchanges. Game theory has formalized the problem of cooperation at the most elementary level as a mixed-motive two-person game with two choices: cooperate and defect. These choices intersect at four possible outcomes, abbreviated as CC, CD, DD, and DC. Each outcome has an associated payoff: R (reward), S (sucker), P (punishment), and T (temptation), respectively. Using these payoffs, we defined

a two-person social dilemma as any ordering of these payoffs such that mutual cooperation is collectively optimal yet may be undermined by the temptation to cheat (if $T > R$) or by the fear of being cheated (if $P > S$), or by both. In the game of Stag Hunt, the problem is “fear” but not “greed” ($R > T > P > S$), and in the game of Chicken, the problem is “greed” but not “fear” ($T > R > S > P$). The problem is most challenging when both fear and greed are present, that is, when $T > R$ and $P > S$. Given the assumption that $R > P$, there is only one way this can happen: if $T > R > P > S$, the celebrated game of Prisoner’s Dilemma (PD).

For social exchange situations, the games of Stag Hunt, Chicken, and PD correspond to different forms of interdependence. As Rousseau (who invented the game) noted, Stag Hunt games may arise when contributions by all participants are necessary to produce a common good that everyone values highly. The prototypical example is a work team that has the opportunity to receive a substantial bonus payment for all team members, but only when a production target is met for which great effort on the part of all members is needed. The Chicken game models a situation where the bonus may be obtained if at least some members shoulder the burden, but it will certainly be lost if at least a certain fraction of the group fails to pull its weight. Group members may prefer to free-ride, but if they feel that the bonus may be lost due to others free-riding, they “give in” and work hard to avoid the worst. Finally, in the PD game, contributions do not sufficiently reduce the chances of obtaining the bonus to compensate for the cost of effort. Even in PD, however, universal defection is suboptimal, because all group members prefer to work and get the bonus than lose it because of universal free-riding.

To compare predictions for long-term exchange relationships, we assume that actors in an exchange relation play the underlying game repeatedly and learn the outcomes after every round of mutual decision making. For simplicity, we further confine our analysis to symmetrical games in which the payoffs R , T , P , and S are equal for both players. With respect to learning theory, we draw on results that we elaborated elsewhere in formal computational experiments (Flache and Macy, 2002; Macy, 1991; Macy and Flache, 2002). In these studies, we used the Bush-Mosteller stochastic learning model, a mathematical formalization of reinforcement learning. Figure 3.1 provides a schematic overview of the learning mechanism in our computational model.

FIGURE 3.1. The Reinforcement Learning Mechanism

The first step in Figure 3.1 is the decision made by each player whether to cooperate or defect. This decision is probabilistic, based on the player's current propensity to cooperate. The resulting outcome then generates payoffs (R, S, P, or T) that the players evaluate as satisfactory or unsatisfactory relative to their aspiration levels. Satisfactory payoffs present a positive stimulus (or reward) and unsatisfactory payoffs present a negative stimulus (or punishment). The stimulus modifies the probability of repeating the associated choice, such that satisfactory choices become more likely to be repeated, while repetition of unsatisfactory choices becomes less likely. For a formal specification of learning dynamics, we refer interested readers to our previous publications (e.g., Flache and Macy, 2002; Macy and Flache, 2002).

Unlike reinforcement learning theory, framing theory has not yet been fully formalized in such a way that model dynamics are directly comparable. To make a comparison possible, we derived from framing theory informally stylized facts about the effects of game structures and game parameters. We compared these facts to the implications of a computational model of reinforcement learning.

Framing Predictions

The first step is to make assumptions about players' frames at the outset of a repeated game, when the relationship does not yet have a history. Framing theory assumes that decision making is "forward-looking" in the sense that actors' mental images of a relationship shape their initial behavior and aspirations. Mental images, in turn, depend on the social context of a relationship (see Lindenberg et al., this volume). This highlights an important theoretical difference between framing and reinforcement learning: While assumptions about the initial perceptions of the relationship are endogenous in framing theory, reinforcement learning theory treats initial aspirations and behavior as exogenously given and independent of the particular situation. For our analysis of the three abstract social dilemma games, however, the game structures as such do not provide information about the social context of the interaction. Accordingly, for a framing analysis, we need to use assumptions that are exogenous to the theory. Such assumptions are drawn from experimental data on social dilemma games.

Experimental data about social dilemma games seem to be most consistent with the assumption that subjects frame the exchange situation initially in terms of weak solidarity (cf. Davis and Holt, 1993). Two highly robust results from the literature support this interpretation. First, across a wide range of social dilemma games, subjects exhibited in experiments have a large proportion of cooperative choices in the first iterations of repeated games, an observation that is at odds with

the notion that subjects are exclusively gain oriented. At the same time, cooperation rates steadfastly declined over time in the experiments (Andreoni, 1988), a clear indication that the solidarity motive at best tempers but does not dominate gain considerations.

Given initial weak solidarity, we find three stylized facts implied by framing that can be compared with learning predictions. First, framing theory suggests that cooperation rates will be lowest in PD, highest in Stag Hunt, and between these extremes in Chicken. The payoff inequalities of PD, Chicken, and Stag Hunt differ in two dimensions that are salient for the framing explanation: the temptation to engage in opportunism and the importance of loss avoidance. The greater the temptation to engage in opportunism, that is, the larger the gains that an actor may attain when he or she unilaterally deviates from reciprocal solidarity, the more salient the motive of gain relative to the normative frame that stabilizes mutual cooperation. The temptation to unilaterally defect from mutual cooperation is lowest in Stag Hunt, where players prefer mutual cooperation to cheating, and is similar in Chicken and PD, where both players prefer exploitation of the partner to mutual cooperation. Loss avoidance works in the opposite direction. The threat of loss to a player owing to deterioration of the relationship may actually strengthen solidary behavior. Such losses are highest in Chicken (where mutual defection is the least preferred outcome), and they are higher in Stag Hunt than in PD (because the difference between mutual cooperation and mutual defection tends to be larger in Stag Hunt). Taken together, from a framing perspective, conditions for reciprocal solidarity are least favorable in PD (high temptation, low loss from opportunism), and they are most favorable in Stag Hunt (low temptation, medium loss), with the Chicken game between these extremes.

The second stylized fact implied by framing theory is the gradual decline of cooperation rates over time in games with a high temptation to defect (PD and Chicken). Lindenberg (1998) argues that ongoing exposure to such temptation may gradually weaken actors' normative frames such that, at some point, the relationship may "turn sour" and degrade into mutual defection. He also points out that consistent and repeated relational signals of cooperative intentions by both parties may prevent the decline. In the simple social dilemma games that we analyzed, however, the only interaction between players was in their decision to cooperate or defect. Hence, the only relational signal an actor can give after occasional unilateral cheating is subsequent unconditional cooperation. Clearly, the same temptation that leads an actor to cheat in the first place may also prevent the actor from giving this costly signal. Accordingly, framing theory suggests that exchange relations have a tendency to eventually degrade into mutual defection, more so in games with higher temptation to defect (PD, Chicken) and less so in Stag Hunt.

Clearly our list of stylized facts derived from framing is far from exhaustive and may be extended in future research. For the predictions we derive here, we explored whether similar conditions and dynamics of reciprocal solidarity would be obtained from reinforcement learning.

Framing Predictions and a Computational Model of Learning

We use a set of learning assumptions that make reciprocal solidarity based on stochastic collusion possible but not trivial. We assume that at the outset of the games players will randomize between cooperation and defection, reflecting the assumption that no stimuli have yet been experienced that favor choices in one direction or the other. Furthermore, we set the rate of behavioral change following stimuli relatively high, approaching a “win-stay, lose-change” heuristic, in which choices are always repeated when rewarded and always changed to the alternative (C or D) when punished. To formalize the three games, we use payoffs ordered from the set [4, 3, 1, 0] for each of the three social dilemma payoff inequalities. We assume an aspiration level of $A = 2$ that corresponds to the payoff expected when behavioral propensities are uninformed by prior experience and all players randomize such that all four payoffs are equiprobable. With this aspiration level, mutual cooperation is the unique outcome in all three games that simultaneously satisfy both players. There is no guarantee, however, that mutual cooperation will arise, since players will also be punished for cooperation should the partner defect, and they will be rewarded for defection should the partner cooperate (Flache and Macy, 2002; Macy and Flache, 2002).

Figure 3.2 shows single replications of the learning dynamics that we obtained in all three social dilemma games using these baseline assumptions. The figure charts the change in the probability of cooperation (PC) for one of two players with statistically identical probabilities.

FIGURE 3.2. Stochastic Collusion in Three Social Dilemma Games($p = [4,3,1,0]$, $A = 2$, initial probability of cooperation = 50%, high learning rate)

Figure 3.2 shows that dissatisficing learning players initially wander about in an unstable equilibrium with a low probability of cooperation, but eventually escape the social trap by random walk (or what we call “stochastic collusion”). The figure also reveals differences between the games. Mutual cooperation stabilizes most readily in Stag Hunt and least readily in Prisoner’s Dilemma. To test the robustness of this difference, we performed 1,000 replications of the experiment and measured the proportion of runs that stabilized on mutual cooperation within 250 iterations. The results confirmed the differences between the games. These differences reflect subtle but important interactions between aspiration levels and the type of social dilemma—the relative importance of fear (the problem in Stag Hunt) and greed (the problem in Chicken).

The findings also show that satisficing is equally important, at least in the Prisoner’s Dilemma and in the Chicken game. In these games, appreciation that the payoff for mutual cooperation is “good enough” motivates players to stay the course despite the temptation to cheat (given $T > R$). Otherwise, mutual cooperation would not be self-reinforcing. In Stag Hunt, satisficing is less needed in the long run, because there is no temptation to cheat ($R > T$). Despite the absence of greed, however, the findings reveal that, even in Stag Hunt, fear may inhibit stochastic collusion if high aspirations limit satisficing.

With respect to framing theory, the results shown in Figure 3.2 demonstrate that our simple reinforcement learning model generates the same qualitative differences between games as are predicted by the first stylized fact we derived from framing theory. Interestingly, the underlying mechanisms also seem very similar. In learning, it is the proper balance between the punishment for defection and the reward for cooperation that drives the emergent reciprocity in exchange relations. In framing, reciprocity thrives on the proper balance between resistance to the temptation to cheat and the motivation to avoid losses caused by mutual sanctioning.

However, when we turn to the second stylized fact derived from framing—gradual decline of cooperation—we find a clear difference between the predictions. As Figure 3.2 shows, reinforcement learning implies a robust tendency of exchange relationships to recover from occasional violations of the reciprocity norm, even when these violations are quite frequent, as, for example, between iterations 20 and 60 of the PD experiment. Framing, on the other hand, suggests that after too many violations of normative expectations, relations decay into mutual defection without the possibility to recover (cf. Lindenberg, 1998). The latter pattern seems more consistent with experimental results from social dilemma games (Andreoni, 1988) than the consistent recovery predicted by learning models.

To further test this difference between the theories, we added to the learning model an additional learning principle that may explain,

from a learning perspective, why reciprocal solidarity can become unstable. This learning principle, called “habituation” (Sokolov, 1963), assumes a decline in the tendency to respond to stimuli that have become familiar through repeated exposure. Technically, we operationalize habituation as the tendency to adapt aspiration levels to experienced payoffs. Habituation can lead to desensitization to a recurrent stimulus, whether reward or punishment, and to increased sensitivity to change in the stimulus. Thus, habituation to reward decreases sensitivity to further reward but increases sensitivity to punishment. We model habituation as the tendency of aspirations gradually to float toward the average payoff experienced in recent interactions. In addition, we assume the same start conditions for the experiment as before. Figure 3.3 shows the results.

Figure 3.3 shows the destabilizing effects of habituation on the learning dynamics. All three graphs show that cooperative reciprocity eventually obtains, as in the baseline experiment shown in Figure 3.2. However, unlike the earlier experiment, we now see that cooperation soon destabilizes and deteriorates. Consistent with the differences between games that we found in the first experiment, cooperative periods seem to be shortest in PD and cooperation seems to be more stable in Chicken and Stag Hunt. This pattern was confirmed using statistical tests. The dynamics for Chicken also show that in this game the strong punishment for mutual defection serves to suppress habituation in favor of the social costs of failure to exchange. As Figure 3.3 shows, the Chicken dynamics reveal no periods of stable mutual defection, unlike in Stag Hunt or PD, where habituation may make players temporarily immune to the low payoffs associated with the PD outcome.

The dynamics of habituation in the learning model resemble the pattern suggested by framing theory, but only to a point. Both models predict that cooperative reciprocity eventually degrades into mutual defection. Moreover, consistent with the second stylized fact we derived from framing theory, this decline seems to be more frequent in

FIGURE 3.3. Change in PC Over 500 Iterations With Floating Aspirations ($p = [4, 3, 1, 0]$, initial $A = 2.0$, high learning rate, initial cooperation rate $PC, 1 = 0.5$)

PD (but not in the Chicken game) than in the game of Stag Hunt. A clear difference, however, is that, according to framing theory, this decline is irreversible after some point. Learning theory predicts that, after the decline, habituation will lead to increased sensitivity of the players to the rewards associated with mutual cooperation. As a consequence, learning actors have the ability to recover reciprocity even after long periods of exchange failures and even without the possibility to exchange relational signals other than cooperation or defection.

Discussion and Conclusion

Both framing theory and learning theory have three important properties that render them attractive as explanations of reciprocity in exchange relations. First, they do not trivialize the problem of opportunism. Second, they take into account individual discretion to deviate from normative obligations or genetic programs. Third, they avoid heroic assumptions about individual cognitive capacities and perfect information.

In order to compare the two approaches, we derived from framing theory stylized predictions about the dynamics of reciprocity in repeated 2×2 games and compared these to the implications of a computational model of stochastic learning in identical games. We conclude from the results that a simple learning model can explain two key observations about solidarity in ongoing exchanges that are consistent with framing predictions:

- stable ongoing reciprocity despite occasional mishaps and moderate rewards for opportunistic behavior;
- deterioration of reciprocity relationships as a consequence of strong rewards for opportunism or habituation to the rewards for mutual cooperation.

Our analysis also revealed testable differences between the learning model and the framing approach in their predictions about recovery of reciprocity from collapse of mutual cooperation. Although framing suggests that norm violations eventually lead to irreversible collapse, learning theory implies that recovery is possible if the learning rate is sufficiently high. With a low learning rate, however, the learning model also predicts difficulty recovering from the collapse of mutual cooperation. Laboratory experiments are needed to test the relative explanatory power of the two theories as the rates of learning and habituation are manipulated to generate discrepant predictions.

Although we have not explored all possible implications of the two theories, we tentatively conclude that learning theory may provide a more parsimonious explanation of the dynamics of ongoing reciprocity, based on elementary principles that remain largely

implicit in framing theory. Conversely, framing theory addresses explicitly how actors' perceptions of the exchange situation may shape their initial aspirations and behavioral propensities, a crucial element for relational dynamics that is left exogenous in reinforcement learning. Clearly, each theory may benefit from a more explicit elaboration of those elements that are underspecified in its counterpart.

References

- Andreoni, J. (1988). Why free ride? Strategies and learning in public goods experiments. *Journal of Public Economics*, 37, 291–304.
- Aunger, R. (2001). *Darwinizing culture: The status of memetics as a science*. Oxford: Oxford University Press.
- Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
- Bandura, A. (1977). *Social learning theory*. Englewood Cliffs, N.J.: Prentice Hall.
- Buunk, B.P., and Schaufeli, W.B. (1999). Reciprocity in interpersonal relationships: An evolutionary perspective on its importance for health and well-being. In W. Stroebe and M. Hewstone (Eds.), *European review of social psychology* (Vol. 10, pp. 259–291). Chichester: Wiley.
- Coleman, J.S. (1990). *Foundations of social theory*. Cambridge, Mass.: Harvard University Press.
- Cosmides, L., and Tooby, J. (1992). Cognitive adaptations for social exchange. In J.L. Barkow, L. Cosmides, and J. Tooby (Eds.), *The adapted mind* (pp. 163–228). Oxford: Oxford University Press.
- Davis, D.D., and Holt, C.A. (1993). *Experimental economics*. Princeton, N.J.: Princeton University Press.
- Flache, A., and Macy, M.W. (2002). Stochastic collusion and the power law of learning: A general reinforcement learning model of cooperation. *Journal of Conflict Resolution*, 46(50), 629–653.
- Friedman, J.W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38, 1–12.
- Fudenberg, D., and Levine, D. (1998). *The theory of learning in games*. Boston: MIT Press.
- Gouldner, A. (1960). The norm of reciprocity: A preliminary statement. *American Sociological Review*, 25, 161–178.
- Ligthart, P.E.M. (1995). *Solidarity in economic transactions. An experimental study of framing effects in bargaining and contracting*. Amsterdam: Thesis Publishers.
- Lindenberg, S.M. (1988). Contractual relations and weak solidarity: The behavioral basis of restraints on gain maximization. *Journal of Institutional and Theoretical Economics*, 144, 39–58.
- Lindenberg, S.M. (1998). Solidarity: Its microfoundations and macro-dependence. A framing approach. In P. Doreian and T.J. Fararo (Eds.), *The problem of solidarity: Theories and models* (pp. 61–112). Amsterdam: Gordon and Breach.
- Lindenberg, S.M. (2001). Social rationality versus rational egoism. In J.H. Turner (Ed.), *Handbook of sociological theory* (pp. 635–668). New York: Kluwer.
- Macy, M.W. (1991). Learning to cooperate: Stochastic and tacit collusion in social exchange. *American Journal of Sociology*, 97, 808–843.
- Macy, M.W., and Flache, A. (2002). Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences USA*, May 14, 99(10), 7229–7236.
- March, J.G., and Simon, H.A. (1958). *Organizations*. New York: Wiley.
- Maynard-Smith, J. (1982). *Evolution and the theory of games*. Cambridge: Cambridge University Press.

- Mühlau, P. (2000). *The governance of the employment relation. A relational signaling perspective*. Amsterdam: Thela Thesis.
- Orbell, J.M., and Dawes, R.M. (1991). A “cognitive miser” theory of cooperators’ advantage. *American Political Science Review*, 85, 515–528.
- Rapoport, A., and Chammah, A.M. (1965). *Prisoner’s dilemma: A study in conflict and cooperation*. Ann Arbor: Michigan University Press.
- Scharpf, F.W. (1990). Games real actors could play. *Rationality and Society*, 2, 471–494.
- Simon, H. (1992). Decision making and problem solving. In M. Zey (Ed.), *Decision making: Alternatives to rational choice models* (pp. 32–53). London: Sage.
- Sokolov, Y.N. (1963). *Perception and the conditioned reflex*. Oxford: Pergamon.
- Thorndike, E.L. (1911). *Animal Intelligence*. New York: Macmillan.