2 Conditional Cooperation: Behavioral Regularities from the Lab and the Field and Their Policy Implications

Simon Gächter

2.1 The Problem of Voluntary Cooperation

A well-known fact from the theory of public goods is that voluntary provision will lead to an inefficient undersupply (Samuelson 1954). The reason is the famous free rider problem: since, by definition of a public good, an agent can benefit from it even if he or she has not contributed to it, everyone has an incentive to hope that others will provide the public good. More specifically, a rational and selfish agent will equate only his or her private marginal benefits to the marginal costs of the public good, whereas efficiency requires that the sum of marginal benefits should equal the marginal costs. Thus there exists a tension between individual and collective rationality, which is prototypical for many cooperation problems. This tension lies at the heart of the matter in such diverse areas as warfare, environmental protection, management of commons, tax compliance, corruption, voting, participation in collective actions like demonstrations and strikes, donations to charities, teamwork, collusion between firms, embargoes and consumer boycotts, and so on.

While the logic of self-interest is straightforward, the data seem to be at odds with the free rider hypothesis derived under the joint assumptions of rationality and selfishness. The fact that people vote even in anonymous situations, take part in collective actions, often do not overuse common resources, care for the environment, mostly do not evade taxes on a large scale, and donate to public radio and charities suggests that the strict self-interest hypothesis is inconsistent with the degree of voluntary cooperation we observe around us.

How can we explain this? What are the implications for public policy and management? This chapter outlines some possible answers to both these questions. Our main sources of information are controlled

laboratory and field experiments.[1] As I will show in this chapter, the main finding from a large body of experiments conducted in a variety of settings in the last three decades is that there is much more cooperation than predicted by standard theory (Ledyard 1995). Yet the experiments also show that voluntary cooperation is fragile in the sense that in repeatedly played public goods games cooperation declines over time.

How can we explain (the fragility of) voluntary cooperation? One important explanation is that people have "warm-glow" preferences; in other words, they have some positive utility simply from the act of contributing (e.g., Andreoni 1990). A second explanation is that many people have altruistic preferences—they want to benefit others. A third reason is errors—people make mistakes (e.g., Anderson, Goeree, and Holt 1998). In a clever design, Palfrey and Prisbrey (1997) test for warm-glow, altruism, and errors and find that altruism does not explain contributions, but some people have warm-glow preferences. Errors are important as well and explain why in repeated experiments contribution rates typically decline.

It should be noted that both motives—altruism and warm glow—are independent from other people's cooperation behavior. A set of recent experiments has cast doubt on this assumption. A large number of people are "conditionally cooperative"—they cooperate if they believe others cooperate as well. Yet a significant fraction of people is best characterized as free riders. In summary, recent evidence suggests that there is considerable heterogeneity with respect to people's cooperation preferences; in other words, there are types of players.

In section 2.3 I take up the issue of preference heterogeneity and discuss four of its predicted consequences:

(1) *Voluntary cooperation is fragile* This holds in particular without further institutional remedies, like possibilities for communication, punishment, or assortative interactions. The reason is that conditional cooperators who experience free riding will stop cooperating themselves.

(2) *Social interaction effects exist in voluntary cooperation* This means conditional cooperators will adapt their behavior to the group they are in. If other group members shirk, they shirk as well; if others cooperate, they cooperate as well. These social interaction effects mean that people's behavior is influenced by their group.

(3) *Group composition with respect to types matters for voluntary cooperation* For instance, if conditional cooperators know the other group members are cooperators as well, then they should be able to maintain high cooperation levels. The team spirit of like-minded cooperators should suffice to maintain high cooperation. Similarly, free riders who know that others are free rider types as well are predicted to defect.

(4) *Belief management matters for voluntary cooperation* Conditional cooperators cooperate by definition, if they believe others cooperate as well. Hence, any factor influencing beliefs will affect cooperation behavior.

I present evidence from new experiments designed to test these predictions. The evidence from these experiments unequivocally supports the importance of conditional cooperation and preference heterogeneity in understanding cooperation behavior. I see the experiments as behavioral models that may help us understand important field phenomena. In section 2.4 I therefore interpret field evidence on tax evasion, bribery, welfare fraud, attitudes toward the welfare state, charitable giving, and work morale in the light of the four behavioral models.

These findings on the importance of conditional cooperation and preference heterogeneity have consequences for theory and policy. If people are largely motivated by warm-glow preferences, and if the decay in contributions is due to reduced errors, then the modeling approach might be different than if people were free riders or conditional cooperators whose interaction explained the decay in contributions. In the former case, a modeling approach where errors figure prominently might be the preferable one (see, e.g., Anderson, Goeree, and Holt 1998). In the latter case, a theory of social preferences might be chosen (see, e.g., Camerer 2003; Fehr and Schmidt 2003; and Sobel 2005 for surveys of models, and Tyran and Sausgruber 2006 for a policy application). The findings also have consequences for public policy and management. I discuss them in section 2.5. Section 2.6 concludes.

2.2 Conditional Cooperation in the Lab and the Field

I start by presenting some stylized facts from laboratory experiments (section 2.2.1). This will only be a sketch and the interested reader may wish to consult Ledyard (1995) and Gächter and Herrmann (2005)

for more complete accounts of important results from economic experiments. Dawes (1980) discusses evidence from social psychological experiments. I will discuss recent field experiments that are consistent with the lab findings in section 2.2.2. Section 2.2.3 presents evidence that behavior in the lab is consistent with naturally occurring field behavior.

2.2.1 Evidence from the Laboratory

The linear public goods game (or voluntary contribution mechanism) has proved extremely useful for testing the free rider hypothesis in the lab. In a typical linear public goods experiment, n people form a group. All group members are endowed with z tokens. Each subject i has to decide independently how many tokens (between 0 and z) to contribute to a common project (the public good). The contributions of the whole group are summed up. The experimenter then multiplies the sum of contributions by $\alpha > 1$ and distributes the resulting amount equally among the four group members. Thus each subject i's payoff is

$$\pi_i = z - g_i + \frac{\alpha}{n}\sum_{j=1}^{n} g_j, \quad j = 1, \ldots, n,\ \alpha > 1,\ \alpha/n < 1. \tag{1}$$

The first term $(z - g_i)$ indicates the payoff from the tokens not contributed to the public good (the "private payoff"). The second term is the payoff from the public good. Each token contributed to the public good becomes worth $\alpha > 1$ tokens. The resulting amount is distributed equally among the n group members—irrespective of how much an individual has contributed. Thus an individual benefits from the contributions of other group members, even if he or she has contributed nothing to the public good. A rational and selfish individual therefore has an incentive to keep all tokens for him- or herself, since his or her return per token from the public good is only $\alpha/n < 1$, whereas it is 1 if he or she keeps the token. By contrast, since $\alpha > 1$, the group as a whole is best off if everybody contributes all z tokens.

Figure 2.1 depicts a typical finding of a public goods experiment where the exact same game is repeated ten times. Subjects, who play in groups of four, know about the repetition. In each period each subject receives 20 tokens and decides how many of them to keep or contribute to the public good. After each round subjects are informed about what the other three group members have contributed. Figure 2.1 shows the resulting cooperation patterns in a "stranger" condition,

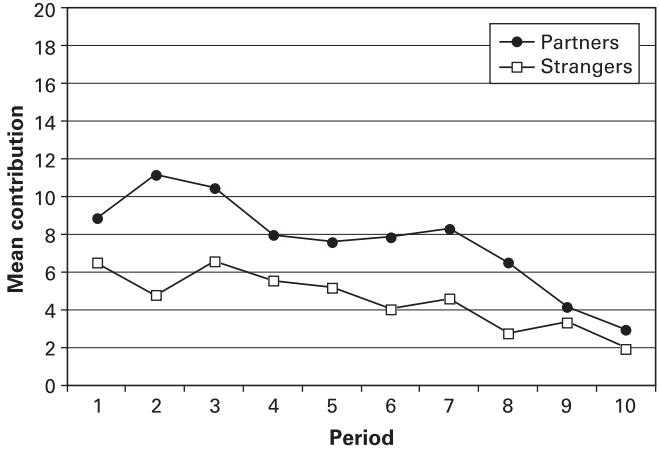

Figure 2.1
Contributions to a public good in constant groups (partners) and randomly changing groups (strangers) over ten repetitions. *Source*: Fehr and Gächter (2000).

where group members change randomly from round to round, and a "partner" condition, in which groups stay constant for all rounds.

Figure 2.1 illustrates two stylized facts from dozens of public goods experiments. First, people contribute substantially more than theoretically predicted. In most experiments, partners contribute more than strangers (see Keser and van Winden 2000 and Andreoni and Croson 2008 for an overview). The significance of this and related findings is that people are immediately able to distinguish whether they are in a situation requiring strategic cooperation (the partner condition) or not (the stranger condition) and to adapt their behavior accordingly.

The second stylized fact is that cooperation is very fragile and tends to collapse with repeated interactions. Why is this so? One explanation is that people have altruistic or warm-glow preferences, but also have to learn how to play this game. Since errors can only go in one direction, any erroneous decision looks like a contribution. Palfrey and Prisbrey (1997) test these explanations and find that the data are inconsistent with altruism. They find some evidence for warm-glow preferences but also conclude that people learn and commit fewer errors over time, which is why contributions decline.

Notice that warm glow, altruism, and errors are motivations that are independent of others' contributions. Psychologists have long argued that people's cooperation behavior depends on what others do (e.g.,

Kelley and Stahelski 1970). Using the methodology of experimental economics, Keser and van Winden (2000) were among the first economists to argue for the prevalence of conditional cooperation. Croson (2002) went one decisive step further by eliciting beliefs about other group members' contributions. She found a very high and statistically significant correlation of beliefs and contributions: subjects who expected others to contribute a lot were more likely to contribute high amounts than were subjects who expected others to free ride. This observation clearly suggests that people's contribution behavior is *not* independent of what they expect others to do. Thus, Croson's findings are consistent with conditional cooperation.

Croson (2002) did not look at individual behavior. Her observation is that, on average, people behave conditionally cooperatively in that their contributions and beliefs are positively correlated. Fischbacher and Gächter (2006) also elicited beliefs and replicated Croson's finding of a positive correlation between beliefs and contributions. At the individual level they find subjects who show a positive correlation between beliefs and contributions, whereas other subjects contribute zero even if they believe that others contribute positive amounts.

There are at least three problems with using the correlation between beliefs and contributions as an indicator of conditional cooperation. First, beliefs evolve endogenously in the experiment and are thus beyond the control of the experimenter. Second, a free rider who believes others contribute zero and actually contributes nothing him- or herself is observationally equivalent to a pessimistic conditional cooperator who only contributes a little because he or she believes others will free ride. Third, people may project their behavioral tendencies unto others; in other words beliefs may reflect a "false consensus effect" (see, e.g., Kelley and Stahelski 1970; Orbell and Dawes 1993).

Fischbacher, Gächter, and Fehr (2001) and Fischbacher and Gächter (2006) circumvent these problems by using a revealed preference method in their public goods games to infer people's contribution preferences as a function of other group members' contributions. Therefore, the subjects in their experiment do not choose one contribution but a contribution as a *function* of other group members' average contribution. The public goods game is played in groups of four subjects and the payoff function is again the same as in (2.1). The game is played just once to avoid confounds with strategic considerations. Every subject has to indicate a contribution conditional on others' aver-

age contribution; in other words, for each of the twenty-one possible values of the average of others' contribution, subjects have to enter the number of tokens they want to contribute.

Fischbacher, Gächter, and Fehr (2001) and Fischbacher and Gächter (2006) classify their subjects according to their contribution function (for details see their papers). A subject is called a free rider if and only if he or she contributes zero in all twenty-one cases. A subject is called a conditional cooperator if the contribution schedule is a positive function of the others' average contribution. A somewhat peculiar type is the triangle contributor, whose contribution increases when others' contributions are low and decreases for higher levels of others' contributions. Figure 2.2 illustrates the average contribution function of the different types in the experiments of Fischbacher and Gächter.

More than half of all subjects are conditional cooperators. Twenty-three percent are free riders. The rest are either triangle contributors or nonclassifiable others. Fischbacher, Gächter, and Fehr (2001) got a very similar distribution of types and even of average contribution patterns.

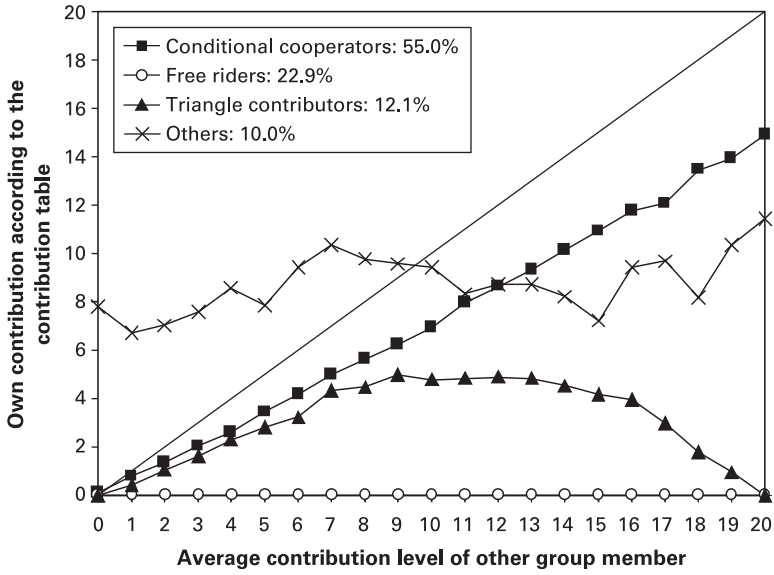

Figure 2.2
Average contribution function of types: Free riders, Conditional cooperators, Triangle contributors, and Others. Observations on the diagonal would correspond to a perfect conditional cooperator. *Source*: Fischbacher and Gächter (2006).

Ockenfels (1999), Bardsley and Moffatt (forthcoming), Burlando and Guala (2005), Muller et al. (2005), Ones and Putterman (forthcoming), and Page, Putterman, and Unel (2005) also find evidence for heterogeneous cooperation preferences in related experimental designs. These studies differ in so many details that a straightforward comparison of the distribution of the different types is not possible. Yet in almost all studies most subjects are classified as free riders or conditional cooperators, with the latter constituting the majority.

In summary, the evidence from the laboratory unambiguously shows that there is much more cooperation than is predicted by standard theory. Moreover, we find strong evidence that many people's attitude toward voluntary cooperation is conditional on other people's cooperation. This suggests that warm glow is not a dominant motivation. Furthermore, many people contribute more the more others contribute. This fact speaks against pure altruism explanations, which predict that people reduce their own contributions when informed that others already contribute to the public good.

A second important finding is that people's contribution preferences are heterogeneous. While a large number of people seem to be conditional cooperators, a significant fraction of subjects is best characterized as free riders. Some others show more complicated patterns. In section 2.3 I will discuss experiments that test directly for implications of preference heterogeneity. Before I do so, I will discuss evidence from the field.

2.2.2 Evidence from Field Experiments

Field experiments offer a great opportunity to test the behavioral relevance of laboratory findings in naturally occurring contexts (see also Harrison and List 2004). In this section I discuss a few field experiments that present results consistent with the lab evidence.

A first interesting study is by Frey and Meier (2004). Their subjects are University of Zurich students. Each semester each student is asked upon registering whether, in addition to the tuition fee, he or she would like to donate to two funds—one that helps needy students with cheap loans, and one supporting foreign students. A donation to the loans fund costs 7 Swiss francs (roughly €4.70), while one to the foreign student fund is 5 Swiss francs. Students can either donate these fixed amounts or not donate; intermediate donations are not possible. The data set comprises 37,624 students. For the field experiment, 2,500 nonfreshmen students were randomly selected; 2,000 of them received

information about what others did. One thousand students received the information that a high percentage of others (64 percent) made a donation in the past; the remaining 1,000 students got the information that a relatively low number (46 percent) made a donation in the past.[2] Using 500 students, Frey and Meier elicited expectations about the fraction of students who make a donation.

The results are consistent with theories of conditional cooperation. First, students who expect a larger number of others to donate are more likely to donate. The correlation between expressed expectations and actual donation is 0.34 ($p < 0.001$). Second, a logit analysis shows that those students who received the information that 64 percent of others had donated in the past are more likely to donate than those who received the information that only 46 percent donated.

Heldt (2005) uses a similar idea as Frey and Meier (2004) to test for conditional cooperation. In his natural field experiment, subjects are tourists who use a cross-country skiing slope. They are then asked to make a donation for the slope's preparation. Heldt also manipulates the information people get. He finds that those who are informed that 70 percent of other tourists donated to the preparation of the slope contributed significantly more than those who did not get that information. Thus this behavior is consistent with conditional cooperation.

The study by Martin and Randal (2005) is similar in spirit. In their natural field experiment, conducted in a museum in New Zealand, visitors could donate to the museum by putting money into a transparent box. The experimenters manipulated whether there was money in the box or not. Consistent with conditional cooperation, they found that people donate significantly more when there is money in the box than when it is empty.

Shang and Croson (2005) conducted a field experiment on donations to a public radio station, which is a naturally occurring public good. The study was similar in spirit to Frey and Meier (2004). In a fundraising drive, people who called in to make a donation (to renew their membership) were confronted with what others had donated in the past. Specifically, in the experimental condition (but not in the control condition) the experimenter read the following sentence: "We had another member, they contributed $75 [$180 or $300]," and right after that "How much would you like to pledge today?" Then the callers could make their pledge (any amount they wished). In total, 538 members called to make a donation. The benchmark for donation decision is the previous year's fund drive, in which the average amount

donated was \$135 and the median amount, \$75. The amounts used as the treatments correspond to the 50th percentile (\$75), the 85th percentile (\$180), and the 90th percentile (\$300) in the previous fund drive. The results again support conditional cooperation. Callers who were confronted with a previous pledge of \$300 donated significantly more than people in the control condition who were not confronted with that information; callers who received the \$75 or \$180 information, respectively, also contributed more than the control group, but this effect is not significant.[3]

In summary, the results from field experiments support the importance of conditional cooperation in the field. In the next section I briefly discuss a study that tests to what extent the same person behaves conditionally cooperatively inside and outside of the lab. This is an interesting question, because lab experiments are sometimes criticized for their lack of external validity.

2.2.3 Connections between the Lab and the Field

To gather information about the connection between lab and field behavior, the subjects in Benz and Meier (2005) took part in a lab experiment where they made a donation decision. The same subjects were observed in a naturally occurring environment—the donation decisions to two student support funds as described above and analyzed by Frey and Meier (2004). In one experiment $(n = 99)$, called "social funds," the donation was to exactly the same funds as in the naturally occurring situation; in a second experiment $(n = 83)$, called "charities," the donation was to another charity unrelated to the university.

The results show that lab and naturally occurring behavior are correlated. In the social funds experiment, the correlation between the average donation in the experiment and the average donation in the past four semesters is 0.28 $(p < 0.01)$. In the charities experiment the correlation is very similar (0.27; $p < 0.01$). A more refined statistical analysis that controls for sociodemographic variables in a multivariate regression supports the main findings. Thus, although the lab is an artificial environment, one can observe behavior also triggered in a naturally occurring environment.

A second interesting study on the connection between lab and field behavior was done by Carpenter and Seki (2005), who combined the advantages of both environments in a very innovative way. The subjects of their study were Japanese fishermen who took part in a lab experiment, but who were also observed in their daily fishing activities.

Specifically, Carpenter and Seki collected data from fishing hauls, which they related to measures of the fishermen's social preferences. Carpenter and Seki use a finitely repeated public goods experiment with and without opportunities for social disapproval to statistically derive five measures of social preferences for each fisherman: his level of unconditional cooperation; his conditional cooperation; the propensity to disapprove; the fisherman's response to received social disapproval; and finally, the level of the unconditional response to disapproval. The results show that fishing productivity is significantly related to the experimentally derived measures of social preferences.

In my view, the results by Benz and Meier (2005) and Carpenter and Seki (2005) strongly underscore the complementarity between the lab and the field. In both the lab and the field we observe real behavior. In the lab we observe behavior in an artificial environment, whereas in a naturally occurring situation behavior takes place in a context-rich environment. Depending on the research question, context-richness and artificiality are either drawbacks or advantages. The lab's advantage is that we can observe motivations and behavioral patterns with a degree of clarity most often not feasible outside the lab. The fact that we have observed conditional cooperation in tightly controlled lab experiments supports the interpretation of the field results as stemming from conditional cooperation. The observation of conditional cooperation in the field tells us that the psychology of conditional cooperation carries over from the lab to the field.

In the following section I will use the power of the lab to test the implications of conditional cooperation and preference heterogeneity. I see these experiments as four behavioral models that might help us interpret naturally occurring field situations in policy-relevant domains like tax morale or welfare state policies, but also in managerial domains such as workplace behavior. The four models will also help me guide my discussion of consequences for public policy and management.

2.3 Four Consequences of Conditional Cooperation and Preference Heterogeneity

I will present four experiments in this section that test four implications of conditional cooperation and preference heterogeneity in general. The testable consequences are that (1) in groups where group members are randomly selected *voluntary cooperation is fragile*; (2) *there*

are group interaction effects, meaning that people adapt their cooperation behavior to the relevant group they belong to; (3) *group composition matters*—in groups composed of like-minded types (groups composed of either cooperators or free riders) we should see starkly different cooperation patterns; and (4) *belief management matters*—in other words, factors that shift the belief about how much others contribute will influence contribution behavior. I discuss these four hypotheses and their experimental support in turn.

2.3.1 Voluntary Cooperation Is Fragile

I provide evidence in this section that heterogeneous motivations in randomly composed groups will lead to fragile cooperation. The reason is that free riders presumably do not contribute to the public good, while the conditional cooperators' contributions might be nonminimal, depending on their belief about other group members' contributions. Subjects learn the other team members' contributions during the repeated interaction. The free riders have no reason to react to that information. The conditional cooperators, on the other hand, will update their beliefs. Given that the average conditional cooperator does not fully match the others' contribution, the reaction will most likely be a decrease in contributions. There is no reason to expect that the remaining types (triangle contributors and others) will behave in a way that offsets the negative trend.

To test this argument rigorously, Fischbacher and Gächter (2006) combined the elicitation of contribution functions described above with a standard ten-period public goods game. The experiment was conducted in the stranger mode, meaning in every period the groups of four were formed randomly out of all twenty-four subjects present in a session. As predicted, contributions actually fell over time in all six sessions (from 40 percent initially to 10 percent on average by the last period).

Is this decline actually due to the interaction of heterogeneously motivated types? Stringent support for this conjecture comes from using the elicited contribution functions for predicting contributions. Recall that the strategies asked subjects to indicate how much they were prepared to contribute to the public good for all feasible average contribution levels of the other group members. In the standard ten-period public goods game Fischbacher and Gächter (2006) also elicited in each period each subject's *belief* about the other group members' contributions. Therefore, we can—given a stated belief about other

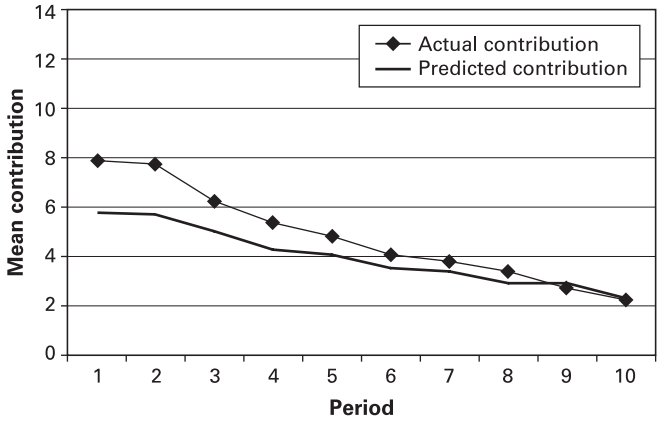

Figure 2.3
Average actual contributions and predicted contributions derived from beliefs and schedules. *Source*: Fischbacher and Gächter (2006).

group members' average contribution—*predict* what a subject should contribute to the public good if he or she would be perfectly consistent with his or her elicited contribution function. Figure 2.3 depicts the actual average contributions in the ten rounds of the public goods game and the predicted contributions as a result of stated beliefs and contribution schedules.

Although average predicted contributions are too low compared with actual contributions, we find that predicted contributions, which are derived from the contribution functions and the elicited beliefs, decline and converge to the actual pattern. This result therefore supports the argument that preference heterogeneity leads to unstable cooperation.

2.3.2 There Are Social Interaction Effects in Cooperation

If people are motivated by conditional cooperation, this may give rise to a social interaction effect, which occurs if an individual changes his or her behavior as a function of his or her respective group members' behavior. Identifying social or group interaction effects (often also called "neighborhood" or "peer effects") is notoriously difficult (Manski 2000). The ideal data set would observe the same individual at the same time in different groups, which are identical—apart from having different group members. Obviously this is impossible in the field. By contrast, in the lab it is possible to come very close to this counterfactual state. In an experiment, one is able to observe decisions of the

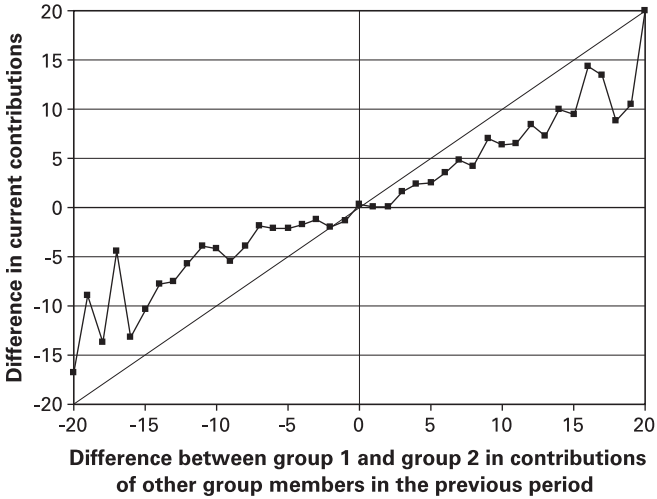

Figure 2.4
Social interaction effects: difference in own contribution as a function of the group members' contributions in the two groups. *Source*: Falk, Fischbacher, and Gächter (2005).

same subject at the same time in two economically identical environments. Social interactions—the fact that a person is systematically affected by the behavior of his or her group members—are the only reason to behave differently in these two environments. Falk, Fischbacher, and Gächter (2005) test this idea in a design where every subject is simultaneously a member of two groups, group 1 and group 2, which provide two independent public goods. The two groups consist of three group members each and are identical except that for each subject the other two group members in both groups are different people. Group composition stays constant for the twenty periods of the game. Falk, Fischbacher, and Gächter speak of a social interaction effect if the following holds: the larger the *difference* in contributions of group members in group 1 and group 2 in the previous period, the larger is the *difference* in current contributions of a group member to the two groups. Figure 2.4 provides the evidence from the 126 subjects who participated in this experiment.

The results provide unambiguous support for the social interaction hypothesis. In a given period a majority of subjects contributes more to the group that has contributed more in the previous period. This result holds for all fourteen independent units of observations, a result that is very unlikely to be due to chance ($p < 0.00007$).

2.3.3 Group Composition Matters

We have seen that a mixture of conditional cooperators and free riders is unfavorable for reaching cooperation in the public goods game. According to our third conjecture, conditional cooperators would presumably prefer to play the game with like-minded cooperators. Cooperation should be easy if the team players know they are among like-minded group members. Similarly, if the "true game" subjects are playing is a game where cooperation is one of the equilibria (free riding being another one), then knowing that others are like-minded cooperators should make it easy for subjects to coordinate on cooperation and to prevent free riding. Likewise, if free rider types know they are among other free riders, free riding should be paramount.

Gächter and Thöni (2005) conducted an experiment where subjects (105 in their version) play in groups of like-minded people. Like-mindedness refers to the type of subject according to classification as a free rider or a cooperator. The experiment starts with a three-person one-shot public goods game. When all subjects have chosen their contribution the subjects are ranked according to that contribution. Then the subjects are reassigned to new groups of three subjects. The reassignment works as follows: the three subjects with the highest contribution in the one-shot public goods game constitute the first group. The subjects with the fourth- to sixth-highest contribution are in the second group, and so on. Finally, the three least cooperative subjects find themselves in the last group. The subjects are informed about the reassignment procedure only after they finish the first game. Then the subjects learn the contributions their new group members chose in the one-shot public goods game. In the new group subjects play a ten-period public goods game. It is also important to note that the subjects do not know what the reassignment mechanism will be when choosing their contribution in the one-shot public goods game. Therefore, a high contribution in this first game credibly reveals a cooperative attitude.

The left panel of figure 2.5 shows the results of the main treatment. The maximal contribution in this game is 20. For expositional ease the groups are divided into three classes (top, middle, and low) according to their average contribution in the one-shot public goods game. The three graphs show the average contribution during the ten periods, separated by class. The unconnected dots in period zero depict the average contribution in the one-shot public goods game, which determines the group composition. The classes remain clearly separated over all periods. The groups in the top class consist to a large degree

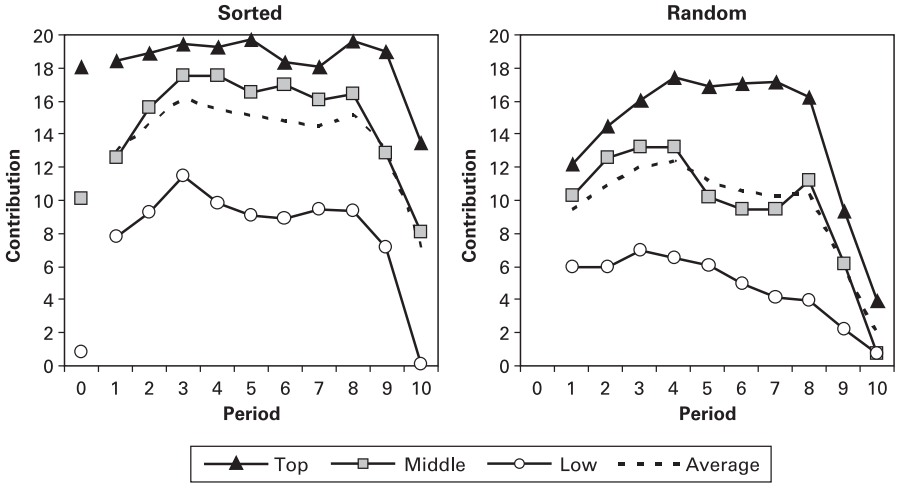

Figure 2.5
Left panel: average contributions over the ten periods for the top, middle, and low class in the sorted treatment. The unconnected dots in period zero are the average contributions in the ranking treatment. Right panel: average contribution of the most, intermediate, and least cooperative groups over the ten periods. *Source*: Gächter and Thöni (2005).

of subjects who contributed their entire endowment in the one-shot public goods game. These groups manage to maintain almost full cooperation until the penultimate period. The contributions of the middle class (consisting of subjects with intermediate contributions in the one-shot public goods game) show a similar pattern on a somewhat lower level. Surprisingly, the subjects in the low class, who almost all chose a contribution of zero in the one-shot public goods game, also manage to reach a certain level of cooperation in the repeated game. There are two explanations for this observation. First, if uncooperative subjects know that they are among fellow uncooperatives then it is clear there are no cooperative subjects to free ride on. This presumably motivates even uncooperative subjects to contribute in order to encourage the other free riders to contribute as well. A second related reason is that, in contrast to a one-shot game, a ten-period repeated game induces even free riders to strategically feign cooperation. Yet by the final period feigning cooperation does not pay off anymore, and consequently the contributions of these free rider subjects drop to zero.

The right panel of figure 2.5 shows the results from a control experiment. Groups are formed randomly in this experiment, meaning there is no reassignment according to cooperativeness. In order to make the

two treatments comparable, the data is still separated into the three classes of the top, middle, and lowest third with respect to their mean contribution levels. The separation now merely reflects the fact that there is variance in the contributions. Subjects in these control experiments are able to maintain a high level of contributions in all terciles until period 8; only in the penultimate and final periods do contributions drop to rather low levels. This "endgame effect" is typical for repeated public goods experiments in which groups are fixed for a finite number of periods (see, e.g., Keser and van Winden 2000).

Cooperation in the top class of the sorted treatment is much higher than the average contribution in the random treatment (dotted line in the right panel). However, the real value of the sorting mechanism becomes clear if we compare the top class with the most cooperative third of the groups in the random treatment. The average contribution of the top class of like-minded groups is significantly higher than the average contribution of the most cooperative third of the groups in the random treatment.

In summary, to be among like-minded people strongly affects cooperation behavior of all types. Related experiments suggest a similar conclusion. In Gunnthorsdottir, Houser, and McCabe (2007), subjects were regrouped as a function of their contributions but subjects were not aware of this. In Ones and Putterman (forthcoming) and Page, Putterman, and Unel (2005) subjects learned about others' contributions and were then regrouped according to the subjects' preferences. In all experiments regrouping made a significant difference relative to random groupings. Thus, for reasons of preference heterogeneity the "ecology of collective action," as Ones and Putterman aptly put it, matters a lot for the efficiency of voluntary cooperation.[4]

2.3.4 Belief Management Matters

Since the belief about others' contributions is important for conditional cooperators, our fourth conjecture says that any factor that alters these beliefs will influence cooperation. In the experiments of Fischbacher and Gächter (2006), for instance, beliefs evolved endogenously and mimicked the decline in cooperation. To test how beliefs can be influenced, Gächter and Renner (2005) developed a leader-follower design in a group of four players who stayed together for ten rounds (the number of rounds was known to the subjects). Specifically, one group member was designated as the leader. All group members had the same payoff function (see formula 2.1). The sole difference between the

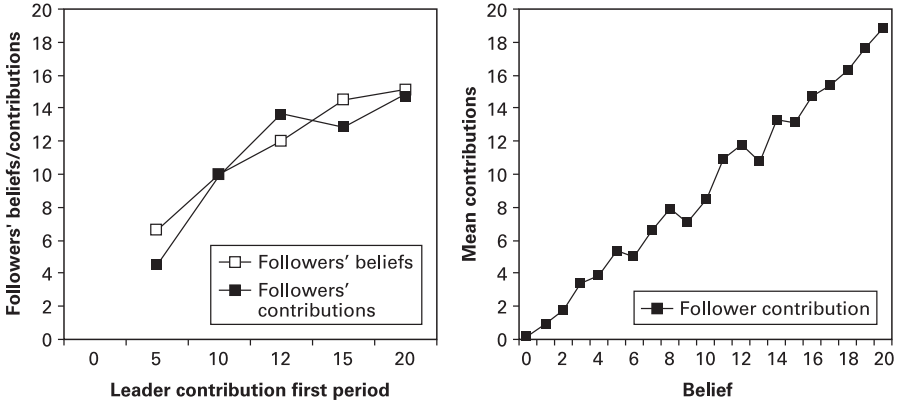

Figure 2.6
Left panel: leader's contribution in the first period and followers' beliefs and actual contributions in the first period. Right panel: relationship between beliefs and followers' actual contributions over all rounds. *Source*: Gächter and Renner (2005).

leader and the followers was that the leader made the first contribution decision. The followers observed the leader's contribution before they decided simultaneously about their own contributions. Gächter and Renner also elicited the followers' beliefs about the other followers' contributions. This allowed them to determine how the leader's contribution influences the beliefs about other followers' contributions.

The line with the open squares in the left panel of figure 2.6 shows that the leader's contribution in the first period positively influences the followers' beliefs about other followers' contributions. The first period is particularly interesting because the followers have not yet made any observation about the other followers' actual contributions. The more the leader contributes in the first period, the higher are the followers' beliefs about what other followers will contribute. This is the main and most direct evidence that a leader manages the followers' beliefs. In their actual contributions followers match their beliefs quite closely (see the line with the filled squares).

Using the data from all periods, the right panel of figure 2.6 shows that followers' beliefs and actual contributions are highly positively correlated. An econometric analysis reveals that these beliefs result from two sources: in a given period $t > 1$, beliefs are highly significantly positively correlated with the leader's contribution in this period. Yet beliefs are also highly significantly positively correlated with what the other followers contributed in the *previous* period $(t - 1)$.

Moreover, quantitatively, the followers' contributions in $t-1$ are more important than the leader's contribution for the followers' beliefs about other followers' contributions in period t. Thus there is an important path dependency in contributions. If the leader contributed little in the first period, followers are likely to contribute a small amount as well. This observation will—in addition to the leader's contribution—shape beliefs about other followers' contributions. In turn beliefs are—as the right panel of figure 2.6 shows—positively correlated with actual contributions. In other words, a bad start will make it very hard for the leader to lead his group by good example to high contribution levels. By contrast, a bold leader who sets a good example right from the beginning will positively influence followers' beliefs and contributions.

In summary, in this section I presented four experiments testing four implications of conditional cooperation and preference heterogeneity in general. As discussed earlier, I see these experiments as behavioral models that reveal something of the behavioral logic of conditional cooperation and preference heterogeneity. In the final two sections, I will therefore use these behavioral models to look at field phenomena and to discuss implications for public policy and management.[5]

2.4 Understanding Field Phenomena

2.4.1 Charitable Giving

During the war in former Yugoslavia three Austrian charity organizations set up the fund-raising campaign "Nachbar in Not" to finance food, clothes, and medical aid for the war victims. People donated more than 950 million Austrian schillings (approximately 70 million Euro) during the three years of the campaign to "Nachbar in Not" alone—donations to other charity organizations are not included. "Licht ins Dunkel" by the Austrian Broadcasting Corporation (ORF) is another example of a very successful and very large charitable fund-raising campaign that has for many years run around Christmas.

In both campaigns it was standard practice to list the names, hometowns, and donated amount of *all* donors who supported the campaigns, either on television or in newspapers. Donations by well-known politicians and celebrities were particularly prominently featured. The results from the field experiments discussed in section 2.2 and the lab results on how leader contributions can shape followers' contributions suggest that fund-raising organizers did not only rely on people's feelings of altruism, compassion, and warm glow, but also

on conditional cooperation. Seed money effects are a related phenomenon that at least in part exploits the psychology of conditional cooperation (List and Lucking-Reiley 2002). Likewise, fundraisers often make a symbolic gift to the donor. Reciprocity as a form of conditional cooperation predicts that nicer gifts will lead to higher donations. Falk (2004) tests this prediction in a field experiment and finds it unambiguously supported.

Conditional cooperation is of course not the only reason why people donate to charities (see Andreoni 2006 and Vesterlund 2006 for extensive reviews). People certainly also contribute for signaling reasons (Glazer and Konrad 1996), social approval (e.g., Andreoni and Petrie 2004; Soetevent 2005), or because observing others provides information about the charity (Romano and Yildirim 2001; Vesterlund 2003). Our results suggest that genuine conditional cooperation may be an important determinant of people's philanthropy, in addition to all other motivations.

2.4.2 Tax Morale, Benefit Fraud, and Corruption

Norms of reciprocity and conditional cooperation might also influence tax morale. Tax morale is an interesting case because taxes are typically used to finance public goods from which one benefits even if one has not paid taxes. Indeed, there is evidence both from the field and the lab that people pay more taxes than the standard economic model of tax evasion predicts (e.g., Andreoni, Erard, and Feinstein 1998; Webley et al. 1991; Torgler 2002). Our results suggest that, controlling for detection probabilities, conditional cooperators will be more likely to evade taxes (or falsely claim welfare benefits) if they have the impression that many others do the same. Too many cheaters can spoil tax morale. The evidence is consistent with this prediction. People are less likely to cheat on their taxes or to commit benefit fraud if others behave honestly (e.g., Cialdini 1989; Slemrod 1992; Andreoni, Erard, and Feinstein 1998; Rothstein 2000). Frey and Torgler (2004) provide the most direct evidence on the relevance of conditional cooperation for tax morale. They use data from the European Values Survey and conduct a multivariate analysis across 30 countries (with at least 1000 individuals per country). Frey and Torgler find a positive correlation between people's tax morale (measured by a question about whether cheating on tax is justified if you have the chance) and people's perception of how many others cheat on taxes.[6] While Frey and Torgler cannot prove causation in their data, the results from the strategy method

experiments by Fischbacher, Gächter, and Fehr (2001) and Fischbacher and Gächter (2006) suggest that causality goes from beliefs about others' cheating to their own cheating rather than vice versa.

The prevalence of corruption also seems to be influenced by motivations similar to those of conditional cooperation (see Abbink, Irlenbusch, and Renner 2002 for an experiment and further references to the literature). There are also important social interaction effects in these phenomena (Bertrand, Luttmer, and Mullainathan 2000; van der Klaauw and van Ours 2003), which is predicted by conditional cooperation and our model on these social interaction effects (section 2.3.2).

A particularly interesting observation is that the perception of the fairness of the tax system matters (Seidl and Traub 2001). Likewise, treatment by authorities apparently is an important determinant for people's tax morale (Pommerehne and Weck-Hannemann 1996; Frey 1997; Goette and Kucher 1998; Scholz and Lubell 1998; Feld and Frey 2002; Torgler 2003; Cummings et al. 2005; Alm and Torgler 2006). For instance, Cummings et al. (2005) present results from laboratory experiments they conducted in Botswana and South Africa. The experiments demonstrate that differences in the fairness of tax administration, perceived fiscal exchange, and attitudes toward the government can explain observed differences in compliance. Cummings et al. show that the experimental results are robust by replicating them for the same countries using survey responses measuring tax compliance.

How can our models explain such findings? First, there may be a direct effect from the concerned individual, who may reciprocate unfair treatment by authorities and/or the tax system with lower tax morale, simply because the taxpayer resents the unfair treatment (Smith 1992). Second, much like in the leadership experiments discussed in section 2.3.4, which showed that the leader strongly shapes the beliefs followers hold about other followers' behavior, tax authorities may have an indirect effect via beliefs about other taxpayers' behavior. The reason is that if many people share similar feelings and experiences, then this will lower the belief that others have a high tax morale, further undermining tax morale. Similarly, the government's trust in the honesty of its citizens may lead to a direct effect of "trust breeds trust" (Feld and Frey 2002), presumably because people like to be considered trustworthy. Again, if such feelings are widespread, they may shape beliefs about other citizens' tax morale and hence reinforce taxpayer morale.

A further interesting observation is that tax evasion at the Swiss cantonal level is higher in cantons where citizens have more direct democratic rights (e.g., Torgler 2005). According to our models, direct democratic procedures may positively influence tax morale. This is because direct democracy may affect beliefs about other people's tax morale when a tax law is passed in a referendum. A referendum signals people's opinion about a topic, and the dissemination of opinions via the result of a referendum may shape people's beliefs about others' behavior. Feld and Tyran (2002) tested this intuition in an experiment and found support for it.

2.4.3 Solidarity and Support for the Welfare State

Observers of welfare state policies (e.g., Wax 2000; Fong 2001; Fong, Bowles, and Gintis 2005; Lindbeck, Nyberg, and Weibull 1999) point out that many people hold reciprocity norms akin to the conditional cooperation observed in our experiments. Fong, Bowles, and Gintis (2002) even argue that "people support the welfare state because it conforms to deeply held norms of reciprocity and conditional obligations to others." There is evidence that people resent certain welfare policies if they think the recipient is a free rider who could earn his or her own living (Wax 2000; Fong, Bowles, and Gintis 2005). In their paper on tax payer resentment (i.e., the resentment against financing welfare payments), Besley and Coate (1992, 175) quote a notable British columnist, Lynda-Lee Porter, who neatly expresses the psychology of such resentment: "Our bronzed, healthy, young hedonistic army of self-unemployed are holidaying by the sea at our expense this year and, yes I do resent it. I resent working to support the idle loafers who have a laugh at our expensively generous system which allows them to get away with legalised plunder."

2.4.4 Work Morale

Business practitioners agree that "work morale" (i.e., loyalty, initiative, creativity, helping others, zest for the job, etc.) is crucial for productivity (Bewley 1999, 2005). Our models predict that work morale is strongly shaped by the behavior of management and coworkers. First, there may be social interaction effects in that people adapt their work morale to those of their peers. Empirical evidence supports this prediction (Ichino and Maggi 2000; Falk and Ichino 2006).

Additionally, our leadership model, discussed in section 2.3.4, and further experiments on leadership (e.g., Potters, Sefton, and Vesterlund

2004; Güth et al. 2004) suggest that managers may strongly influence morale and voluntary cooperation. To our knowledge there is no systematic evidence available, but some telling anecdotal evidence supports the point. For instance, Lawrence Weinstein, the head of Unisys, said in the wake of the Enron scandal, "Once you as a CEO go over the line, then people think it's okay to go over the line themselves."[7] This quote clearly expresses the conviction that leading by example matters for the ethical behavior of employees. Moreover, our results from section 2.3.4 suggest that a CEO's behavior may have long-lasting consequences on company morale and culture because of path-dependency effects.

Finally, our finding from section 2.3.3 that group composition matters may explain why companies sometimes fire workers, despite knowing that firing looks like a policy of management by threats. Yet Bewley (1999) notes that companies fire shirkers and incompetents to reestablish the work morale of the rest. Our models can explain this. Recall that the experimental findings reported above suggest that in heterogeneous groups contributions decline to low levels because the conditional cooperators stop cooperating once they experience free riding. If conditional cooperators know that they are among like-minded cooperators, cooperation can be established at very high levels. In a company context, this may mean that even a few shirkers can undermine work morale. Motivated workers may prefer that free riders are fired because they do not like being taken advantage of by their colleagues and because it reestablishes beliefs about others' team spirit.

2.5 Consequences for Public Policy and Management

In this section I briefly discuss policy implications that follow from the experimental findings and the four behavioral models discussed earlier. I first look at implications for public policy (section 2.5.1) and then at consequences for management (section 2.5.2).

2.5.1 Public Policy

Public policy is relevant mainly in the domains covered in sections 2.4.2 and 2.4.3. A first observation is that behavior by leaders—politicians and top officials—may matter strongly for citizens' morale. Leaders are belief managers, among other things. Leading by example strongly shapes beliefs about what others are doing, as the experiments in section 2.3.4 show. Therefore, there is a "multiplier effect," because a bad

example (dishonesty in tax matters, corruption, or unethical behavior in other domains) may not only have direct effects on the concerned individual, but may also have indirect belief effects on how others will react. Moreover, there may be strong path-dependency effects, which may adversely affect morale in the long run. Leaders should thus be role models for whom there are higher moral standards than for normal citizens. Leaders in particular should be forced to resign quickly if there is confirmed evidence of dishonesty and inappropriate behavior.

Belief management happens not only through leaders, but also through things like the perceived fairness of the tax system, fair treatment by authorities, and direct democratic participation rights. The experimental results discussed above suggest that these factors are very important and should be strengthened. Tax reforms should improve the fairness of the tax system (based on careful evidence on how fairly the tax system is perceived) not only because fairness is desirable in its own right, but also because of its indirect effect on beliefs about other citizens' tax morale. A similar conclusion holds for the reform of tax authorities. How tax authorities publicly deal with tax evasion may strongly shape people's beliefs about the prevalence of tax evasion and thereby, as shown by Frey and Torgler (2004), influence tax morale (see also Kahan 2005). For instance, tax authorities should not only put tax evaders in the limelight, but they should also communicate that the large majority of citizens pay their dues.[8] Direct democratic participation rights may also have a strong effect on tax morale (see, e.g., Feld and Frey 2002; Feld and Tyran 2002; Torgler 2005; Torgler and Schaltegger 2005). People value participation for reasons of procedural fairness (Benz 2005). Also, the referenda results communicate people's norms and values for many issues and thereby shape people's beliefs about others' norms and values. For constitutional reasons, granting direct democratic rights is admittedly not an easy task in representative democracies.

The experimental results from sections 2.3.1 and 2.3.3 suggest that free riders trigger reduced cooperation. Cooperation unravels when free riders are not punished because the conditional cooperators reduce their cooperation as well. Experiments have shown that this result can be overturned if targeted punishment of free riders is possible (e.g., Fehr and Gächter 2000) or if the free riders are excluded from the group (Gächter and Thöni 2005, section 3.3; Cinyabuguma, Page, and Putterman 2005). If there is punishment, free riders have an incentive

to cooperate and cooperators do not feel cheated. Cooperators therefore are happy to cooperate. This suggests that policy should aim to punish free riding (i.e., tax evasion, benefit fraud, and corruption). The experiments described above suggest that the goal should be to punish the free riders and at the same time to maintain the cooperators' optimistic beliefs by reassuring them that they will not be duped by the free riders. Thus they will continue to uphold their morale together with other like-minded cooperators.

Yet, apart from the legal implementation (which might be relatively simple), this is no easy task at all given the behavioral regularities discussed above. Punishment may entail monitoring and a general distrust of citizens. This is problematic for two reasons. First, there is evidence that monitoring may crowd out intrinsic motivation and reciprocal behavior (Frey 1993, 1997; Bohnet, Frey, and Huck 2001; Fehr and Gächter 2002). Second, monitoring may express distrust (Falk and Kosfeld 2006), which, in addition to the crowding-out effect, may have detrimental effects on beliefs about the tax morale of other taxpayers. Thus, in order to avoid the negative side effect of distrusting most citizens, policies should aim to punish big offenders severely and treat mild offenders (provided they are not serial offenders) mildly (by not using the full force of penal law, for instance). This has two advantages. First, strong sanctions have a deterrence effect, and they also reassure the honest citizens that large-scale antisocial behavior will be punished, which reduces the so-called sucker effect. Second, by trusting citizens and by fostering the fairness of the tax system and the tax authorities, the possible crowding out of intrinsic motivation and voluntary cooperation may be avoided.

The problem is complicated by the possibility that the game people actually play is one with multiple equilibria (see also Kahan 2005). Endemic cheating is an equilibrium, since conditional cooperators will also cheat if everyone else cheats. With multiple equilibria different policies may be required depending on the equilibrium currently in effect. A society with a good equilibrium of high trust, good tax morale, and low corruption must secure this equilibrium through policies that selectively punish the cheaters and maintain the conditional cooperators' good faith. If a society is trapped in a bad equilibrium, straightforward penalties and monitoring may be required to improve. Much more research is yet needed to understand what an optimal policy looks like in the presence of preference heterogeneity and multiple equilibria.

2.5.2 Management

The conclusions for management are very similar to those for public policy. First, managers, especially top managers, should be aware that they are role models who set an example and may strongly shape corporate cultures through path dependency in behaviors. Like politicians, they should therefore be held to high ethical standards.

Next, the problem of punishing shirkers in an organization is similar to the problem of how to treat antisocial behavior in the public policy domain. Management by threats will not create loyalty and may undermine intrinsic motivation and voluntary cooperation. Therefore, firing shirkers according to procedurally fair standards (Bewley 1999; Benz 2005) may help maintain high work morale among a team-spirited workforce.

Last, since group composition effects matter strongly for cooperative behavior, hiring team-spirited people is crucial if teamwork is important on the job. Composing teams of like-minded team players can help maintain high cooperation levels without any threat or negative side effects of monitoring and distrust.

2.6 Concluding Remarks

I have discussed experimental evidence from the lab and the field that shows many people are conditional cooperators, whereas others are best characterized as free riders. I believe that this sort of preference heterogeneity helps us better understand important phenomena in the field, like tax morale and attitudes toward the welfare state. Since, if many people are conditional cooperators, beliefs about others' behavior are highly relevant for voluntary cooperation, policy should not only take into account the incentive effects on an individual's behavior, but also how policy affects the beliefs and behavior of the majority of citizens, who are conditional cooperators. The evidence discussed in this chapter can only be considered a starting point. Much more research is needed for a proper understanding of the policy consequences of conditional cooperation and preference heterogeneity.

Notes

This chapter is part of the MacArthur Foundation Network on Economic Environments and the Evolution of Individual Preferences and Social Norms. Helpful comments by the seminar participants at the CESifo Summer Institute on Economics and Psychology in Venice, 2005, and by the editors and two anonymous referees are highly appreciated.

1. The laboratory allows for a degree of control not often feasible in a naturally occurring field situation. In all the experiments I will discuss below, participants earned considerable amounts of money that depended on their decisions. Thus, the laboratory allows observation of real economic behavior under controlled circumstances and also permits causal inferences often not feasible from naturally occurring data. See Kagel and Roth (1995) and Camerer (2003) for excellent overviews of experiments in economics and game theory and Guala (2005) for a discussion of the methodology of experimental economics.

2. No deception was involved because real frequencies (resulting from different time periods) were used.

3. A referee of this paper suggested that a potential problem might be callers' concern about their self-image and how they look in the eyes of the receiver of the call.

4. See Ones and Putterman (forthcoming) and Gächter and Thöni (2005) for a further discussion of the related literature.

5. See Falk (2003), Fehr and Fischbacher (2002), and Kahan (2005) for related discussions and further examples.

6. Cheaters may also entertain a self-serving belief about how many others cheat on their taxes, to justify their own misbehavior. Thus, causality may not run from beliefs about the prevalence of cheating in the population, but cheating may induce self-serving beliefs. I am grateful to a referee for suggesting this possibility.

7. Quoted from *The Economist*, July 27, 2002, p. 58.

8. An anonymous referee suggested, citing the following anecdotal evidence, that communication might be very important. India's 1997 tax amnesty has been seen as a financial success (it raised $2.5 billion from over 350,000 individuals). The tax amnesty was accompanied by intensive media activity. Celebrities such as sport and film stars promoted participation in the amnesty program, which contributed greatly to its success.

References

Abbink, K., B. Irlenbusch, and E. Renner. 2002. "An Experimental Bribery Game." *Journal of Law, Economics and Organization* 18, no. 2: 428–454.

Alm, J., and B. Torgler. 2006. "Culture Differences and Tax Morale in the United States and in Europe." *Journal of Economic Psychology* 27, no. 2: 224–246.

Anderson, S. P., J. K. Goeree, and C. Holt. 1998. "A Theoretical Analysis of Altruism and Decision Error in Public Goods Games." *Journal of Public Economics* 70: 297–323.

Andreoni, J. 1990. "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving." *Economic Journal* 100: 464–477.

Andreoni, J. 2006. "Philanthropy." In *Handbook on the Economics of Giving, Reciprocity and Altruism*, ed. L.-A. Gerard-Varet, S.-C. Kolm, and J. M. Ythier, 1201–1269. Amsterdam: Elsevier.

Andreoni, J., and R. Croson. 2008. "Partners versus Strangers: Random Rematching in Public Goods Experiments." Forthcoming in *Handbook of Experimental Economic Results*, ed. C. Plott and V. Smith. Amsterdam: Elsevier.

Andreoni, J., B. Erard, and J. Feinstein. 1998. "Tax Compliance." *Journal of Economic Literature* 36: 818–860.

Andreoni, J., and R. Petrie. 2004. "Public Goods Experiments without Confidentiality: A Glimpse into Fund-Raising." *Journal of Public Economics* 88: 1605–1623.

Bardsley, N., and P. Moffatt. Forthcoming. "The Experimetrics of Public Goods: Inferring Motivations from Contributions." *Theory and Decision*.

Benz, M. 2005. "The Relevance of Procedural Utility for Economics." Mimeo., University of Zurich.

Benz, M., and S. Meier. 2005. "Do People Behave in Experiments as in Real Life? Evidence from Donations." Mimeo., University of Zurich.

Bertrand, M., E. Luttmer, and S. Mullainathan. 2000. "Network Effects and Welfare Cultures." *Quarterly Journal of Economics* 115, no. 3: 1019–1055.

Besley, T., and S. Coate. 1992. "Understanding Taxpayer Resentment." *Journal of Public Economics* 48: 165–183.

Bewley, T. 1999. *Why Don't Wages Fall in a Recession*? Cambridge, MA: Harvard University Press.

Bewley, T. 2005. "Fairness, Reciprocity and Wage Rigidity." In *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*, ed. H. Gintis, S. Bowles, R. Boyd, and E. Fehr, 303–338. Cambridge, MA: MIT Press.

Bohnet, I., B. S. Frey, and S. Huck. 2001. "More Order with Less Law: On Contract Enforcement, Trust, and Crowding." *American Political Science Review* 95: 131–144.

Burlando, R., and F. Guala. 2005. "Heterogeneous Agents in Public Goods Experiments." *Experimental Economics* 8, no. 1: 35–54.

Camerer, C. 2003. *Behavioral Game Theory*. Princeton: Princeton University Press.

Carpenter, J. P., and E. Seki. 2005. "Do Social Preferences Increase Productivity? Field Experimental Evidence from Fishermen in Toyama Bay." IZA Discussion Paper No. 1697.

Cialdini, R. 1989. "Social Motivations to Comply: Norms, Values and Principles." In *Tax Payer Compliance*, ed. J. A. Roth, J. T. Scholz, and A. Dryden Witte, 200–227. Philadelphia: University of Pennsylvania Press.

Cinyabuguma, M., T. Page, and L. Putterman. 2005. "Cooperation under the Threat of Expulsion in a Public Goods Experiment." *Journal of Public Economics* 89: 1421–1435.

Croson, R. 2002. *Theories of Altruism and Reciprocity: Evidence from Linear Public Goods Games*. Mimeo., University of Pennsylvania.

Cummings, R., J. Matinez-Vazquez, M. McKee, and B. Torgler. 2005. "Effects of Tax Morale on Tax Compliance: Experimental and Survey Evidence." Working Paper No. 05-16, International Studies Program, Andrew Young School of Policy Studies, Georgia State University.

Dawes, R. 1980. "Social Dilemmas." *Annual Review of Psychology* 31: 169–193.

Falk, A. 2003. "*Homo Oeconomicus* versus *Homo Reciprocans*: Ansätze für ein neues Wirtschaftspolitisches Leitbild?" *Perspektiven der Wirtschaftspolitik* 4, no. 1: 141–172.

Falk, A. 2004. "Charitable Giving as a Gift Exchange: Evidence from a Field Experiment." Discussion Paper IZA-DP No. 1148, Bonn.

Falk, A., U. Fischbacher, and S. Gächter. 2005. "'Living in Two Neighborhoods'—Social Interaction in the Lab." Working Paper No. 150, Institute for Empirical Research in Economics, University of Zurich.

Falk, A., and A. Ichino. 2006. "Clean Evidence on Peer Pressure." *Journal of Labor Economics* 24, no. 1: 39–57.

Falk, A., and M. Kosfeld. 2006. "The Hidden Costs of Control." *American Economic Review* 96, no. 5: 1611–1630.

Fehr, E., and U. Fischbacher. 2002. "Why Social Preferences Matter—The Impact of Non-Selfish Motives on Competition, Cooperation and Incentives." *Economic Journal* 112: C1–C33.

Fehr, E., and S. Gächter. 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review* 90, no. 4: 980–994.

Fehr, E., and S. Gächter. 2002. "Do Incentive Contracts Undermine Voluntary Cooperation?" Working Paper No. 34, Institute for Empirical Research in Economics, University of Zurich.

Fehr, E., and K. Schmidt. 2003. "Theories of Fairness and Reciprocity: Evidence and Economic Applications." In *Advances in Economics and Econometrics. Theory and Applications*. Vol. I, ed. Matthias Dewatripont, Lars Hansen, and Stephen Turnovsky, 208–257. Cambridge: Cambridge University Press.

Feld, L., and B. S. Frey. 2002. "Trust Breeds Trust: How Taxpayers are Treated." *Economics of Governance* 3: 87–99.

Feld, L., and J.-R. Tyran. 2002. "Tax Evasion and Voting: An Experimental Analysis." *Kyklos* 55, no. 2: 197–222.

Fischbacher, U., S. Gächter, and E. Fehr. 2001. "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment." *Economics Letters* 71: 397–404.

Fischbacher, U., and S. Gächter. 2006. "Heterogeneous Social Preferences and the Dynamics of Free Riding. CeDEx Discussion Paper Series No. 2006-01, University of Nottingham.

Fong, C. 2001. "Social Preferences, Self-Interest, and the Demand for Redistribution." *Journal of Public Economics* 82: 225–246.

Fong, C., S. Bowles, and H. Gintis. 2005. "Reciprocity and the Welfare State." In *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*, ed. H. Gintis, S. Bowles, R. Boyd, and E. Fehr, 277–302. Cambridge, MA: MIT Press.

Frey, B. S. 1993. "Shirking or Work Morale? The Impact of Regulating." *European Economic Review* 37: 1523–1532.

Frey, B. S. 1997. "A Constitution for Knaves Crowds out Civic Virtues." *Economic Journal* 107: 1043–1053.

Frey, B. S., and S. Meier. 2004. "Social Comparisons and Pro-Social Behavior. Testing 'Conditional Cooperation' in a Field Experiment." *American Economic Review* 94, no. 5: 1717–1722.

Frey, B. S., and B. Torgler. 2004. "Taxation and Conditional Cooperation." CREMA Working Paper 2004-20, Center for Research in Economics, Management and the Arts, Basel, Switzerland.

Gächter, S., and B. Herrmann. 2005. "Human Cooperation from an Economic Perspective." In *Cooperation in Primates and Humans: Mechanisms and Evolution*, ed. P. M. Kappeler and C. P. van Schaik, 279–303. Heidelberg: Springer.

Gächter, S., and E. Renner. 2005. "Leading by Example in the Presence of Free Rider Incentives." Mimeo., University of Nottingham.

Gächter, S., and C. Thöni. 2005. "Social Learning and Voluntary Cooperation among Like-Minded People." *Journal of the European Economic Association* 3: 303–314.

Glazer, A., and K. A. Konrad. 1996. "A Signaling Explanation for Private Charity." *American Economic Review* 86, no. 4: 1019–1028.

Goette, L., and M. Kucher. 1998. "Trust Me: An Empirical Analysis of Taxpayer Honesty." *Finanzarchiv* 55: 429–444.

Guala, F. 2005. *The Methodology of Experimental Economics*. Cambridge: Cambridge University Press.

Gunnthorsdottir, A., D. Houser, and K. McCabe. (2007). "Dispositions, History and Contributions in Public Goods Experiments." *Journal of Economic Behavior and Organization* 62: 304–315.

Güth, W., V. Levati, M. Sutter, and E. van der Heijden. 2004. "Leadership and Cooperation in Public Goods Experiments." Mimeo., Max Planck Institute, Jena.

Harrison, G., and J. List. 2004. "Field Experiments." *Journal of Economic Literature* 42: 1009–1055.

Heldt, T. 2005. "Conditional Cooperation in the Field: Cross-Country Skiers' Behavior in Sweden." Mimeo., Dlarna University and Uppsala University, Sweden.

Ichino, A., and G. Maggi. 2000. "Work Environment and Individual Background: Explaining Regional Shirking Differential in a Large Italian Firm." *Quarterly Journal of Economics* 115, no. 3: 1057–1090.

Kagel, J., and A. E. Roth, eds. 1995. *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.

Kahan, D. 2005. "The Logic of Reciprocity: Trust, Collective Action, and Law." In *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*, ed. H. Gintis, S. Bowles, R. Boyd, and E. Fehr, 339–378. Cambridge, MA: MIT Press.

Kelley, H., and A. Stahelski. 1970. Social Interaction Basis of Cooperators' and Competitors' Beliefs about Others." *Journal of Personality and Social Psychology* 16: 190–197.

Keser, C., and F. van Winden. 2000. "Conditional Cooperation and Voluntary Contributions to Public Goods." *Scandinavian Journal of Economics* 102, no. 1: 23–39.

Ledyard, J. 1995. "Public Goods: A survey of Experimental Research." In *Handbook of Experimental Economics*, ed. J. Kagel and A. E. Roth, 111–194. Princeton, NJ: Princeton University Press.

Lindbeck, A., S. Nyberg, and J. Weibull. 1999. "Social Norms and Economic Incentives in the Welfare State." *Quarterly Journal of Economics* 114, no. 1: 1–35.

List, J., and D. Lucking-Reiley. 2002. "The Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign." *Journal of Political Economy* 110, no. 1: 215–233.

Manski, C. 2000. "Economic Analysis of Social Interactions." *Journal of Economic Perspectives* 14, no. 3: 115–136.

Martin, R., and J. Randal. 2005. "Voluntary Contributions to a Public Good: A Natural Field Experiment." Mimeo., Victoria University, Wellington, New Zealand.

Muller, L., M. Sefton, R. Steinberg, and L. Vesterlund. 2005. "Strategic Behavior and Learning in Repeated Voluntary-Contribution Experiments." CeDEx Working Paper No. 2005-13, University of Nottingham.

Ockenfels, A. 1999. "Fairness, Reziprozität und Eigennutz—Ökonomische Theorie und Experimentelle Evidenz." *Die Einheit der Gesellschaftswissenschaften* 108: 111–194.

Orbell, J. M., and R. M. Dawes. 1993. "Social Welfare, Cooperators' Advantage, and the Option of Not Playing the Game." *American Sociological Review* 58, no. 6: 787–800.

Ones, U., and L. Putterman. Forthcoming. "The Ecology of Collective Action: A Public Goods and Sanctions Experiment with Controlled Group Formation." *Journal of Economic Behavior and Organization*.

Page, T., L. Putterman, and B. Unel. 2005. "Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency." *Economic Journal* 116: 1032–1052.

Palfrey, T. R., and J. E. Prisbrey. 1997. "Anomalous Behavior in Public Goods Experiments: How Much and Why?" *American Economic Review* 87: 829–846.

Pommerehne, W., and H. Weck-Hannemann. 1996. "Tax Rates, Tax Administration and Income Tax Evasion in Switzerland." *Public Choice* 88: 161–170.

Potters, J., M. Sefton, and L. Vesterlund. 2004. "Leading-by-Example in Voluntary Contribution Games: An Experimental Study." CeDEx Working Paper No. 2004-01, University of Nottingham.

Romano, R., and H. Yildirim. 2001. "Why Charities Announce Donations: A Positive Perspective." *Journal of Public Economics* 81, no. 3: 423–447.

Rothstein, B. 2000. "Trust, Social Dilemmas and Collective Memories." *Journal of Theoretical Politics* 12, no. 4: 447–501.

Samuelson, P. 1954. "The Pure Theory of Public Expenditures." *Review of Economics and Statistics* 36, no. 4: 387–389.

Scholz, J. T., and M. Lubell. 1998. "Trust and Taxpayers: Testing the Heuristic Approach to Collective Action." *American Journal of Political Science* 45: 160–178.

Seidl, C., and S. Traub. 2001. "Taxpayers' Attitudes, Behaviour, and Perception of Fairness." *Pacific Economic Review* 6, no. 2: 255–267.

Shang, J., and R. Croson. 2005. "Field Experiments in Charitable Contribution: The Impact of Social Influence on the Voluntary Provision of Public Goods." Mimeo., Wharton University.

Slemrod, J., ed. 1992. *Why People Pay Taxes: Tax Compliance and Enforcement*. Ann Arbor: University of Michigan Press.

Sobel, J. 2005. "Interdependent References and Reciprocity." *Journal of Economic Literature* 43: 392–436.

Soetevent, A. 2005. "Anonymity in Giving in a Natural Context—A Field Experiment in 30 Churches. *Journal of Public Economics* 89: 2301–2323.

Smith, K. W. 1992. "Reciprocity and Fairness: Positive Incentives for Tax Compliance." In *Why People Pay Taxes: Tax Compliance and Enforcement*, ed. J. Slemrod, 223–250. Ann Arbor: University of Michigan Press.

Torgler, B. 2002. "Speaking to Theorists and Searching for Facts: Tax Morale and Tax Compliance in Experiments." *Journal of Economic Surveys* 16, no. 5: 657–683.

Torgler, B. 2003. "Tax Morale, Rule-Governed Behavior and Trust." *Constitutional Political Economy* 14: 119–140.

Torgler, B. 2005. "Tax Morale and Direct Democracy." *European Journal of Political Economy* 21: 525–531.

Torgler, B., and C. Schaltegger. 2005. "Tax Amnesties and Political Participation." *Public Finance Review* 33, no. 3: 403–431.

Tyran, J.-R., and R. Sausgruber. 2006. "A Little Fairness May Induce a Lot of Redistribution in Democracy." *European Economic Review* 50, no. 2: 469–485.

van der Klaauw, B., and J. van Ours. 2003. "From Welfare to Work: Does the Neighborhood Matter?" *Journal of Public Economics* 87: 957–985.

Vesterlund, L. 2003. "The Informational Value of Sequential Fundraising." *Journal of Public Economics* 87: 627–657.

Vesterlund, L. 2006. "Why Do People Give?" In *The Nonprofit Sector*, 2nd ed., ed. R. Steinberg and W. W. Powell, 568–588. New Haven, CT: Yale University Press.

Wax, A. 2000. "Rethinking Welfare Rights: Reciprocity Norms, Reactive Attitudes, and the Political Economy of Welfare Reforms." *Law and Contemporary Problems* 63: 257–298.

Webley, P., H. Robben, H. Elffers, and D. Hessing. 1991. *Tax Evasion: An Experimental Approach*. Cambridge: Cambridge University Press.