Collective Action and the Evolution of Social Norms

Elinor Ostrom

With the publication of *The Logic of Collective Action* in 1965, Mancur Olson challenged a cherished foundation of modern democratic thought that groups would tend to form and take collective action whenever members jointly benefitted. Instead, Olson (1965, p. 2) offered the provocative assertion that no self-interested person would contribute to the production of a public good: "[U]nless the number of individuals in a group is quite small, or unless there is coercion or some other special device to make individuals act in their common interest, *rational, self-interested individuals will not act to achieve their common or group interests.*" This argument soon became known as the "zero contribution thesis."

The idea that rational agents were not likely to cooperate in certain settings, even when such cooperation would be to their mutual benefit, was also soon shown to have the structure of an n-person prisoner's dilemma game (Hardin 1971, 1982). Indeed, the prisoner's dilemma game, along with other social dilemmas, has come to be viewed as the canonical representation of collective action problems (Lichbach, 1996). The zero contribution thesis underpins the presumption in policy textbooks (and many contemporary public policies) that individuals cannot overcome collective action problems and need to have externally enforced rules to achieve their own long-term self-interest.

The zero contribution thesis, however, contradicts observations of everyday life. After all, many people vote, do not cheat on their taxes, and contribute effort

■ *Elinor Ostrom is Co-Director, Workshop in Political Theory and Policy Analysis, Center for the Study of Institutions, Population, and Environmental Change, Indiana University, Bloomington, Indiana.*

to voluntary associations. Extensive fieldwork has by now established that individuals in all walks of life and all parts of the world voluntarily organize themselves so as to gain the benefits of trade, to provide mutual protection against risk, and to create and enforce rules that protect natural resources.[1] Solid empirical evidence is mounting that governmental policy can frustrate, rather than facilitate, the private provision of public goods (Montgomery and Bean, 1999). Field research also confirms that the temptation to free ride on the provision of collective benefits is a universal problem. In all known self-organized resource governance regimes that have survived for multiple generations, participants invest resources in monitoring and sanctioning the actions of each other so as to reduce the probability of free riding (Ostrom, 1990).

While these empirical studies have posed a severe challenge to the zero contribution theory, these findings have not yet been well integrated into an accepted, revised theory of collective action. A substantial gap exists between the theoretical prediction that self-interested individuals will have extreme difficulty in coordinating collective action and the reality that such cooperative behavior is widespread, although far from inevitable.

Both theorists and empirical researchers are trying to bridge this gap. Recent work in game theory—often in a symbiotic relationship with evidence from experimental studies—has set out to provide an alternative micro theory of individual behavior that begins to explain anomalous findings (McCabe, Rassenti and Smith, 1996; Rabin, 1993; Fehr and Schmidt, 1999; Selten, 1991; Bowles, 1998). On the empirical side, considerable effort has gone into trying to identify the key factors that affect the likelihood of successful collective action (Feeny et al., 1990; Baland and Platteau, 1996; Ostrom, forthcoming).

This paper will describe both avenues of research on the underpinnings of collective action, first focusing on the experimental evidence and potential theoretical explanations, and then on the real-world empirical evidence. This two-pronged approach to the problem has been a vibrant area of research that is yielding many insights. A central finding is that the world contains multiple types of individuals, some more willing than others to initiate reciprocity to achieve the benefits of collective action. Thus, a core question is how potential cooperators signal one another and design institutions that reinforce rather than destroy conditional cooperation. While no full-blown theory of collective action yet exists, evolutionary theories appear most able to explain the diverse findings from the lab and the field and to carry the nucleus of an overarching theory.

---

[1] See Milgrom, North and Weingast (1990) and Bromley et al. (1992). An extensive bibliography by Hess (1999) on diverse institutions for dealing with common pool resources can be searched on the web at ⟨http://www.indiana.edu/~workshop/wsl/wsl.html⟩ or obtained on a CD-ROM disk.

Laboratory Evidence on Rational Choice in Collective Action Situations

Most studies by political economists assume a standard model of rational individual action—what I will call a rational egoist. A wide range of economic experiments have found that the rational egoist assumption works well in predicting the outcome in auctions and competitive market situations (Kagel and Roth, 1995). While subjects do not arrive at the predicted equilibrium in the first round of market experiments, behavior closely approximates the predicted equilibrium by the end of the first five rounds in these experiments. One of the major successes of experimental economics is to demonstrate the robustness of microeconomic theory for explaining market behavior.

In regard to collective action situations, on the other hand, the results are entirely different. Linear public good experiments are widely used for examining the willingness of individuals to overcome collective action problems. In a linear public good experiment, each individual is endowed with a fixed set of assets and must decide how many of these assets to contribute to a public good. When an individual makes a contribution of, say, 10 units to the public good, each of the participants in the group, including that individual, receive a benefit of, say, five units apiece. In this setting, the optimal outcome for the group of players as a whole is for everyone to contribute all of their endowments to provide the public good (if a group of 10 people, each individual contribution of 10 will have a social payoff of 50!). However, the unique equilibrium for rational egoists in a single-shot game is that everyone contributes zero, since each individual has access to benefits of the public good funded by the contributions of others, without paying any costs.[2]

If the public goods game is played for a finite number of rounds, zero is also the predicted equilibrium for every round. Rational egoists will reason that zero contribution is the equilibrium in the last round, and because they expect everyone to contribute zero in the last round, they also expect everyone to contribute zero in the second-to-last round, and eventually by backward induction they will work their way to the decision not to contribute to the public good in the present. Of course, these predictions are based on the assumptions that all players are fully rational and interested only in their own immediate financial payoff, that all players

---

[2] In a linear public good game, utility is a linear function of individual earnings,

$$U_i = U_i[(E - x_i) + A \cdot P(\sum x_i)],$$

where E is an individual endowment of assets, x_i is the amount of this endowment contributed to provide the good, A is the allocation formula used to distribute the group benefit to individual players, and P is the production function. In a linear public good game, A is specified as $1/N$ and $0 < 1/N < P < 1$ (but both of these functions vary in other types of collective action). So long as $P < 1$, contributing to the collective good is never an optimal strategy for a fully self-interested player.

understand the structure of the game fully and believe that all other players are fully rational, and that no external actor can enforce agreements between the players.

Since the first public good experiments were undertaken by Dawes, McTavish and Shaklee (1977), a truly huge number of such experiments has been undertaken under various conditions (see Davis and Holt, 1993; Ledyard, 1995; and Offerman, 1997, for an overview). By now seven general findings have been replicated so frequently that these can be considered the core facts that theory needs to explain.

1) Subjects contribute between 40 and 60 percent of their endowments to the public good in a one-shot game as well as in the first round of finitely repeated games.

2) After the first round, contribution levels tend to decay downward, but remain well above zero. A repeated finding is that over 70 percent of subjects contribute nothing in the announced last round of a finitely repeated sequence.

3) Those who believe others will cooperate in social dilemmas are more likely to cooperate themselves. A rational egoist in a public good game, however, should not in any way be affected by a belief regarding the contribution levels of others. The dominant strategy is a zero contribution no matter what others do.

4) In general, learning the game better tends to lead to more cooperation, not less. In a clear test of an earlier speculation that it just took time for subjects to learn the predicted equilibrium strategy in public good games, Isaac, Walker and Williams (1994) repeated the same game for 10 rounds, 40 rounds, and 60 rounds with experienced subjects who were specifically told the end period of each design. They found that the rate of decay is inversely related to the number of decision rounds. In other words, instead of learning *not* to cooperate, subjects learn how to cooperate at a moderate level for ever-longer periods of time!

5) Face-to-face communication in a public good game—as well as in other types of social dilemmas—produces substantial increases in cooperation that are sustained across all periods including the last period (Ostrom and Walker, 1997).[3] The strong effect of communication is not consistent with currently accepted theory, because verbal agreements in these experiments are not enforced. Thus, communication is only "cheap talk" and makes no difference in predicted outcomes in social dilemmas. But instead of using this opportunity to fool others into cooperating, subjects use the time to discuss the optimal joint strategy, to extract promises from one another, and to give verbal tongue-lashings when aggregate contributions fall below promised levels. Interestingly, when communication is implemented by allowing subjects to signal promises to cooperate through their

---

[3] Even more startling, Bohnet and Frey (1999) find that simply allowing subjects to see the other persons with whom they are playing greatly increases cooperation as contrasted to completely anonymous situations. Further, Frank, Gilovich and Regan (1993) find that allowing subjects to have a face-to-face discussion enables them to predict who will play cooperatively at a rate significantly better than chance.

computer terminals, much less cooperation occurs than in experiments allowing face-to-face communication.

6) When the structure of the game allows it, subjects will expend personal resources to punish those who make below-average contributions to a collective benefit, including the last period of a finitely repeated game. No rational egoist is predicted to spend anything to punish others, since the positive impact of such an action is shared equally with others whether or not they also spend resources on punishing. Indeed, experiments conducted in the United States, Switzerland, and Japan show that individuals who are initially the least trusting are more willing to contribute to sanctioning systems and are likely to be transformed into strong cooperators by the availability of a sanctioning mechanism (Fehr and Gächter, forthcoming). The finding that face-to-face communication is more efficacious than computerized signaling is probably due to the richer language structure available and the added intrinsic costs involved in hearing the intonation and seeing the body language of those who are genuinely angry at free riders (Ostrom, 1998a).

7) The rate of contribution to a public good is affected by various contextual factors including the framing of the situation and the rules used for assigning participants, increasing competition among them, allowing communication, authorizing sanctioning mechanisms, or allocating benefits.

These facts are hard to explain using the standard theory that all individuals who face the same objective game structure evaluate decisions in the same way![4] We cannot simply resort to the easy criticism that undergraduate students are erratic. Increasing the size of the payoffs offered in experiments does not appear to change the broad patterns of empirical results obtained.[5] I believe that one is forced by these well-substantiated facts to adopt a more eclectic (and classical) view of human behavior.

Building a Theory of Collective Action with Multiple Types of Players

From the experimental findings, one can begin to put together some of the key assumptions that need to be included in a revised theory of collective action.

---

[4] Although the discussion here focuses on collective action and public good games in particular, a broader range of experiments exists in which the rational egoist's prediction pans out badly. These include the ultimatum game, the dictator game, the trust game, and common-pool resources games with communication.

[5] Most of these experiments involve ultimatum games but the findings are quite relevant. Cameron (1995), for example, conducted ultimatum experiments in Indonesia and thereby was able to use sums that amounted to three months' wages. In this extremely tempting situation, she still found that 56 percent of the Proposers allocated between 40 and 50 percent of this very substantial sum to the Responder.

Assuming the existence of two types of "norm-using" players—"conditional cooperators" and "willing punishers"—in addition to rational egoists, enables one to start making more coherent sense out of the findings of the laboratory experiments on contributions to public goods.

Conditional cooperators are individuals who are willing to initiate cooperative action when they estimate others will reciprocate and to repeat these actions as long as a sufficient proportion of the others involved reciprocate. Conditional cooperators are the source of the relatively high levels of contributions in one-shot or initial rounds of prisoner's dilemma and public good games. Their initial contributions may encourage some rational egoists to contribute as well, so as to obtain higher returns in the early rounds of the game (Kreps et al., 1982). Conditional cooperators will tend to trust others and be trustworthy in sequential prisoner's dilemma games as long as the proportion of others who return trust is relatively high. Conditional cooperators tend to vary, however, in their tolerance for free riding. Some are easily disappointed if others do not contribute, so they begin to reduce their own contributions. As they reduce their contributions, they discourage other conditional cooperators from further contributions. Without communication or institutional mechanisms to stop the downward cascade, eventually only the most determined conditional cooperators continue to make positive contributions in the final rounds.

The first four findings are consistent with an assumption that conditional cooperators are involved in most collective action situations. Conditional cooperators are apparently a substantial proportion of the population, given the large number of one-shot and finitely repeated experiments with initial cooperation rates ranging from 40 to 60 percent. Estimating that others are likely to cooperate should increase their willingness to cooperate. Further, knowing the number of repetitions will be relatively long, conditional cooperators can restrain their disappointment with free riders and keep moderate levels of cooperation (and joint payoffs) going for ever-longer periods of time.

The fifth and sixth findings depend on the presence of a third type of player who is willing, if given an opportunity, to punish presumed free riders through verbal rebukes or to use costly material payoffs when available. Willing punishers may also become willing rewarders if the circle of relationships allows them to reward those who have contributed more than the minimal level. Some conditional cooperators may also be willing punishers. Together, conditional cooperators and willing punishers create a more robust opening for collective action and a mechanism for helping it grow. When allowed to communicate on a face-to-face basis, willing punishers convey a considerable level of scorn and anger toward others who have not fully reciprocated their trust and give substantial positive encouragement when cooperation rates are high. Even more important for the long-term sustainability of collective action is the willingness of some to pay a cost to sanction others. The presence of these norm-using types of players is hard to dispute given the

empirical evidence. The key question now is: How could these norm-using types of players have emerged and survived in a world of rational egoists?

Emergence and Survival of Multiple Types of Players in Evolutionary Processes

Evolutionary theories provide useful ways of modeling the emergence and survival of multiple types of players in a population. In a strict evolutionary model, individuals inherit strategies and do not change strategies in their lifetime. In this approach, those carrying the more successful strategies for an environment reproduce at a higher rate. After many iterations the more successful strategies come to prominence in the population (Axelrod, 1986). Such models are a useful starting point for thinking about competition and relative survival rates among different strategies.[6]

Human evolution occurred mostly during the long Pleistocene era that lasted for about 3 million years, up to about 10,000 years ago. During this era, humans roamed the earth in small bands of hunter-gatherers who were dependent on each other for mutual protection, sharing food, and providing for the young. Survival was dependent not only on aggressively seeking individual returns but also on solving many day-to-day collective action problems. Those of our ancestors who solved these problems most effectively, and learned how to recognize who was deceitful and who was a trustworthy reciprocator, had a selective advantage over those who did not (Barkow, Cosmides and Tooby, 1992).

Evolutionary psychologists who study the cognitive structure of the human brain conclude that humans do not develop general analytical skills that are then applied to a variety of specific problems. Humans are not terribly skilled at general logical problem solving (as any scholar who has taught probability theory to undergraduates can attest). Rather, the human brain appears to have evolved a domain-specific, human-reasoning architecture (Clark and Karmiloff-Smith, 1991). For example, humans use a different approach to reasoning about deontic relationships—what is forbidden, obligated, or permitted—as contrasted to reasoning about what is true and false. When reasoning about deontic relationships, humans tend to check for violations, or cheaters (Manktelow and Over, 1991). When reasoning about whether empirical relationships are true, they tend to use a confirmation strategy (Oaksford and Chater, 1994). This deontic effect in human reasoning has repeatedly been detected even in children as young as three years old and is not associated with overall intelligence or educational level of the subject (Cummins, 1996).

Thus, recent developments in evolutionary theory and supporting empirical research provide strong support for the assumption that modern humans have inherited a propensity to learn social norms, similar to our inherited propensity to learn grammatical rules (Pinker, 1994). Social norms are shared understandings

---

[6] For examples of strict evolutionary models involving collective action, see Nowak and Sigmund (1998), Sethi and Somanathan (1996) and Epstein and Axtell (1996).

about actions that are obligatory, permitted, or forbidden (Crawford and Ostrom, 1995). Which norms are learned, however, varies from one culture to another, across families, and with exposure to diverse social norms expressed within various types of situations. The intrinsic cost or anguish that an individual suffers from failing to use a social norm, such as telling the truth or keeping a promise, is referred to as guilt, if entirely self-inflicted, or as shame, when the knowledge of the failure is known by others (Posner and Rasmusen, 1999).

The Indirect Evolutionary Approach to Adaptation Through Experience

Recent work on an *indirect* evolutionary approach to the study of human behavior offers a rigorous theoretical approach for understanding how preferences—including those associated with social norms—evolve or adapt (Güth and Yaari, 1992; Güth, 1995). In an indirect evolutionary model, players receive objective payoffs, but make decisions based on the transformation of these material rewards into intrinsic preferences. Those who value reciprocity, fairness, and being trustworthy add a subjective change parameter to actions (of themselves or others) that are consistent or not consistent with their norms. This approach allows individuals to start with a predisposition to act in a certain way—thus, they are not rational egoists who only look forward—but it also allows those preferences to adapt in a relatively short number of iterations given the objective payoffs they receive and their intrinsic preferences about those payoffs.

Social dilemmas associated with games of trust, like sequential prisoner's dilemma games, are particularly useful games for discussing the indirect evolutionary approach. In such games, if two players trust each other and cooperate, they can both receive a moderately high payoff. However, if one player cooperates and the other does not, then the one who did not cooperate receives an even higher payoff, while the other receives little or nothing. For a rational egoist playing this game, the choice is not to trust, because the expectation is that the other player will not trust, either. As a result, both players will end up with lower payoffs than if they had been able to trust and cooperate. When considering such games, it is useful to remember that most contractual relationships—whether for private or public goods—have at least an element of this basic structure of trying to assure mutual trust. An indirect evolutionary approach explains how a mixture of norm-users and rational egoists would emerge in settings where standard rational choice theory assumes the presence of rational egoists alone.

In this approach, social norms may lead individuals to behave differently in the same objective situation depending on how strongly they value conformance with (or deviance from) a norm. Rational egoists can be thought of as having intrinsic payoffs that are the same as objective payoffs, since they do not value the social norm of reciprocity. Conditional cooperators (to take only one additional type of player for now) would be modeled as being trustworthy types and would have an additional parameter that adds value to the objective payoffs when reciprocating trust with trustworthiness. By their behavior and resulting interaction, however,

different types of players are likely to gain differential objective returns. In a game of trust where players are chosen from a population that initially contains some proportion of rational egoists and conditional cooperators, the level of information about player types affects the relative proportion of rational egoists and conditional cooperators over time. With complete information regarding types, conditional cooperators playing a trustworthy strategy will more frequently receive the higher payoff, while rational egoists will consistently receive a lower payoff, since others will not trust them.

Only the trustworthy type would survive in an evolutionary process with complete information (Güth and Kliemt, 1998, p. 386). Viewed as a cultural evolutionary process, new entrants to the population would be more likely to adopt the preference ordering of those who obtained the higher material payoffs in the immediate past (Boyd and Richerson, 1985). Those who were less successful would tend to learn the values of those who had achieved higher material rewards (Börgers and Sarin, 1997).[7] Where a player's type is common knowledge, rational egoists would not survive. Full and accurate information about all players' types, however, is a very strong assumption and unlikely to be met in most real world settings.

If there is no information about player types for a relatively large population, preferences will evolve so that only rational egoists survive.[8] If information about the proportion of a population that is trustworthy is known, and no information is known about the type of a specific player, Güth and Kliemt (1998) show that first players will trust second players as long as the expected return of meeting trustworthy players and receiving the higher payoff exceeds the payoff obtained when neither player trusts the other. In such a setting, however, the share of the population held by the norm-using types is bound to decline. On the other hand, if there is a noisy signal about a player's type that is at least more accurate than random, trustworthy types will survive as a substantial proportion of the population. Noisy signals may result from seeing one another, face-to-face communication, and various mechanisms that humans have designed to monitor each other's behavior.

Evidence Testing the Indirect Evolutionary Approach

An indirect evolutionary approach is able to explain how a mixture of contingent cooperators and rational egoists would emerge in settings where traditional

---

[7] Eshel, Samuelson and Shaked (1998) develop a learning model where a population of Altruists who adopt a strategy of providing a local public good interacts in a local neighborhood with a population of Egoists who free ride. In this local interaction setting, Altruists' strategies are imitated sufficiently often in a Markovian learning process to become one of the absorbing states. Altruists interacting with Egoists outside a circular local neighborhood are not so likely to survive.

[8] This implies that, in a game where players know only their own payoffs and not the payoffs of others, they are more likely to behave like rational egoists. McCabe and Smith (1999) show that players tend to evolve toward the predicted, subgame perfect outcomes in experiments where they have only private information of their own payoffs and to cooperative outcomes when they have information about payoffs and the moves made by other players (see also McCabe, Rassenti and Smith, 1996).

game theory predicts that only rational egoists should prevail. The first six of the seven core findings summarized above were in part the stimulus for the development of the indirect evolutionary theory and the seventh is not inconsistent (see below for further discussion of it). Given the recent development of this approach, direct tests of this theory are not extensive. From the viewpoint of an indirect evolutionary process, participants in a collective action problem would start with differential, intrinsic preferences over outcomes due to their predispositions toward norms such as reciprocity and trust. Participants would learn about the likely behavior of others and shift their behavior in light of the experience and the objective payoffs they have received. Several recent experiments provide evidence of these kinds of contingent behaviors and behavioral shifts.[9]

In a one-shot, sequential, double-blind prisoner's dilemma experiment, for example, the players were asked to rank their preferences over the final outcomes after they had made their own choice, but before they knew their partner's decision. Forty percent of a pool of 136 subjects ranked the cooperative outcome (C,C) higher than the outcome if they defect while the other cooperates (D,C), and 27 percent were indifferent between these outcomes, even though their individual payoff was substantially higher for them in the latter outcome (Ahn, Ostrom and Walker, 1998).[10] This finding confirms that not all players enter a collective action situation as pure forward-looking rational egoists who make decisions based solely on individual outcomes. Some bring with them a set of norms and values that can support cooperation.

On the other hand, preferences based on these norms can be altered by bad experiences. After 72 subjects had played 12 rounds of a finitely repeated prisoner's dilemma game where partners were randomly matched each round, rates of cooperation were very low and many players had experienced multiple instances where partners had declined to cooperate, only 19 percent of the respondents ranked (C,C) above (D,C), while 17 percent were indifferent (Ahn et al., 1999). In this setting, the norms supporting cooperation and reciprocity were diminished, but not eliminated, by experience.

In another version of the prisoner's dilemma game, Cain (1998) first had players participate in a "dictator game"—in which one player divides a sum of money and the other player must accept the division, whatever it is—and then a prisoner's dilemma game. Stingy players, defined as those who retained at least 70 percent of their endowment in the earlier dictator game, tended to predict that all players would defect in the prisoner's dilemma game. Nice players, defined as those

---

[9] Further, Kikuchi, Watanabe and Yamagishi (1996) have found that those who express a high degree of trust are able to predict others' behavior more accurately than those with low levels of trust.

[10] To examine the frequency of nonrational egoist preferences, a group of 181 undergraduates was given a questionnaire containing a similar payoff structure on the first day of classes at Indiana University in January 1999. They were asked to rank their preferences. In this nondecision setting, 52 percent reflected preferences that were not consistent with being rational egoists; specifically, 27 percent ranked the outcome (C,C) over (D,C) and 25 percent were indifferent.

that gave away at least 30 percent of their endowment, tended to predict that other nice players would cooperate and stingy players would defect. Before playing the prisoner's dilemma game, players were told whether their opponent had been "stingy" or "nice" in the dictator game. Nice players chose cooperation in the prisoner's dilemma game 69 percent of the time when they were paired with other nice players and 39 percent of the time when they were paired with stingy players.

Finally, interesting experimental (as well as field) evidence has accumulated that externally imposed rules tend to "crowd out" endogenous cooperative behavior (Frey, 1994). For example, consider some paradoxical findings of Frohlich and Oppenheimer (1996) from a prisoner's dilemma game. One set of groups played a regular prisoner's dilemma game, some with communication and some without. A second set of groups used an externally imposed, incentive-compatible mechanism designed to enhance cooperative choices. In the first phase of the experiment, the second set gained higher monetary returns than the control groups, as expected. In the second phase of the experiment, both groups played a regular prisoner's dilemma game. To the surprise of the experimenters, a higher level of cooperation occurred in the control groups that played the regular prisoner's dilemma in both phases, especially for those who communicated on a face-to-face basis. The greater cooperation that had occurred due to the exogenously created incentive-compatible mechanism appeared to be transient. As the authors put it (p. 180), the removal of the external mechanism "seemed to undermine subsequent cooperation and leave the group worse off than those in the control group who had played a regular . . . prisoner's dilemma."

Several other recent experimental studies have confirmed the notion that external rules and monitoring can crowd out cooperative behavior.[11] These studies typically find that a social norm, especially in a setting where there is communication between the parties, can work as well or nearly as well at generating cooperative behavior as an externally imposed set of rules and system of monitoring and sanctioning. Moreover, norms seem to have a certain staying power in encouraging a growth of the desire for cooperative behavior over time, while cooperation enforced by externally imposed rules can disappear very quickly. Finally, the worst of all worlds may be one where external authorities impose rules but are only able to achieve weak monitoring and sanctioning. In a world of strong external monitoring and sanctioning, cooperation is enforced without any need for internal norms to develop. In a world of no external rules or monitoring, norms can evolve to support cooperation. But in an in-between case, the mild degree of external monitoring discourages the formation of social norms, while also making it attrac-

[11] Bohnet, Frey and Huck (1999) set up a sequential prisoner's dilemma, but add a regulatory regime where a "litigation process" is initiated if there is a breach of performance. Cardenas, Stranlund and Willis (2000) describe an experiment based on harvesting from a common-pool resource conducted in three rural villages in Columbia where exogenous but imperfect rule enforcement generated less cooperation than allowing face-to-face communication.

tive for some players to deceive and defect and take the relatively low risk of being caught.

The Evolution of Rules and Norms in the Field

Field studies of collective action problems are extensive and generally find that cooperation levels vary from extremely high to extremely low across different settings. (As discussed above, the seventh core finding from experimental research is that contextual factors affect the rate of contribution to public goods.) An immense number of contextual variables are also identified by field researchers as conducive or detrimental to endogenous collective action. Among those proposed are: the type of production and allocation functions; the predictability of resource flows; the relative scarcity of the good; the size of the group involved; the heterogeneity of the group; the dependence of the group on the good; common understanding of the group; the size of the total collective benefit; the marginal contribution by one person to the collective good; the size of the temptation to free ride; the loss to cooperators when others do not cooperate; having a choice of participating or not; the presence of leadership; past experience and level of social capital; the autonomy to make binding rules; and a wide diversity of rules that are used to change the structure of the situation (see literature cited in Ostrom, forthcoming).

Some consistent findings are emerging from empirical field research. A frequent finding is that when the users of a common-pool resource organize themselves to devise and enforce some of their own basic rules, they tend to manage local resources more sustainably than when rules are externally imposed on them (for example, Tang, 1992; Blomquist, 1992; Baland and Platteau, 1996; Wade, 1994). Common-pool resources are natural or humanly created systems that generate a finite flow of benefits where it is costly to exclude beneficiaries and one person's consumption subtracts from the amount of benefits available to others (Ostrom, Gardner and Walker, 1994). The users of a common-pool resource face a first-level dilemma that each individual would prefer that others control their use of the resource while each is able to use the resource freely. An effort to change these rules is a second-level dilemma, since the new rules that they share are a public good. Thus, users face a collective action problem, similar in many respects to the experiments discussed above, of how to cooperate when their immediate best-response strategies lead to suboptimal outcomes for all. A key question now is: How does evolutionary theory help us understand the well-established finding that many groups of individuals overcome both dilemmas? Further, how can we understand how self-organized resource regimes, that rarely rely on external third-party enforcement, frequently outperform government-owned resource regimes that rely on externally enforced, formal rules?

The Emergence of Self-Organized Collective Action

From evolutionary theory, we should expect some individuals to have an initial propensity to follow a norm of reciprocity and to be willing to restrict their own use of a common pool resource so long as almost everyone reciprocates. If a small core group of users identify each other, they can begin a process of cooperation without having to devise a full-blown organization with all of the rules that they might eventually need to sustain cooperation over time. The presence of a leader or entrepreneur, who articulates different ways of organizing to improve joint outcomes, is frequently an important initial stimulus (Frohlich, Oppenheimer and Young, 1971; Varughese, 1999).[12]

If a group of users can determine its own membership—including those who agree to use the resource according to their agreed-upon rules and excluding those who do not agree to these rules—the group has made an important first step toward the development of greater trust and reciprocity. Group boundaries are frequently marked by well-understood criteria, like everyone who lives in a particular community or has joined a specific local cooperative. Membership may also be marked by symbolic boundaries and involve complex rituals and beliefs that help solidify individual beliefs about the trustworthiness of others.

Design Principles of Long-Surviving, Self-Organized Resource Regimes

Successful self-organized resource regimes can initially draw upon locally evolved norms of reciprocity and trustworthiness and the likely presence of local leaders in most community settings. More important, however, for explaining their long-term survival and comparative effectiveness, resource regimes that have flourished over multiple generations tend to be characterized by a set of design principles. These design principles are extensively discussed in Ostrom (1990) and have been subjected to extensive empirical testing.[13] Evolutionary theory helps to explain how these design principles work to help groups sustain and build their cooperation over long periods of time.

We have already discussed the first design principle—the presence of clear boundary rules. Using this principle enables participants to know who is in and who is out of a defined set of relationships and thus with whom to cooperate. The second design principle is that the local rules-in-use restrict the amount, timing,

---

[12] Empirical studies of civil rights movements, where contributions can be very costly, find that organizers search for ways to assure potential participants of the importance of shared internal norms and that many others will also participate (Chong, 1991). Membership in churches and other groups that jointly commit themselves to protests and other forms of collective action is also an important factor (Opp, Voss and Gern, 1995).

[13] The design principles that characterize long-standing common-pool resource regimes have now been subject to considerable further empirical studies since they were first articulated (Ostrom, 1990). While minor modifications have been offered to express the design principles somewhat differently, no empirical study has challenged their validity, to my knowledge (Morrow and Hull, 1996; Asquith, 1999; Bardhan, 1999; Lam, 1998).

and technology of harvesting the resource; allocate benefits proportional to required inputs; and are crafted to take local conditions into account. If a group of users is going to harvest from a resource over the long run, they must devise rules related to how much, when, and how different products are to be harvested, and they need to assess the costs on users of operating a system. Well-tailored rules help to account for the perseverance of the resource itself. How to relate user inputs to the benefits they obtain is a crucial element of establishing a fair system (Trawick, 1999). If some users get all the benefits and pay few of the costs, others become unwilling to follow rules over time.

In long-surviving irrigation systems, for example, subtly different rules are used in each system for assessing water fees used to pay for maintenance activities, but water tends to be allocated proportional to fees or other required inputs (Bardhan, 1999). Sometimes water and responsibilities for resource inputs are distributed on a share basis, sometimes on the order in which water is taken, and sometimes strictly on the amount of land irrigated. No single set of rules defined for all irrigation systems in a region would satisfy the particular problems in managing each of these broadly similar, but distinctly different, systems (Tang, 1992; Lam, 1998).

The third design principle is that most of the individuals affected by a resource regime can participate in making and modifying their rules. Resource regimes that use this principle are both able to tailor better rules to local circumstances and to devise rules that are considered fair by participants. The Chisasibi Cree, for example, have devised a complex set of entry and authority rules related to the fish stocks of James Bay as well as the beaver stock located in their defined hunting territory. Berkes (1987, p. 87) explains that these resource systems and the rules used to regulate them have survived and prospered for so long because effective "social mechanisms ensure adherence to rules which exist by virtue of mutual consent within the community. People who violate these rules suffer not only a loss of favor from the animals (important in the Cree ideology of hunting) but also social disgrace." Fair rules of distribution help to build trusting relationships, since more individuals are willing to abide by these rules because they participated in their design and also because they meet shared concepts of fairness (Bowles, 1998).

In a study of 48 irrigation systems in India, Bardhan (1999) finds that the quality of maintenance of irrigation canals is significantly lower on those systems where farmers perceive the rules to be made by a local elite. On the other hand, those farmers (of the 480 interviewed) who responded that the rules have been crafted by most of the farmers, as contrasted to the elite or the government, have a more positive attitude about the water allocation rules and the rule compliance of other farmers. Further, in all of the villages where a government agency decides how water is to be allocated and distributed, frequent rule violations are reported and farmers tend to contribute less to the local village fund. Consistent with this is the finding by Ray and Williams (1999) that the deadweight loss from upstream farmers stealing water on government-owned irrigation systems in Maharashtra,

India, approaches one-fourth of the revenues that could be earned in an efficient water allocation and pricing regime.

Few long-surviving resource regimes rely only on endogenous levels of trust and reciprocity. The fourth design principle is that most long-surviving resource regimes select their own monitors, who are accountable to the users or are users themselves and who keep an eye on resource conditions as well as on user behavior. Further, the fifth design principle points out that these resource regimes use *graduated sanctions* that depend on the seriousness and context of the offense. By creating official positions for local monitors, a resource regime does not have to rely only on willing punishers to impose personal costs on those who break a rule. The community legitimates a position. In some systems, users rotate into this position so everyone has a chance to be a participant as well as a monitor. In other systems, all participants contribute resources and they hire monitors jointly. With local monitors, conditional cooperators are assured that someone is generally checking on the conformance of others to local rules. Thus, they can continue their own cooperation without constant fear that others are taking advantage of them.

On the other hand, the initial sanctions that are imposed are often so low as to have no impact on an expected benefit-cost ratio of breaking local rules (given the substantial temptations frequently involved). Rather, the initial sanction needs to be considered more as information both to the person who is "caught" and to others in the community. Everyone can make an error or can face difficult problems leading them to break a rule. Rule infractions, however, can generate a downward cascade of cooperation in a group that relies only on conditional cooperation and has no capacity to sanction (for example, Kikuchi et al., 1998). In a regime that uses graduated punishments, however, a person who purposely or by error breaks a rule is notified that others notice the infraction (thereby increasing the individual's confidence that others would also be caught). Further, the individual learns that others basically continue to extend their trust and want only a small token to convey a recognition that the mishap occurred. Self-organized regimes rely more on what Margaret Levi calls "quasi-voluntary" cooperation than either strictly voluntary or coerced cooperation (Levi, 1988). A real threat to the continuance of self-organized regimes occurs, however, if some participants break rules repeatedly. The capability to escalate sanctions enables such a regime to warn members that if they do not conform they will have to pay ever-higher sanctions and may eventually be forced to leave the community.

Let me summarize my argument to this point. When the users of a resource design their own rules (Design Principle 3) that are enforced by local users or accountable to them (Design Principle 4) using graduated sanctions (Design Principle 5) that define who has rights to withdraw from the resource (Design Principle 1) and that effectively assign costs proportionate to benefits (Design Principle 2), collective action and monitoring problems are solved in a reinforcing manner (Agrawal, 1999).

Individuals who think a set of rules will be effective in producing higher joint

benefits and that monitoring (including their own) will protect them against being a sucker are willing to undertake conditional cooperation. Once some users have made contingent self-commitments, they are then motivated to monitor other people's behavior, at least from time to time, to assure themselves that others are following the rules most of the time. Conditional cooperation and mutual monitoring reinforce one another, especially in regimes where the rules are designed to reduce monitoring costs. Over time, further adherence to shared norms evolves and high levels of cooperation are achieved without the need to engage in very close and costly monitoring to enforce rule conformance.

The operation of these principles is then bolstered by the sixth design principle that points to the importance of access to rapid, low-cost, local arenas to resolve conflict among users or between users and officials. Rules, unlike physical constraints, have to be understood to be effective. There are always situations in which participants can interpret a rule that they have jointly made in different ways. By devising simple, local mechanisms to get conflicts aired immediately and resolutions that are generally known in the community, the number of conflicts that reduce trust can be reduced. If individuals are going to follow rules over a long period of time, some mechanism for discussing and resolving what constitutes a rule infraction is necessary to the continuance of rule conformance itself.

The capability of local users to develop an ever-more effective regime over time is affected by whether they have minimal recognition of the right to organize by a national or local government. This is the seventh design principle. While some resource regimes have operated for relatively long times without such rights (Ghate, 2000), participants have had to rely almost entirely on unanimity as the rule used to change rules. (Otherwise, any temporarily disgruntled participant who voted against a rule change could go to the external authorities to threaten the regime itself!) Unanimity as a decision rule for changing rules imposes high transaction costs and prevents a group from searching for better matched rules at relatively lower costs.

Users frequently devise their own rules without creating formal, governmental jurisdictions for this purpose. In many in-shore fisheries, for example, local fishers devise extensive rules defining who can use a fishing ground and what kind of equipment can be used (Acheson, 1988; Schlager, 1994). As long as external governmental officials give at least minimal recognition to the legitimacy of such rules, the fishers themselves may be able to enforce the rules. But if external governmental officials presume that only they can make authoritative rules, then it is difficult for local users to sustain a self-organized regime (Johnson and Libecap, 1982).

When common pool resources are somewhat larger, an eighth design principle tends to characterize successful systems—the presence of governance activities organized in multiple layers of nested enterprises. The rules appropriate for allocating water among major branches of an irrigation system, for

example, may not be appropriate for allocating water among farmers along a single distributory channel. Consequently, among long-enduring self-governed regimes, smaller-scale organizations tend to be nested in ever-larger organizations. It is not unusual to find a large, farmer-governed irrigation system, for example, with five layers of organization each with its own distinct set of rules (Yoder, 1992).

Threats to Sustained Collective Action

All economic and political organizations are vulnerable to threats, and self-organized resource-governance regimes are no exception. Both exogenous and endogenous factors challenge their long-term viability. Here we will concentrate on those factors that affect the distribution of types of participants within a regime and the strength of the norms of trust and reciprocity held by participants. Major migration (out of or into an area) is always a threat that may or may not be countered effectively. Out-migration may change the economic viability of a regime due to loss of those who contribute needed resources. In-migration may bring new participants who do not trust others and do not rapidly learn social norms that have been established over a long period of time. Since collective action is largely based on mutual trust, some self-organized resource regimes that are in areas of rapid settlement have disintegrated within relatively short times (Baland and Platteau, 1996).

In addition to rapid shifts in population due to market changes or land distribution policies, several more exogenous and endogenous threats have been identified in the empirical literature (Sengupta, 1991; Bates, 1987; and literature cited in Ostrom, 1998b; Britt, 2000). These include: 1) efforts by national governments to impose a single set of rules on all governance units in a region; 2) rapid changes in technology, in factor availability, and in reliance on monetary transactions; 3) transmission failures from one generation to the next of the operational principles on which self-organized governance is based; 4) turning to external sources of help too frequently; 5) international aid that does not take account of indigenous knowledge and institutions; 6) growth of corruption and other forms of opportunistic behavior; and 7) a lack of large-scale institutional arrangements that provide fair and low-cost resolution mechanisms for conflicts that arise among local regimes, educational and extension facilities, and insurance mechanisms to help when natural disasters strike at a local level.

Contextual variables are thus essential for understanding the initial growth and sustainability of collective action as well as the challenges that long-surviving, self-organized regimes must try to overcome. Simply saying that context matters is not, however, a satisfactory theoretical approach. Adopting an evolutionary approach is the first step toward a more general theoretical synthesis that addresses the question of how context matters. In particular, we need to address how context affects the presence or absence of conditional cooperators and willing punishers

and the likelihood that the norms held by these participants are adopted and strengthened by others in a relevant population.

Conclusion

Both laboratory experiments and field studies confirm that a substantial number of collective action situations are resolved successfully, at least in part. The old-style notion, pre-Mancur Olson, that groups would find ways to act in their own collective interest was not entirely misguided. Indeed, recent developments in evolutionary theory—including the study of cultural evolution—have begun to provide genetic and adaptive underpinnings for the propensity to cooperate based on the development and growth of social norms. Given the frequency and diversity of collective action situations in all modern economies, this represents a more optimistic view than the zero contribution hypothesis. Instead of pure pessimism or pure optimism, however, the picture requires further work to explain why some contextual variables enhance cooperation while others discourage it.

Empirical and theoretical work in the future needs to ask how a large array of contextual variables affects the processes of teaching and evoking social norms; of informing participants about the behavior of others and their adherence to social norms; and of rewarding those who use social norms, such as reciprocity, trust, and fairness. We need to understand how institutional, cultural, and biophysical contexts affect the types of individuals who are recruited into and leave particular types of collective action situations, the kind of information that is made available about past actions, and how individuals can themselves change structural variables so as to enhance the probabilities of norm-using types being involved and growing in strength over time.

Further developments along these lines are essential for the development of public policies that enhance socially beneficial, cooperative behavior based in part on social norms. It is possible that past policy initiatives to encourage collective action that were based primarily on externally changing payoff structures for rational egoists may have been misdirected—and perhaps even crowded out the formation of social norms that might have enhanced cooperative behavior in their own way. Increasing the authority of individuals to devise their own rules may well result in processes that allow social norms to evolve and thereby increase the probability of individuals better solving collective action problems.

References

Acheson, James M. 1988. *The Lobster Gangs of Maine.* Hanover, NH: University Press of New England.

Agrawal, Arun. 1999. *Greener Pastures: Politics, Markets, and Community among a Migrant Pastoral People.* Durham, NC: Duke University Press.

Ahn, Toh-Kyeong, Elinor Ostrom, David Schmidt and James Walker. 1999. "Dilemma Games: Game Parameters and Matching Protocols." Bloomington: Indiana University, Workshop in Political Theory and Policy Analysis, Working paper.

Ahn, Toh-Kyeong, Elinor Ostrom and James Walker. 1998. "Trust and Reciprocity: Experimental Evidence from PD Games." Bloomington: Indiana University, Workshop in Political Theory and Policy Analysis, Working paper.

Asquith, Nigel M. 1999. *How Should the World Bank Encourage Private Sector Investment in Biodiversity Conservation?* A Report Prepared for Kathy MacKinnon, Biodiversity Specialist, The World Bank, Washington, D.C. Durham, North Carolina: Sanford Institute of Public Policy, Duke University.

Axelrod, Robert. 1986. "An Evolutionary Approach to Norms." *American Political Science Review.* December 80:4, pp. 1095–1111.

Baland, Jean-Marie and Jean-Philippe Platteau. 1996. *Halting Degradation of Natural Resources: Is There a Role for Rural Communities?* Oxford: Clarendon Press.

Bardhan, Pranab. 1999. "Water Community: An Empirical Analysis of Cooperation on Irrigation in South India." Berkeley: University of California, Department of Economics, Working paper.

Barkow, Jerome H., Leda Cosmides and John Tooby, eds. 1992. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture.* Oxford: Oxford University Press.

Bates, Robert H. 1987. *Essays on the Political Economy of Rural Africa.* Berkeley: University of California Press.

Berkes, Fikret. 1987. "Common Property Resource Management and Cree Indian Fisheries in Subarctic Canada," in *The Question of the Commons: The Culture and Ecology of Communal Resources.* Bonnie J. McCay and James Acheson, eds. Tucson: University of Arizona Press, pp. 66–91.

Blomquist, William. 1992. *Dividing the Waters: Governing Groundwater in Southern California.* San Francisco, CA: ICS Press.

Bohnet, Iris and Bruno S. Frey. 1999. "The Sound of Silence in Prisoner's Dilemma and Dictator Games." *Journal of Economic Behavior and Organization.* January, 38:1, pp. 43–58.

Bohnet, Iris, Bruno S. Frey and Steffen Huck. 1999. "More Order with Less Law: On Contract Enforcement, Trust, and Crowding." Cambridge, MA: Harvard University, Working paper.

Börgers, Tilman and Rajiv Sarin. 1997. "Learning Through Reinforcement and Replicator Dynamics." *Journal of Economic Theory.* 77, pp. 1–14.

Bowles, Samuel. 1998. "Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions." *Journal of Economic Literature.* March, 36, pp. 75–111.

Boyd, Robert and Peter J. Richerson. 1985. *Culture and the Evolutionary Process.* Chicago: University of Chicago Press.

Britt, Charla. 2000. "Forestry and Forest Policies." Bloomington: Indiana University, Workshop in Political Theory and Policy Analysis, Working paper.

Bromley, Daniel W. et al., eds. 1992. *Making the Commons Work: Theory, Practice, and Policy.* San Francisco, CA: ICS Press.

Cain, Michael. 1998. "An Experimental Investigation of Motives and Information in the Pris-

oner's Dilemma Game." *Advances in Group Processes.* 15, pp. 133–60.

Cameron, Lisa. 1995. "Raising the Stakes in the Ultimatum Game: Experimental Evidence from Indonesia." Princeton, NJ: Princeton University, Discussion paper.

Cardenas, Juan-Camilo, John K. Stranlund and Cleve E. Willis. 2000. "Local Environmental Control and Institutional Crowding-Out." *World Development.* Autumn, forthcoming.

Chong, Dennis. 1991. *Collective Action and the Civil Rights Movement.* Chicago: University of Chicago Press.

Clark, Andy and Annette Karmiloff-Smith. 1991. "The Cognizer's Innards: A Psychological and Philosophical Perspective on the Development of Thought." *Mind and Language.* Winter, 8:4, pp. 487–519.

Crawford, Sue E. S. and Elinor Ostrom. 1995. "A Grammar of Institutions." *American Political Science Review.* September, 89:3, pp. 582–600.

Cummins, Denise D. 1996. "Evidence of Deontic Reasoning in 3- and 4-Year-Olds." *Memory and Cognition.* 24, pp. 823–29.

Davis, Douglas D. and Charles A. Holt. 1993. *Experimental Economics.* Princeton, NJ: Princeton University Press.

Dawes, Robyn M., Jeanne McTavish and Harriet Shaklee. 1977. "Behavior, Communication, and Assumptions about Other People's Behavior in a Commons Dilemma Situation." *Journal of Personality and Social Psychology.* 35:1, pp. 1–11.

Epstein, Joshua M. and Robert Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up.* Cambridge, MA: MIT Press.

Eshel, Ilan, Larry Samuelson and Avner Shaked. 1998. "Altruists, Egoists, and Hooligans in a Local Interaction Model." *American Economic Review.* March, 88:1, pp. 157–79.

Feeny, David, Fikret Berkes, Bonnie J. McCay and James M. Acheson. 1990. "The Tragedy of the Commons: Twenty-Two Years Later." *Human Ecology.* 18:1, pp. 1–19.

Fehr, Ernst and Simon Gächter. Forthcoming. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review.* 90:1.

Fehr, Ernst and Klaus Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics.* 114:3, pp. 817–68.

Frank, Robert H., Thomas Gilovich and Dennis T. Regan. 1993. "The Evolution of One-Shot Cooperation: An Experiment." *Ethology and Sociobiology.* July, 14, pp. 247–56.

Frey, Bruno S. 1994. "How Intrinsic Motivation is Crowded Out and In." *Rationality and Society.* 6, pp. 334–52.

Frohlich, Norman and Joe A. Oppenheimer. 1996. "Experiencing Impartiality to Invoke Fairness in the N-PD: Some Experimental Results." *Public Choice.* 86, pp. 117–35.

Frohlich, Norman, Joe A. Oppenheimer and Oran Young. 1971. *Political Leadership and Collective Goods.* Princeton, NJ: Princeton University Press.

Ghate, Rucha. 2000. "The Role of Autonomy in Self-Organizing Process: A Case Study of Local Forest Management in India." Bloomington: Indiana University, Workshop in Political Theory and Policy Analysis, Working paper.

Güth, Werner. 1995. "An Evolutionary Approach to Explaining Cooperative Behavior by Reciprocal Incentives." *International Journal of Game Theory.* 24, pp. 323–44.

Güth, Werner and Hartmut Kliemt. 1998. "The Indirect Evolutionary Approach: Bridging the Gap between Rationality and Adaptation." *Rationality and Society.* August, 10:3, pp. 377–99.

Güth, Werner and Menahem Yaari. 1992. "An Evolutionary Approach to Explaining Reciprocal Behavior in a Simple Strategic Game," in *Explaining Process and Change. Approaches to Evolutionary Economics.* Ulrich Witt, ed. Ann Arbor: University of Michigan Press, pp. 23–34.

Hardin, Russell. 1971. "Collective Action as an Agreeable n-Prisoners' Dilemma." *Science.* September-October, 16, pp. 472–81.

Hardin, Russell. 1982. *Collective Action.* Baltimore, MD: Johns Hopkins University Press.

Hess, Charlotte. 1999. *A Comprehensive Bibliography of Common Pool Resources.* CD-ROM. Bloomington: Indiana University, Workshop in Political Theory and Policy Analysis.

Isaac, R. Mark, James Walker and Arlington W. Williams. 1994. "Group Size and the Voluntary Provision of Public Goods: Experimental Evidence Utilizing Large Groups." *Journal of Public Economics.* May, 54:1, pp. 1–36.

Johnson, Ronald N. and Gary D. Libecap. 1982. "Contracting Problems and Regulation: The Case of the Fishery." *American Economic Review.* December, 27:5, pp. 1005–1023.

Kagel, John and Alvin Roth, eds. 1995. *The Handbook of Experimental Economics.* Princeton, NJ: Princeton University Press.

Kikuchi, Masao, Yoriko Watanabe and Toshio Yamagishi. 1996. "Accuracy in the Prediction of Others' Trustworthiness and General Trust: An Experimental Study." *Japanese Journal of Experimental Social Psychology.* 37:1, pp. 23–36.

Kikuchi, Masao, Masako Fujita, Esther Marciano and Yujiro Hayami. 1998. "State and Community in the Deterioration of a National Irrigation System." Paper presented at the World Bank-EDI Conference on "Norms and Evolution in the Grassroots of Asia," Stanford University, February 6–7.

Kreps, David M., Paul Milgrom, John Roberts and Robert Wilson. 1982. "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma." *Journal of Economic Theory.* 27, pp. 245–52.

Lam, Wai Fung. 1998. *Governing Irrigation Systems in Nepal: Institutions, Infrastructure, and Collective Action.* Oakland, CA: ICS Press.

Ledyard, John. 1995. "Public Goods: A Survey of Experimental Research," in *The Handbook of Experimental Economics.* John Kagel and Alvin Roth, eds. Princeton, NJ: Princeton University Press, pp. 111–94.

Levi, Margaret. 1988. *Of Rule and Revenue.* Berkeley: University of California Press.

Lichbach, Mark Irving. 1996. *The Cooperator's Dilemma.* Ann Arbor: University of Michigan Press.

Manktelow, Ken I. and David E. Over. 1991. "Social Roles and Utilities in Reasoning with Deontic Conditionals." *Cognition.* 39, pp. 85–105.

McCabe, Kevin A., Stephen J. Rassenti and Vernon L. Smith. 1996. "Game Theory and Reciprocity in Some Extensive Form Experimental Games." *Proceedings of the National Academy of Sciences.* November, 93, 13,421–13,428.

McCabe, Kevin A. and Vernon L. Smith. 1999. "Strategic Analysis by Players in Games: What Information Do They Use." Tucson: University of Arizona, Economic Research Laboratory, Working paper.

Milgrom, Paul R., Douglass C. North and Barry R. Weingast. 1990. "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs." *Economics and Politics.* March, 2:1, pp. 1–23.

Montgomery, Michael R. and Richard Bean. 1999. "Market Failure, Government Failure, and the Private Supply of Public Goods: The Case of Climate-Controlled Walkway Networks." *Public Choice.* June, 99:3/4, pp. 403–37.

Morrow, Christopher E. and Rebecca Watts Hull. 1996. "Donor-Initiated Common Pool Resource Institutions: The Case of the Yanesha Forestry Cooperative." *World Development.* 24:10, pp. 1641–1657.

Nowak, Martin A. and Karl Sigmund. 1998. "Evolution of Indirect Reciprocity by Image Scoring." *Nature.* 393:6685, pp. 573–77.

Oaksford, Mike and Nick Chater. 1994. "A Rational Analysis of the Selection Task as Optimal Data Selection." *Psychological Review.* 101:4, pp. 608–31.

Offerman, Theo. 1997. *Beliefs and Decision Rules in Public Goods Games: Theory and Experiments.* Dordrecht, the Netherlands: Kluwer Academic Publishers.

Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups.* Cambridge, MA: Harvard University Press.

Opp, Karl-Dieter, Peter Voss and Christiana Gern. 1995. *Origins of Spontaneous Revolution.* Ann Arbor: University of Michigan Press.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action.* New York: Cambridge University Press.

Ostrom, Elinor. 1998a. "A Behavioral Approach to the Rational Choice Theory of Collective Action." *American Political Science Review.* March, 92:1, pp. 1–22.

Ostrom, Elinor. 1998b. "Institutional Analysis, Design Principles, and Threats to Sustainable Community Governance and Management of Commons," in *Law and the Governance of Renewable Resources: Studies from Northern Europe and Africa.* Erling Berge and Nils Christian Stenseth, eds. Oakland, CA: ICS Press, pp. 27–53.

Ostrom, Elinor. Forthcoming. "Reformulating the Commons," in *The Commons Revisited: An Americas Perspective.* Joanna Burger, Richard Norgaard, Elinor Ostrom, David Policansky, and Bernard Goldstein, eds. Washington, DC: Island Press.

Ostrom, Elinor, Roy Gardner and James Walker. 1994. *Rules, Games, and Common-Pool Resources.* Ann Arbor: University of Michigan Press.

Ostrom, Elinor and James Walker. 1997. "Neither Markets Nor States: Linking Transformation Processes in Collective Action Arenas," in *Perspectives on Public Choice: A Handbook.* Dennis C. Mueller, ed. Cambridge: Cambridge University Press, pp. 35–72.

Pinker, Steven. 1994. *The Language Instinct.* New York: W. Morrow.

Posner, Richard A. and Eric B. Rasmusen. 1999. "Creating and Enforcing Norms, with Special Reference to Sanctions." *International Review of Law and Economics.* September, 19:3, pp. 369–82.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review.* 83, pp. 1281–1302.

Ray, Ishar and Jeffrey Williams. 1999. "Evalu-

ation of Price Policy in the Presence of Water Theft." *American Journal of Agricultural Economics.* November, 81, pp. 928–41.

Schlager, Edella. 1994. "Fishers' Institutional Responses to Common-Pool Resource Dilemmas," in *Rules, Games, and Common-Pool Resources.* Elinor Ostrom, Roy Gardner, and James Walker, eds. Ann Arbor: University of Michigan Press, pp. 247–65.

Selten, Reinhard. 1991. "Evolution, Learning, and Economic Behavior." *Games and Economic Behavior.* February, 3:1, pp. 3–24.

Sengupta, Nirmal. 1991. *Managing Common Property: Irrigation in India and the Philippines.* New Delhi: Sage.

Sethi, Rajiv and Eswaran Somanathan. 1996. "The Evolution of Social Norms in Common Property Resource Use." *American Economic Review.* September, 86:4, pp. 766–88.

Tang, Shui Yan. 1992. *Institutions and Collective Action: Self Governance in Irrigation.* San Francisco, CA: ICS Press.

Trawick, Paul. 1999. "The Moral Economy of Water: 'Comedy' and 'Tragedy' in the Andean Commons." Lexington: University of Kentucky, Department of Anthropology, Working paper.

Varughese, George. 1999. "Villagers, Bureaucrats, and Forests in Nepal: Designing Governance for a Complex Resource." Ph.D. dissertation, Indiana University.

Wade, Robert. 1994. *Village Republics: Economic Conditions for Collective Action in South India.* San Francisco, CA: ICS Press.

Yoder, Robert D. 1992. *Performance of the Chhattis Mauja Irrigation System, a Thirty-five Hundred Hectare System Built and Managed by Farmers in Nepal.* Colombo, Sri Lanka: International Irrigation Management Institute.