

BACKGROUND PAPER TO THE 2010 WORLD DEVELOPMENT REPORT

Implications for Climate-Change Policy of Research on Cooperation in Social Dilemmas

Timothy Irwin

The World Bank
Development Economics
Office of the Senior Vice President and Chief Economist
July 2009



Abstract

The problem of climate change seems to be a tragedy of the commons: despite the global benefits of reducing green-house gas emissions, no individual has any incentive to reduce his or her own emissions. Yet many people are making efforts to reduce emissions and putting pressure on businesses and governments to do the same. Although the size of these efforts is unclear, their very existence might seem puzzling. The efforts are consistent, however, with some theoretical and empirical evidence about the extent of cooperation in other social

dilemmas. This evidence does not imply that greenhouse-gas emissions will be reduced to desirable levels, but it does suggest that the potential for voluntary cooperation should not be ignored. It also suggests that cooperation can be promoted by (i) allowing cooperators to punish defectors without withdrawing their own cooperation; (ii) publicly emphasizing the social benefits and extent of cooperation and the social norms that require it; and (iii) improving the quantity and timeliness of public information about cooperation and defection.

This paper—prepared as a background paper to the World Bank’s *World Development Report 2010: Development in a Changing Climate*—is a product of the Development Economics Vice Presidency. The views expressed in this paper are those of the authors and do not reflect the views of the World Bank or its affiliated organizations. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The author may be contacted at tirwin@lecg.com.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

**Implications for climate-change policy of research on cooperation in
social dilemmas**

Background paper for the
World Development Report 2010 on climate change

Timothy Irwin*

* This paper was written in early 2008, while the author was at the World Bank. He can now be reached at tirwin@lecg.com

Introduction

Climate change presents the world with a social dilemma. We would be better off if we reduced greenhouse-gas emissions, but each of us bears only an infinitesimal share of the cost of our own emissions. So, if we are rational and self-interested, it seems we won't do anything to reduce our emissions. Smaller-scale social dilemmas can be solved if a local or national government imposes a tax or quota on the harmful activity, but the absence of a global government means that the problem of climate change cannot easily be solved in this way. Even the largest countries bear only a small share of the cost of their emissions (about 15 percent) and small countries bear a negligible share.¹ So national governments acting in the interests of rational and self-interested citizens seem unlikely to do much to reduce emissions in their respective countries. Although they may sign international agreements to limit emissions, pressures from rational and self-interested citizens seem likely to discourage them from enforcing any agreement that requires costly changes.

This pessimistic view is consistent with the fact that global greenhouse-gas emissions are continuing to rise, despite a growing consensus that they are dangerous and numerous calls to cut them. Yet it is inconsistent with the fact that many people seem to be voluntarily reducing emissions at some personal cost (e.g., Chafe and French 2008). Some are prepared to pay more for locally grown food. Some prefer to buy more-expensive power from retailers that generate power from renewable sources. Others purchase carbon credits to offset the emissions caused by their air travel. Responding to these concerns, and those of some employees and investors, many businesses are also reducing or offsetting emissions, even when they are under no legal obligation to do so. Some facilitate the purchase of offsets by their customers. Others get certified as carbon neutral. Some governments, too, have taken unilateral actions to reduce emissions. Of course, most of the sacrifices people make are small, and it's hard to know how significant they are in aggregate, but their very existence is puzzling if we assume that self-interest and rationality will lead people to do nothing.

To understand why people might voluntarily reduce emissions and urge businesses and governments to do the same, we can look at research on cooperation in other social dilemmas in which cooperation is socially beneficial but is, or appears to be, individually costly. The last two decades have seen a proliferation of such research in disciplines as diverse as biology, economics, psychology, and anthropology. Although the emphasis of the research varies from discipline to discipline, much of it is united not only by a common theme but also by common use of the language and concepts of game theory—or, as it might better be called, “interactive decision theory” (Aumann 1987).

¹ In 2000, the most recent year for which comprehensive data are available, the largest emitter was the United States, which had an estimated 14.91% of global emissions. The European Union had an estimated 11.23% of global emissions.

Game theory has often been used to analyze climate-change treaties.² The analyses generally assume that the players are rational and self-interested and they generally are pessimistic. This paper reviews a different strand of research, which focuses on the behavior of individuals—who as well as affecting emissions directly also influence the behavior of governments in treaty negotiations. Some of this research asks how much cooperation can be expected of rational and self-interested individuals. But some of it emphasizes imperfect rationality or preferences about what happens to others.³ And some of it assumes only that successful behavior tends to become more common over time and asks when cooperation will evolve despite incentives not to cooperate. Although a few studies in this literature are explicitly about cooperation to reduce climate change, most are not.

The problem of climate change is more complex than the games considered here. The range of decisions about greenhouse-gas emissions is enormous, and there is great uncertainty about the payoffs associated with different combinations of actions. The payoffs also differ greatly among players; if climate change is moderate, some people may benefit. The players themselves are diverse, including individuals, businesses, and governments, each with its own decision-making style. And the number of players is enormous. So we can't assume that the research will yield conclusions that can immediately be applied to the problem of climate change. We can hope only that it captures important features of the problem and that it throws some light on how much cooperation can be expected and what can be done to increase cooperation.

Two-person climate-change games

The best models of the dilemma created by climate change are multiperson games. But two-person games are simpler, and prominent in research, so it is useful to consider them briefly before turning to their more-realistic multiperson counterparts.

The two-person game that best captures the tragic aspect of climate change is the well-known prisoners' dilemma (figure 1), in which cooperation is interpreted as reducing emissions and defection as continuing with business as usual. In a true one-shot prisoners' dilemma, in which payoffs represent preferences, not material outcomes like money received, defection is the only rational option. If we assume that the problem of climate change can be modeled as a true prisoners' dilemma, we assume that each player prefers business-as-usual to reducing emissions whatever the other player does. Thus, business-as-usual dominates emissions reductions and the only Nash equilibrium is for each player to choose businesses-as-usual. Backward induction shows that the same thing

² The IPCC's working group on mitigation discusses this research in Gupta *et al* 2007 and Toth *et al* 2001. Examples include Barrett 2003 and Dutta and Radner 2004.

³ World Bank 2005, ch. 4, reviewed some of this research for a different purpose.

is true of a finitely repeated prisoners' dilemma, in which the individual game is repeated a fixed and known number of times.

Figure 1 Four problems of cooperation

Prisoners' dilemma	Prisoners' delight	Stag hunt	Snowdrift
C D	C D	C D	C D
C 2 0	C 4 3	C 2 0	C 2 1
D 3 1	D 3 2	D 1 1	D 3 0

C stands for cooperate, D for defect. The numbers represent the payoff to the row player, given the choice of the column player. The games are symmetric, so only one player's payoffs need be shown.

Cooperation in one-shot games. People do sometimes cooperate in what seem to be one-shot or finitely repeated prisoners' dilemmas (Colman 1999, ch. 7). One possible explanation is that the players are aren't perfectly rational—backward induction, in particular, requires sophisticated reasoning—or that at least one of the players suspects that the other might not be perfectly rational (Kreps *et al* 1982). Another possible explanation is that the game isn't a true prisoners' dilemma, despite appearances. The material outcomes may parallel the payoffs of a prisoners' dilemma, but the players' preferences need not be determined only by those outcomes. For example, if the two prisoners in the original story of the prisoners' dilemma were brothers who cared as much about the other's prison sentence as about their own, they would be playing "the prisoners' delight" of figure 1 (Binmore 2007, 11). In this game, cooperation is a dominant strategy, and joint cooperation is the only Nash equilibrium.

Social norms. Potentially more important for climate change are social norms that influence choices even when people have no special concern for one another. Consider, for example, a norm that prescribes cooperation with cooperators.⁴ If the norm is strong enough relative to the material benefits of defection, it transforms the prisoners' dilemma into a game in which rational cooperation is possible. For example, if it has the effect of reducing by two units the payoff to defecting when the other player cooperates, it transforms the prisoners' dilemma of figure 1 into the stag hunt of figure 1 (Skyrms 2004). In this game, joint cooperation is the best outcome, but each player risks something by cooperating because it works only if both players cooperate.⁵ Joint cooperation and joint defection are both Nash equilibria of a stag hunt, so cooperation is not assured, but nor is it ruled out.

⁴ How cooperative norms might arise and why they might persist in the face of material incentives to defect is of course a crucial question, which is briefly discussed later. See also Binmore 1994, 1998; Skyrms 1996; Sugden 1986/2004; and Young 1999.

⁵ One definition of a stag hunt requires the joint-defection equilibrium to have a larger basin of attraction than the cooperative equilibrium and thus to be risk-dominant. For this, we can add a positive amount less than one to the payoffs to defection in figure 1.

The potential significance of norms can be seen in experiments with the ultimatum game, in which one player, the proposer, offers part of a given sum of money to the other player, the responder (Güth, Schmittberger, and Schwarz 1982). If the responder accepts the offer, the proposer keeps the rest. If the responder rejects the offer, no one gets anything. The game is not repeated. If the proposer and the responder both want only to maximize their wealth, the proposer should offer the smallest possible amount, anticipating that the responder will realize that any amount is better than nothing. This outcome is not the game's only Nash equilibrium, but it is the only subgame-perfect equilibrium.⁶ In fact, proposers typically offer 40–50 percent of the available sum, and many responders reject offers below 30 percent (Camerer 2003, ch. 1). One interpretation of these results is that proposers offer something partly because they follow a norm that prescribes the sharing of gains and partly because they believe that respondents will follow a norm that prescribes punishing those who don't share.⁷

Changing technology. Games can also be transformed by changes in technology. Proposals for subsidizing research on clean energy can be viewed as proposals to transform a prisoners'-dilemma-like game into one more favorable to cooperation. If technological progress reduces the cost of clean energy sufficiently, reducing emissions becomes cheaper than business as usual, and the emissions game becomes the prisoners' delight. Paying for the research necessary to change the technology is of course a problem of cooperation in its own right, but it may be a less-daunting problem. At the very least, the game has fewer players. It may also be that each of several countries can devote enough funding to the problem to solve it, even though each would prefer that another country bear the budgetary burden. The game would then be like the snowdrift of figure 1, in which the cooperative outcome, clearing the snowdrift, requires that at least one person do some work, but no one wants to be that person.⁸ Joint cooperation is not a Nash equilibrium of this game, but there are two Nash equilibria (in which one player cooperates while the other defects) that get the job done.

Repeated games. Cooperation may also emerge if the apparent prisoners' dilemma is part of indefinitely long series of games that analyzed as a whole is not a prisoners' dilemma. Enlightened self interest can then lead each player to cooperate in the hope of inducing the other player to cooperate—an idea that has been formalized in the so-called folk theorem of game theory (Fudenberg and Maskin 1986). An implication of the theorem is that, if the probability of repetition is great enough and if the future is important enough

⁶ A subgame-perfect equilibrium is, roughly, a Nash equilibrium whose component strategies involve only credible threats.

⁷ Most of the research on ultimatum games has been done in Europe and the United States. Henrich et al 2006 found that the norms that influence decisions vary by culture. In some groups, high offers are rejected as well as low offers. In others, behavior is not far from the subgame-perfect equilibrium for wealth maximizers.

⁸ Doebeli et al 2004. Other names for similar games are chicken and hawk–dove.

to the players, a repeated prisoners' dilemma has numerous equilibria, including one in which both players cooperate.

The best-known demonstration of the possibility of cooperation in a repeated prisoners' dilemma occurred when Robert Axelrod invited people to submit strategies for playing a repeated prisoners' dilemma in computerized round-robin tournaments and found that the winning strategy was tit-for-tat (Axelrod 1984). Tit-for-tat cooperates in the first round of the game and then does whatever the other player did in the previous round. It responds to cooperation with cooperation and to defection with defection. In testing the strategies, Axelrod also made use of evolutionary approach that had recently been developed by biologists interested in game theory (Maynard Smith 1982). In this test, strategies proliferated or dwindled over time in proportion to the payoffs they received. At the end of the test, tit-for-tat was the most common of many surviving strategies. Axelrod further argued that tit-for-tat was an *evolutionarily stable strategy*,⁹ which would mean that a population of tit-for-tat players could resist invasion by *any* mutant strategy, not just those used in his simulation (Axelrod and Hamilton 1981). Axelrod concluded that successful strategies in repeated prisoners' dilemmas were, like tit-for-tat, nice (they cooperated initially), provokable (they retaliated in response to defection), and forgiving (they resumed cooperation in response to renewed cooperation).

Michael Liebreich has recently argued that Axelrod's research shows that the prospects for cooperation in tackling climate change are better than is believed and that the research should inform countries' strategies: "The US needs to start being Nice, Europe needs to learn to Retaliate, and the developing world needs to Forgive" (Liebreich 2007). Axelrod's conclusions cannot be applied to the problem of climate change so quickly, though. On the one hand, we need to look at the multiplayer games before drawing conclusions. On the other, some of Axelrod's conclusions turned out to be wrong.

Tit-for-tat is not actually an evolutionarily stable strategy in the two-person prisoner's dilemma, so there are circumstances in which it can be displaced by other strategies.¹⁰ And although tit-for-tat did better than other strategies in the particular simulations that

⁹ A strategy *A* is an evolutionarily stable if *A* is either a better reply to *A* than is any other strategy or, if there is another strategy *B* that is an equally good reply to *A*, then *A* is a better reply to *B* than is *B* itself (Maynard Smith 1982, 14). A profile of evolutionarily stable strategies is necessarily a Nash equilibrium, but not every Nash equilibrium corresponds to profile of evolutionarily stable strategies.

¹⁰ Unconditional cooperation is as good a reply to tit-for-tat as tit-for-tat is, and in violation of the second part of the definition of an evolutionarily stable strategy (footnote 9) tit-for-tat is not a better response to unconditional cooperation than is unconditional cooperation itself. In practice, this means that random drift in a finite population can lead to an outcome in which tit-for-tat is replaced by unconditional cooperators. Unconditional cooperators are of course vulnerable to unconditional defectors, among other strategies. Moreover, combinations of strategies can invade tit-for-tat (Boyd and Lorberbaum 1987).

Axelrod ran, it doesn't do well in others. In particular, tit-for-tat does badly if the players sometimes make mistakes—if they sometimes defect when they mean to cooperate and vice versa or sometimes misinterpret the other player's action. When two tit-for-tat players play each other, an accidental defection leads to a series of echoing defections.

In simulations that incorporate mistakes, a more generous form of tit-for-tat that forgives some defections does better than pure tit-for-tat (Nowak and Sigmund 1992). In simulations that include a larger class of strategies, Pavlov does even better.¹¹ Pavlov is nasty: it initially defects. It then changes its action if and only if the other player defected in the previous round. If the other player cooperated, it repeats its previous action. Pavlov does poorly against an unconditional defector, because it keeps trying cooperation, but two Pavlovs quickly settle into a routine of joint cooperation, and like generous tit-for-tat, they avoid an accidentally triggered series of defections. And, unlike generous or pure tit-for-tat, Pavlov exploits unconditional cooperators. But even Pavlov's success depends on the details of the set-up. If moves are made sequentially rather than simultaneously, a strategy a bit like generous tit-for-tat reemerges (Freat 1994). Consistent with this diversity, an experiment found that real people play a variety of conditionally cooperative strategies, some similar to Pavlov, others similar to generous tit-for-tat (Wedekind and Milinski 1996). Overall, the evidence confirms that conditional cooperation can emerge in repeated two-person prisoners' dilemmas, but it also shows that there is no universally best strategy and indeed no equilibrium. Among other things, nice strategies don't always win.

Two other issues are worth mentioning before turning to multiplayer games. First, early research on repeated prisoners' dilemmas generally assumed that players encountered each other randomly. In reality, people are more likely to interact with neighbors, colleagues, and people they've successfully cooperated with before. These correlations tend to favor the emergence of cooperation (Bergstrom 2006; Nowak 2006; Skyrms 1996). Second, the cooperation that emerges in the models discussed above relies on direct reciprocity. You help me in the expectation that I will later help you. Cooperation in repeated games can also emerge as a form of *indirect reciprocity* (Nowak and Sigmund 2005). People can establish a reputation for cooperating, and their strategies can take account of whether the players they encounter have a reputation for cooperating. You cooperate with me, because I cooperated with others in the past, and you want others to cooperate with you in the future. Indirect reciprocity can lead to cooperation if people have enough information about each other's reputations.

¹¹ See Nowak and Sigmund 1993 and Binmore 1994, 3.2. Other names for the strategy include simpleton, win-stay-lose-shift, and tat-for-tit.

Social dilemmas

To get an idea of the prospects for cooperation in reducing greenhouse-gas emissions, we need to look at similar games involving many players. In the simplest such game, the multiperson prisoners' dilemma, each of many people must each decide whether to cooperate or to defect. When the game is played once or a fixed and known number of times, defection is again a dominant strategy for each player and joint defection is the only Nash equilibrium. As in the two-person game, however, joint defection is worse for everyone than joint cooperation. In more complex games, the players may have choices that lie between full cooperation and full defection. But the defining feature of this larger class of social dilemmas¹² is still that rational and self-interested individuals choose a level of cooperation that falls short of the social optimum.

Social dilemmas can be framed positively or negatively. In a public-good game, each player chooses whether or not to contribute to a public good. Here, cooperation creates a positive externality. In a public-bad game, each player chooses whether or not to consume something that creates a negative externality. Although the games are framed differently, the payoffs need not differ. In the case of greenhouse-gas emissions, we can think of the public good as reducing emissions below the business-as-usual level and the public bad as increasing emissions above the desired level.

The payoff functions may also differ among social dilemmas. In a *linear public-good game*, the difference between the social and private benefit of contributing does not vary with the level of contributions. The socially optimal outcome is for each player to contribute everything he has, but the rational and self-interested player contributes nothing. In a nonlinear *common-pool-resource game*, players decide how much to take from a forest, fishery, or other resource. For low total levels of taking, the private benefits of taking exceed the social cost, but after some point the private benefits are less than the social cost. The optimal amount of taking is thus positive, but less than the amount chosen in a one-shot game by a group of rational and self-interested players.

The folk theorem says that joint cooperation is an equilibrium of any indefinitely repeated social dilemma, just as it is of the two-person prisoners' dilemma, so long as the probability of repetition is high enough and the players' discount rates are low enough. But multiperson games are less favorable to cooperation than their two-person counterparts. In the multiperson prisoners' dilemma, for example, each player can infer from his own payoff what the other player did in the previous round and he can then reward or punish the other player accordingly. In the multiperson prisoners' dilemma, the players generally cannot infer from their payoffs what each other player did in the

¹² Dawes 1980. This class of games is also called "a problem of collective action" (Olson 1965/1971), "a tragedy of the commons" (Hardin 1968), "a public-goods game" and "multiperson" or "n-person" prisoners' dilemma.

previous round. Nor can they always punish defectors while rewarding cooperators, because in the previous round some players may have cooperated while others defected.¹³

Public-good games. Much of the experimental evidence on cooperation in social dilemmas comes from public-good games. In a typical experiment, each of four players might be given \$20. Each dollar contributed generates a social return of \$2, but this return is divided among four players, so the private return is only 50 cents. The players can keep any money they don't contribute. If no one contributes anything, they all get \$20. If they all invest \$20, they all get \$32. The game may be played just once or a fixed number of times. In the latter case, a single group may play several rounds together or the groups for each round may be formed in such a way that two players never meet twice, which ensures that each game is one-shot.

Several results have emerged from experiments with public-good games. In a one-shot game and in the first round of a finitely repeated game, the average player contributes about 50 percent of his endowment (Croson 2008). There are, however, big differences among players. Perhaps 30–50 percent of the players contribute nothing, while the rest cooperate in varying degrees (Ledyard 1995). Moreover, in finitely repeated games, cooperation tends to decline over the rounds of the game, so that by the last round the average contribution may be only a quarter or so of the endowment (Croson 2008). The decline may occur partly because the players are learning the game, but this doesn't seem to be the whole explanation. Some players cooperate even though they seem to understand the game and apparently stop cooperating because they are disappointed by the other players' contributions (Andreoni 1989, 1995a; Dawes and Thaler 1998).

Common-pool-resource games. The evidence from experimental common-pool-resource games is less promising for cooperation. Elinor Ostrom and her colleagues have found choices that varied apparently randomly but were on average similar to those predicted of rational and self-interested players, or perhaps worse (Ostrom, Gardner, and Walker 1994). Experiments involving price-setting in oligopolistic markets, in which the structure of the payoffs is similar, also find outcomes close to that predicted of selfish and rational players (Andreoni 1995b). This is curious, since the payoffs in common-pool-resource game are similar to those of the public-good game.

By contrast, Elinor Ostrom and her colleagues find that players behave quite cooperatively when they can communicate with each other, even though they cannot enforce any agreements they make. Field studies of real common-pool resources also find that people sometimes cooperate to protect natural resources in circumstances in which a tragedy of the commons might be predicted (Ostrom, Gardner, and Walker 1994). To take

¹³ Boyd and Richerson 1988 show formally that “the conditions that allow the evolution of reciprocal cooperation [in a multiperson prisoners' dilemma] become extremely restrictive as group size increases.”

just one example, Maine lobstermen managed to avoid overfishing by dividing up the coast into territories, each allocated to a different group of men who, despite the absence of legal property rights, were able to keep out others (Acheson 1987). And of course real oligopolies sometimes succeed in avoiding cutthroat competition.

There are also common-pool-resource experiments that do find significant cooperation. Cardenas, Stranlund, and Willis (2000) went to three rural villages in Colombia and had groups of eight villagers choose how much to take from a hypothetical forest, in a set-up in which the players could see that the pursuit of self-interest was socially suboptimal. As in the public-good experiments, the level of cooperation in the games was less than the social optimum, but more than would be predicted of rational and self-interested players.

Threshold public-good games. Public-good games can be modified by stipulating that the public good is provided if and only if the players' collective contributions reach a threshold (or "provision point"). When such a threshold is added to a public-good game, joint defection remains an equilibrium for a group of rational and self-interested players, but all profiles of contributions that sum to the threshold are also equilibria. Thus the game is no longer a social dilemma as defined; it is more like a snowdrift. Groups often succeed in cooperating in such games (Croson 2008).

Particularly interesting in this regard is an experiment that was framed in terms of climate change (Milinski *et al* 2008). Milinski and his colleagues proceed from the assumption that climate change is tolerable if the atmospheric concentration of greenhouse gases remains below some threshold, but may otherwise be very dangerous. They therefore model climate change as a kind of threshold public-good game. If the experimental subjects collectively gave enough to a "climate change" account, they could keep what remained in their private accounts at the end of the game. If they failed, they risked losing all their money. Specifically, 180 students were divided into thirty groups of six, and each student was given 40 euros. Each group played ten rounds of a game in which each player could give 0, 2, or 4 euros to the climate account. If the group's climate account reached 120 euros by the end of the tenth round, everyone kept whatever money they had left. If everyone gave 2 euros in each round, the group would reach the target and everyone would make 20 euros. If the 120-euro target wasn't reached, there was a chance that everything was lost. When the chance was 90 percent, half the groups reached the target and the other half just failed. When the chance was 10 or 50 percent, however, only one group reached the target. Milinski *et al* conclude "that one possible strategy to relieve the collective-risk dilemma in high-risk situations is to convince people that failure to invest enough is very likely to cause grave financial loss to the individual" (p. 2291).

Factors that increase cooperation in social dilemmas

The experiments also show that the extent of cooperation in social dilemmas is sensitive to variations in the nature or context of dilemma, even when the variations do not change

the fact that defection remains a dominant strategy for self-interested players. In particular, cooperation can be increased by allowing people to communicate with each other before they make their choice, even though they cannot make binding agreements (Dawes 1980; Colman 1999, ch. 9); exposing the players to “moralizing” and a discussion of the social benefits of cooperation (Dawes 1980); encouraging people to believe that the other players will cooperate (Seabright 1993; Frank, Gilovich, and Regan 1993); and publicly disclosing the players’ contributions (Dawes 1980).

Group size. Theory suggests that cooperation becomes more difficult as groups become bigger. Some field studies of common-pool resources find cooperation is indeed more likely among relatively small and cohesive groups (Ostrom, Gardner, and Walker 1994). One review of experimental studies of one-shot or finitely repeated public-good games also concludes that cooperation decreases as group size increases, at least when the group has fewer than eight members (Colman 1999, ch. 9). But another review of experimental studies concludes that group size has no effect on cooperation (Ledyard 1995).

Framing. The framing of a social dilemma can also affect cooperation, even though it doesn’t change anyone’s payoffs. James Andreoni finds that people are more cooperative if their choice is framed as whether or not to create a positive externality than if it is framed as whether or not to create a negative externality (Andreoni 1995b). This would explain why there is apparently more cooperation in public-good games than in common-pool-resource games. Thus, if people frame their choice about greenhouse-gas emissions as “should I cut back emissions to avoid harming other people?” they may choose a higher level of emissions than if they frame the choice as “should I help people by choosing a low level of emissions?” Studies of decision making under uncertainty also find framing effects, including loss aversion and risk-seeking in the face of possible losses (Tversky and Kahneman 1981). So the effects of the framing of emissions choices may be complex.

Reputational spillovers. People cooperate more when they can take advantage of their reputation for cooperating. Charities, for example, understand that they can raise more money if they ensure that donors can publicize their gifts (Harbaugh 1998). It has been shown that cooperation in a public-good game is an evolutionarily stable strategy when this game is alternated with an indirect-reciprocity game (Panchanathan and Boyd 2004). And it has been found in an experiment that the alternation of a public-good game and an indirect-reciprocity game does sustain high-levels of cooperation in the public-good game (Milinski *et al* 2002).

Pursuing this idea, Milinski and his colleagues conducted a similar experiment framed in terms of climate change (Milinski *et al* 2006). They tested people’s willingness to contribute to the costs of an advertisement that would provide information about climate change and tell readers how to reduce their emissions—in a set up in which contributions were sometimes made public and could affect behavior in an indirect-reciprocity game. Specifically, 156 students were divided into 26 groups of 6, each of which played a public-

good game alternated with an indirect-reciprocity game. The students were given 12 euros and were able to keep any money that remained in their account when the game ended. The experiment had 20 rounds. In the odd-numbered rounds, the students played the indirect-reciprocity game. In the even-numbered rounds, they played the public-good game. Half the groups were given expert information on climate change.

In the indirect-reciprocity game, each player had one chance to make a gift to one of the other five players, and one chance to receive a gift. The gift cost the giver 1.5 euros, but the experimenters matched it so the recipient received 3 euros. Players' gifts were public in the indirect-reciprocity game and in some rounds of public-good games. Players were told that there were no opportunities for direct reciprocity: if A had a chance to make a gift to B, B would not have a chance to return the gift to A. In the public-good game, each of the six students could contribute 0, 1, or 2 euros, the contributions again being matched by the experimenters. In contrast to a typical public-good game, the contributions were not divided among the students, but were instead used to pay for the advertisement. In every other round of the public-good game, contributions were disclosed to the other players. In the other rounds, the contributions were anonymous.

Many students contributed to the climate fund, even when their contributions were anonymous, and the groups that were given more information about climate change gave more than the others did. But the average contribution was nearly twice as large when contributions were public. Students who contributed to the advertising fund in the public public-good games were also more likely to receive a gift in the subsequent indirect-reciprocity game.

Punishment. Cooperation is also more likely when cooperators can punish defectors other than by withdrawing their own cooperation. Indeed, contributions in public-good games may decline partly because cooperators want to punish defectors and can do so only by not contributing. To test this explanation, Ernst Fehr and Simon Gächter included a "punishment condition" in some public-good experiments (Fehr and Gächter 2000). Specifically, subjects could pay up to ten monetary units to punish another player, each unit costing the other player three units. When punishment was possible, cooperators often punished defectors, and the level of cooperation rose over the rounds of the game instead of falling. By game six, the average contribution was close to the maximum. Another experiment found that, when people can choose between playing a public-good game in a group in which fines are possible and one in which they are not, the group with fines does better and almost everyone eventually migrates to this group (Güerker, Irlenbusch, and Rockenbach 2006). Results such as these have led some researchers to conclude that many people are conditional cooperators and altruistic punishers: they contribute if others contribute and they punish defectors at personal cost (Gintis *et al* 2005). Put differently, they comply with a norm that prescribes cooperation with cooperators and the punishment of defectors.

It is easy to see why people are more likely to cooperate in the presence of altruistic punishers. Yet cooperation remains puzzling, because altruistic punishment is itself socially beneficial but individually costly. Those who cooperate but don't punish receive the same benefits as those who both cooperate and punish, but they don't pay any of the costs of punishing. It would seem that "winners don't punish" (Dreber *et al* 2008). So how could altruistic punishment evolve in competition with a merely cooperative strategy?

One possibility is that altruistic punishment—along with cooperation in a one-shot prisoners' dilemma and rejection of positive offers in the ultimatum game—is just a byproduct of a self-interested and evolutionarily successful tendency to cooperate conditionally in repeated games. Robert Boyd and colleagues give another possible explanation (Boyd *et al* 2003; see also Sethi and Somanathan 1996). They begin by noting that, in a simple public-good game without punishment, the cost disadvantage of cooperating does not vary with the prevalence of defectors: it is the same whether everyone else is a defector or whether no one else is. This is not true of the cost disadvantage of being an altruistic punisher in a public-good game with punishment. If everyone else in the group is a defector, altruistic punishers have a big cost disadvantage relative to mere cooperators. But, if defectors are rare, altruistic punishers have only a slight cost disadvantage. Moreover, if altruistic punishers are common, defectors will be rare. In the language of evolution, within-group selection pressures against altruistic punishers are weak when altruistic punishers are common.

Boyd and his colleagues then argue that this small within-group disadvantage could be outweighed by between-group selection in favor of groups containing many altruistic punishers. They argue that cultural group selection is more likely than genetic group selection to have accounted for the rise of altruistic punishment. Bowles has separately argued, however, that our ancestors lived in conditions favorable to genetic group selection of such behavior (Bowles 2006).

Either way, cooperation is one side of a coin whose other side is hostility to outsiders. In one public-good experiment, subjects were randomly divided into four groups just before playing the game (Dawes and Thaler 1980). Half were told that the money they contributed would go to their own group, the other half that the money would go to another group. Within these groups, half could communicate with each other and half could not. The groups that couldn't communicate contributed on average about 30 percent of their endowment. Allowing communication did not increase contributions when the contributions benefited another group, but raised them to 70 percent when the contribution went to their own (randomly formed) group.

Although punishment can sustain cooperation, it doesn't always do so. Herrmann, Thöni, and Gächter (2008) conducted an experiment in which university undergraduates in each of sixteen cities played a finitely repeated public-good game, with and without punishment. In some cities, the results were like those discussed above. But in others the punishment of *cooperators* was common, and punishment failed to promote cooperation.

Herrmann and his colleagues find that the punishment of cooperators is more common in countries with weak norms of civic cooperation and a weak rule of law. They conclude: “punishment opportunities are socially beneficial only if complemented by strong social norms of cooperation” (p. 1362).

Cooperative norms combined with taxes or quotas

The research reviewed here shows that people are better at cooperating in social dilemmas than might be expected. Some mixture of enlightened self-interest and cooperative norms often allows them to avoid the worst, even when governments haven't imposed taxes or regulations that align individual interests with social interests. Because it is difficult for national governments to address the problem of climate change, the partial efficacy of voluntary individual action might give rise to qualified optimism: perhaps a combination of imperfect voluntary cooperation and imperfect taxes and quotas will together get us close to solving the problem.

But other research shows that the two things need not reinforce each other (Bowles 2008; Fehr and Rockenbach 2008; Frey 1997; Gintis *et al* 2005). In particular, material incentives designed to encourage cooperation may reduce the effectiveness of cooperative norms. A striking illustration in a different domain comes from a study of a group of day-care centers in Israel that imposed a fine on parents who were late collecting their children and found that parents became even less punctual (Gneezy and Rustichini 2000). Parents may have thought of the fine as a price and may have come to view being late as a service that they could purchase, not as a violation of norm against inconveniencing others.

Something similar happened in the common-pool-resource experiment in rural Colombia, the first stage of which was described earlier. The main purpose of the experiment was to test whether an imperfectly enforced regulation improved cooperation. So in a second stage half the groups were told that they should take only the socially optimal amount of the resource and that they would be fined if they were audited and found to have taken more. But the subjects were given information that implied that the probability of being audited was low and that the expected-value-maximizing choice was still to take too much. In the first rounds of the second stage, regulation worked well, but after a few rounds people were taking at least as much as before. Cardenas and his colleagues hypothesized that regulation “crowded out group-regarding behavior in favor of greater self-interest” (p. 1731). The other half of the groups continued to play the game without regulation but now had the ability to communicate with each other before making their choice. They cooperated somewhat more than before, and in the contrast to the case of regulation, this improvement endured.

Thus a government intervention that improves material incentives but doesn't fully align them with the social interest may be worse than no intervention. As John Gowdy writes, “Monetary incentives may actually discourage the kinds of behaviors needed to solve

collective social problems like global climate change.”¹⁴ If an imperfect intervention is worse than nothing, one option is of course to do nothing. Another is to change the intervention so that the material incentives it creates are strong enough to compensate for any reduction in people’s propensity to cooperate voluntarily. A third is to try to design and introduce the intervention in a way that doesn’t suppress voluntary cooperation and perhaps even promotes it. How exactly to do this is unclear, but Bruno Frey argues that “External interventions crowd-out intrinsic motivation if they are perceived to be controlling and they crowd-in intrinsic motivation if they are perceived to be supportive” and, specifically, “laws that prevent free riding by others, and establish fairness and equity serve to maintain or even crowd-in civic virtue” (Frey 1997). Fehr and Bettina Rockenbach find that “Sanctions revealing selfish or greedy intentions destroy altruistic cooperation almost completely, whereas sanctions perceived as fair leave altruism intact” (Fehr and Rockenbach 2003).

Conclusions

In summary, research shows that people often cooperate in social dilemmas despite strong temptations to defect. Sometimes, cooperation results from direct or indirect reciprocity that is consistent with enlightened self-interest. At other times, it seems to depend on a preference for cooperating with other cooperators and for punishing defectors; or to say much the same thing in a different way, on a norm that prescribes cooperation with cooperators and the punishment of defectors. In light of this research, voluntary efforts to reduce greenhouse-gas emissions are not surprising. Nor is pressure on governments to reduce emissions, even when that is not in the national interest.

Climate change, however, presents an especially difficult test for people’s propensity to cooperate. The number of people involved in the dilemma is much greater than in experimental public-good games or in the real-world cases, such as that of the Maine lobstermen, in which people succeed in cooperating. Moreover, the players, and the norms they are influenced by, are much more diverse. Because the problem is truly global, cooperation isn’t supported by correlated encounters or by desires to promote the interest of one’s own group at the expense of other groups. And there may be less time for cooperative institutions to evolve than was available to Maine lobsterman and other groups who succeeded in solving social dilemmas. Humans may be exceptionally cooperative and yet still not cooperate much in slowing climate change.

The negative part of this conclusion lends weight to proposals to emphasize adaptation in strategies for addressing climate change and to try to mitigate climate change by

¹⁴ Gowdy 2007 uses climate change to illustrate the difference between an approach to policy informed by “conventional” economics and an approach informed by behavioral economics and related disciplines. It covers some of the same territory as this paper, but also includes many references to neuroscience and the study of nonhuman animals.

transforming the game that is being played into one in which cooperation faces less-daunting hurdles, for example by increasing funding for research on renewable energy. The positive part of the conclusion suggests that we should not give up on cooperation even in the basic emissions game. Some cooperation can be expected, and the amount of cooperation may be influenced by policymakers.

Sanctions for defectors. The potential importance of sanctions in enforcing international cooperation in tackling climate change has often been recognized. Joseph Stiglitz, for example, has argued that greenhouse-gas-intensive goods exported from developed countries without appropriate domestic taxes or quotas should be subjected to import taxes or quotas by other countries, partly to encourage those governments to impose domestic taxes or quotas (Stiglitz 2007). The research reviewed here suggests that sanctions could increase cooperation in a different way, by allowing cooperators to feel that defectors were being punished. Punishment need not take the form of international trade sanctions, of course. Individuals might, for example, boycott certain domestic and imported goods that were not subject to appropriate taxes or quotas. As we have seen, however, punishment doesn't always promote cooperation. Its effect would depend on how the punished reacted. Fear that trade sanctions might prompt retaliation rather than greater cooperation also finds some support in the research. But the importance of punishment in sustaining cooperation in other social dilemmas suggests that more attention should be paid to designing ways in which cooperators can sanction defectors without withdrawing their own cooperation.

Promoting cooperative norms. A second option is to use public communication to invoke and promote cooperative norms. While policy analysts are used to thinking that changing people's behavior requires changing their incentives or giving them new information, the research reviewed here suggests that moralizing may also be effective. It may help simply to talk about the social benefits of reducing greenhouse-gas emissions. Gowdy 2007 suggests (somehow) "giving people a shared responsibility and appealing directly to a sense of the common good." It may also help to stress examples of cooperation to reduce emissions and evidence that some cooperation can be expected. Likewise, it may be counterproductive to stress evidence that defection is likely (which helps explain why emphasizing adaptation can be unpopular).

The development of greenhouse-gas taxes or quotas should also be sensitive to the risk that the incentives they create may weaken cooperative norms. It may sometimes be possible to enact roughly optimal taxes or quotas that can be effectively enforced, perhaps because they directly affect only a small number of upstream firms, such as producers or importers of fossil fuels. If so, any weakening of cooperative norms doesn't matter. But if the tax or quota is imposed on many businesses—including farmers and forest owners, for example—and cannot easily be enforced, how the tax or quota is perceived matters. Ideally, it would be perceived as formalizing and strengthening an existing social norm, like a law against theft or against driving on the wrong side of the road, not as creating an entirely new obligation. That suggests that efforts to build public support for the taxes

and quotas among affected parties may have a benefit that goes beyond improving the prospects of passing the necessary legislation.

More information on emissions. A third option is to improve the availability of information on greenhouse-gas emissions. Most mechanisms that support cooperation in other social dilemmas depend on the availability of information about others' players behavior: consider strategies that specify cooperating if and only if (sufficiently many) others cooperate; strategies that require separately punishing defectors; strategies that specify cooperating with people that have previously cooperated with others; and strategies that use cooperation in a public-good game to build a reputation valuable in other games. Public information on greenhouse-gas emissions is notably limited. The World Resources Institute's website, which seems to be the best source of comprehensive national information, contains no information on emissions after 2004, and no comprehensive information is available after 2000.¹⁵ In that year, emissions from sources not reported in later years made up 41 percent of the global total. For Indonesia and Brazil, which had the third and fourth highest levels of estimated emissions in 2000, they made up 90 and 85 percent, respectively, of the national totals.

By contrast, information on inflation, gross-domestic product, and other economic phenomena is often available on a quarterly basis, with a delay of only a few months. Transparency International promptly publishes annual indicators of corruption. The World Bank promptly publishes annual indicators of the ease of doing business. All these indicators attract attention, create controversy, and influence policy. Climate change needs the equivalent. Moreover, the best information would not only be national. It would also cover finer groupings, such as provinces, cities, and perhaps large firms, as well as facilitating estimates of emissions for representative individuals.

¹⁵ See <http://cait.wri.org> (accessed August 2008).

References

- Acheson, James M. 1987. The lobster fiefs revisited: Economic and ecological effects of territoriality in Maine lobster fishing. In Bonnie J. McKay and James M. Acheson (eds), *The question of the commons: The culture and ecology of communal resources*. Tucson: Arizona University Press
- Andreoni, James. 1995a. Cooperation in public goods experiments: Kindness or confusion. *American Economic Review* 85(4):891–904.
- Andreoni, James. 1995b. Warm glow versus cold prickle: The effects of positive and negative framing on cooperation in experiments. *Quarterly Journal of Economics* 110(1):1–21.
- Aumann, Robert. 1987. Game theory. In John Eatwell, Murray Milgate, and Peter Newman (eds), *The new Palgrave: Game theory*. London: Macmillan.
- Axelrod, Robert. 1984. *The evolution of cooperation*. Basic Books.
- Axelrod, Robert, and William D. Hamilton. 1981. The evolution of cooperation. *Science* 211:1390–1396.
- Barrett, Scott. 2003. *Environment and statecraft: The strategy of environmental treaty making*. Oxford: Oxford University Press.
- Bergstrom, Theodore. 2006. Naturally Kantian or zoo economicus? Evolutionary theories of selfishness and altruism among men and beasts. In Serge-Christophe Kolm and Jean Mercier Ythier (eds), *The handbook of the economics of giving, altruism, and reciprocity, Volume I: Foundations*. North Holland.
- Binmore, Ken. 1994. *Game theory and the social contract, Volume 1: Playing fair*. Cambridge, MA: MIT Press.
- Binmore, Ken. 1998. *Game theory and the social contract, Volume 2: Just playing*. Cambridge, MA: MIT Press.
- Binmore, Ken. 2007. *Playing for real: A text on game theory*. Oxford: Oxford University Press.
- Bowles, Samuel. 2006. Group competition, reproductive leveling, and the evolution of human altruism. *Science* 314:1569–1572.
- Bowles, Samuel. 2008. Policies designed for self-interested citizens may undermine the “moral sentiments”: Evidence from economic experiments. *Science* 320:1605–1609.
- Boyd, Robert, and Jeffrey P. Lorberbaum. 1987. No pure strategy is evolutionarily stable in the repeated prisoners’ dilemma game. *Nature* 327:58–59.

- Boyd, Robert, and Peter J. Richerson. 1988. The evolution of reciprocity in sizeable groups. *Journal of Theoretical Biology* 132:337–356.
- Boyd, Robert, Herbert Gintis, Samuel Bowles, and Peter J. Richerson. 2003. The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences* 100:3531–35.
- Camerer, Colin F. 2003. *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
- Cardenas, Juan Camilo, John Stranlund, and Cleve Willis. 2000. Local environmental control and institutional crowding-out. *World Development* 28:1719–1733.
- Chafe, Zöe, and Hilary French. 2008. Improving carbon markets. In *State of the World 2008: Toward a sustainable global economy*. WorldWatch Institute.
- Colman, Andrew M. 1999. *Game theory and its applications in the social and biological sciences*. Second edition. London and New York: Routledge.
- Croson, Rachel T. A. 2008. Public goods experiments. In Steven N. Durlauf and Lawrence E. Blume (eds), *The new Palgrave dictionary of economics*. Second edition. Palgrave Macmillan, 2008.
- Dawes, Robin. 1980. Social dilemmas. *Annual Review of Psychology* 31:169–193.
- Dawes, Robyn M., and Richard H. Thaler. Anomalies: Cooperation. *Journal of Economic Perspectives* 2(3):187–197.
- Doebeli, Michael, Christoph Hauert, and Timothy Killingback. 2004. The evolutionary origin of cooperators and defectors. *Science* 306:859–862.
- Dreber, Anna, David G. Rand, Drew Fudenberg, Martin A. Nowak. 2008. Winners don't punish. *Nature* 452:348–351.
- Dutta, Prajit K., and Roy Radner. 2004. Self-enforcing climate-change treaties. *Proceedings of the National Academy of Sciences* 101(14):5174–5179.
- Fehr, Ernst, and Bettina Rockenbach. 2003. Detrimental effects of sanctions on human altruism. *Nature* 422:137–140.
- Fehr, Ernst, and Simon Gächter. 2002. Altruistic punishment in humans. *Nature* 415:137–140.
- Frank, Robert H., Thomas Gilovich, and Dennis T. Regan. 1993. Does studying economics inhibit cooperation? *Journal of Economic Perspectives* 7(2):159–171.
- Frean, Marcus R. 1994. The prisoners' dilemma without synchrony. *Proceeding of the Royal Society London Series B* 257:75–79.

- Frey, Bruno. 1997. A constitution for knaves crowds out civic virtues. *Economic Journal* 107(4):1043–1053.
- Fudenberg, Drew, and Eric Maskin. 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*: 54(3):553–554.
- Gintis, Herbert, Samuel Bowles, Robert Boyd, and Ernst Fehr. 2005. *Moral sentiments and material interests: The foundations of cooperation in economic life*. Cambridge, MA: MIT Press.
- Gneezy, Uri, and Aldo Rustichini. 2000. A fine is a price. *Journal of Legal Studies* 29(1):1–17.
- Gowdy, John M. 2007. Behavioral economics and climate change policy. Rensselaer Working Papers in Economics 0701.
- Gupta, S., et al. 2007. Policies, instruments and co-operative arrangements. In *Climate Change 2007: Mitigation. Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge: Cambridge University Press.
- Gürerk, Özgür, Bernd Irlenbusch, and Bettina Rockenbach. 2006. The competitive advantage of sanctioning institutions. *Science* 312:108–111.
- Güth, Werner, Rolf Schmittberger, Bernd Schwarze. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3:367–388.
- Harbaugh, William T. 1998. The prestige motive for making charitable transfers. *American Economic Review* (Papers and proceedings) 88:277–282.
- Hardin, Garrett. 1968. The tragedy of the commons. *Science* 162:1243–1248.
- Henrich, Joseph, Richard McElreath, Abigail Barr, Jean Ensminger, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwina Gwako, Natalie Henrich, Carolyn Lesorogol, Frank Marlowe, David Tracer, and John Ziker. 2006. Costly punishment across human societies. *Science* 312:1767–1770.
- Herrmann, Benedikt, Christian Thöni, and Simon Gächter. 2008. Antisocial punishment across societies. *Science* 319:1362–1367.
- Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson. 1982. Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory* 27:245–252.
- Ledyard, John O. 1995. Public goods: A survey of experimental research. Chapter 2 of John H. Kagel and Alvin E. Roth (eds), *Handbook of experimental economics*, Princeton, NJ: Princeton University Press.

- Liebreich, Michael. 2007. How to save the planet: Be nice, retaliatory, forgiving, and clear. Available at <http://www.newenergyfinance.com>.
- Maynard Smith, John. 1982. *Evolution and the theory of games*. Cambridge: Cambridge University Press.
- Milinski, Manfred, Dirk Semmann, and Hans-Jürgen Krambeck. 2002. Reputation helps solve the “tragedy of the commons”. *Nature* 415:424–426.
- Milinski, Manfred, Dirk Semmann, Hans-Jürgen Krambeck, and Jochem Marotzke. 2006. Stabilizing the Earth’s climate is not a losing game: Supporting evidence from public goods experiments. *Proceedings of the National Academy of Sciences* 103:3994–3998.
- Milinski, Manfred, Ralf D. Sommerfeld, Hans-Jürgen Krambeck, Floyd A. Reed, and Jochem Marotzke. 2008. The collective-risk social dilemma and the prevention of simulated dangerous climate change. *Proceedings of the National Academy of Sciences* 105(7):2291–2294.
- Nowak, Martin A. 2006. *Evolutionary dynamics: Exploring the equations of life*. Cambridge, MA: Belknap Press.
- Nowak, Martin A., and Karl Sigmund. 1992. Tit for tat in heterogeneous populations. *Nature* 355:250–253.
- Nowak, Martin A., and Karl Sigmund. 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoners’ dilemma game. *Nature* 364:56–58.
- Nowak, Martin A., and Karl Sigmund. 2005. Evolution of indirect reciprocity. *Nature* 437:1291–1298.
- Olson, Mancur. 1965/1971. *The logic of collective action: Public goods and the theory of groups*. Cambridge: Harvard University Press
- Ostrom, Elinor, Roy Gardner, and James Walker. 1994. *Rules, games, and common-pool resources*. Ann Arbor: University of Michigan Press.
- Panchanathan, Karthik, and Robert Boyd. 2004. Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature* 432:499–502.
- Seabright, Paul. 1993. Managing local commons: Theoretical issues in incentive design. *Journal of Economic Perspectives* 7:113–134.
- Sethi, Rajiv, and E. Somanathan. 1996. The evolution of social norms in common property resource use. *American Economic Review*: 86(4):766–788.
- Skyrms, Brian. 1996. *The evolution of the social contract*. Cambridge: Cambridge University Press.

- Skyrms, Brian. 2004. *The stag hunt and the evolution of social structure*. Cambridge: Cambridge University Press.
- Sugden, Robert. 1986/2004. *The economics of cooperation, rights, and welfare*. Second edition. Palgrave.
- Stiglitz, Joseph. 2007. A new agenda for global warming. *Economists' Voice* July.
- Toth, Ferenc L, *et al.* 2001. Decision-making frameworks. In *Climate Change 2001: Mitigation: Contribution of Working Group III to the Third Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge: Cambridge University Press.
- Tversky, Amos, and Daniel Kahneman. 1981. The framing of decisions and the psychology of choice. *Science* 211:453–458.
- Wedekind, Claus, and Manfred Milinski. 1996. Human cooperation in the simultaneous and the alternating prisoners' dilemma: Pavlov versus generous tit-for-tat. *Proceedings of the National Academy of Sciences* 93:2686–2689.
- World Bank. 2005. *World Development Report 2006: Equity and development*. New York: Oxford University Press.
- Young, H. Peyton. 1996. The economics of conventions. *Journal of Economic Perspectives* 10(2):105–122.