



Anomalies: Cooperation

Robyn M. Dawes, Richard H. Thaler

The Journal of Economic Perspectives, Volume 2, Issue 3 (Summer, 1988), 187-197.

Your use of the JSTOR database indicates your acceptance of JSTOR's Terms and Conditions of Use. A copy of JSTOR's Terms and Conditions of Use is available at <http://www.jstor.org/about/terms.html>, by contacting JSTOR at jstor-info@umich.edu, or by calling JSTOR at (888)388-3574, (734)998-9101 or (FAX) (734)998-9113. No part of a JSTOR transmission may be copied, downloaded, stored, further transmitted, transferred, distributed, altered, or otherwise used, in any form or by any means, except: (1) one stored electronic and one paper copy of any article solely for your personal, non-commercial use, or (2) with prior written permission of JSTOR and the publisher of the article or other text.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The Journal of Economic Perspectives is published by American Economic Association. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/aea.html>.

The Journal of Economic Perspectives
©1988 American Economic Association

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2001 JSTOR

Anomalies Cooperation

Robyn M. Dawes and Richard H. Thaler

Economics can be distinguished from other social sciences by the belief that most (all?) behavior can be explained by assuming that agents have stable, well-defined preferences and make rational choices consistent with those preferences in markets that (eventually) clear. An empirical result qualifies as an anomaly if it is difficult to “rationalize,” or if implausible assumptions are necessary to explain it within the paradigm. This column will present a series of such anomalies. Readers are invited to suggest topics for future columns by sending a note with some references to (or better yet copies of) the relevant research. Comments on anomalies printed here are also welcome. The address is: Richard Thaler, c/o *Journal of Economic Perspectives*, Johnson Graduate School of Management, Malott Hall, Cornell University, Ithaca, NY 14853.

Introduction

Much economic analysis—and virtually all game theory—starts with the assumption that people are both rational and selfish. For example, predictions that players will defect in the prisoner’s dilemma game and free ride in public goods environments are based on both assumptions. People are assumed to be clever enough to figure out that defection or free riding is the dominant strategy, and are assumed to care nothing for outcomes to other players; moreover, people are assumed to have no qualms about their failure to do “the right thing.”¹

¹For a modern treatment of the theory of public goods, see Bergstrom, Blume, and Varian (1986).

■ *Robyn M. Dawes is Professor of Psychology and Head, Department of Social and Decision Sciences, Carnegie-Mellon University, Pittsburgh, PA. Richard H. Thaler is Henrietta Johnson Louis Professor of Economics at the Johnson Graduate School of Management, Cornell University, Ithaca, New York.*

The predictions derived from this assumption of rational selfishness are, however, violated in many familiar contexts. Public television successfully raises enough money from viewers to continue to broadcast. The United Way and other charities receive contributions from many if not most citizens. Even when dining at a restaurant away from home in a place never likely to be visited again, most patrons tip the server. And people vote in presidential elections where the chance that a single vote will alter the outcome is vanishingly small. As summarized by Jack Hirshleifer (1985, p. 55), "... the analytically uncomfortable (though humanly gratifying) fact remains: from the most primitive to the most advanced societies, a higher degree of cooperation takes place than can be explained as a merely pragmatic strategy for egoistic man." But why?

In this column and the next one, the evidence from laboratory experiments is examined to see what has been learned about when and why humans cooperate. This column considers the particularly important case of cooperation vs. free riding in the context of public good provision. (The next column is about the "ultimatum game.")

Single Trial Public Goods Experiments

To investigate why people cooperate, it is necessary to examine behavior in both single play and multiple play environments. Does cooperation evolve, for instance, only as individuals repeatedly interacting with each other find it in their interests to cooperate? A typical public goods experiment uses the following procedures. A group of subjects (most often students but sometimes other adult members of the community) is brought to the laboratory. Groups vary in size, but experiments usually have between 4 and 10 subjects. Each subject is given a sum of money, for example, \$5. The money can either be kept and taken home, or some or all of the money can be invested in a public good, often called a "group exchange." Money invested in the group exchange for the n participants is multiplied by some factor k , where k is greater than 1.0 but less than n . The money invested, with its returns, is distributed equally among all group members. Thus, while the entire group's monetary resources are increased by each contribution (because $k > 1$), each individual's share of one such contribution is less than the amount she or he invests (because $k < n$). Suppose $k = 2$ and $n = 4$. Then if everyone contributes all \$5 to the public good, each ends up with \$10. This is the unique Pareto efficient allocation: no other solution can make everyone better off. But the dominant Nash strategy is to contribute nothing, because in exchange for a player's \$5 contribution, that player receives only \$2.50, while the rest of the payoff (\$7.50) goes to the other players. The rational selfish strategy is to contribute nothing and hope that the other players decide to invest their money in the group exchange. If one player contributes nothing while all the others contribute \$5, then that player will end up with \$12.50, while the other players end up with \$7.50. These conditions constitute a true social dilemma played with real money.

What does economic theory predict will happen in such a game? One prediction, called the *strong free rider hypothesis*, is that everyone will choose the dominant strategy,

that is, nothing will be contributed to the public good. This is surely the outcome predicted by the selfish rational model. A less extreme prediction, called the *weak free rider hypothesis*, is that some people will free ride while others will not, yielding a suboptimal level of the public good, though not necessarily zero. The weak free rider hypothesis obviously does not yield very precise predictions.

The results of single play (“one shot”) public goods experiments lend little support to the strong free rider hypothesis. While not everyone contributes, there is a substantial number of contributors, and the public good is typically provided at 40–60 percent of the optimal quantity. That is, on average, the subjects contribute 40–60 percent of their stake to the public good. In a study by Marwell and Ames (1981), these results held in many conditions: for subjects playing the game for the first time, or after a previous experience; for subjects who believed they were playing in groups of 4 or 80; and for subjects playing for a range of monetary stakes.² Indeed, Marwell and Ames found only one notable exception to this 40–60 percent contribution rate. When the subjects were a group of University of Wisconsin economics graduate students, the contribution rate fell to 20 percent, leading them to title their article “Economists Free Ride: Does Anyone Else?”³ (Interestingly, economists told about the experiments predicted on the average a rate of about 20 percent—but for all participants, not just their students.)

Multiple Trial Experiments

A natural question to ask about the surprisingly high level of cooperation observed by Marwell and Ames is what would happen if the same players repeated the game several times. This question has been investigated by Kim and Walker (1984), Isaac, Walker, and Thomas (1984), and Isaac, McCue, and Plott (1985). The experimental design in these papers is similar to Marwell and Ames, except that there are usually 10 repetitions of the game. Two major conclusions emerge from these papers. First, on the initial trial, cooperation is observed at rates similar to those obtained by Marwell and Ames. For example, across 9 different experiments with varying designs, Isaac, McCue and Plott obtained a 53 percent contribution rate to the public good. Second, within a few repetitions, cooperation declines sharply. After 5 trials, the contributions to the public good were only 16 percent of the optimum. The experiments by Isaac, Walker and Thomas also obtained a decline in the contribution rate over time, though the decline was not as abrupt.⁴

²In the experiments with the highest stakes, contribution rates were somewhat lower, in the 28–35 percent range.

³This result has never been replicated, and so should be treated as preliminary. We wonder, however, whether economists *are* different. Do economists as a group donate less to charity than other similar groups? Are they less likely to leave tips in out-of-town restaurants?

⁴For experiments with a high return to contributing to the public good, the initial contribution rate was 52 percent, which fell to 32 percent on trial 10. In versions with low returns to contributing, the initial rate was 40 percent and the final rate was 8 percent.

Why does the contribution rate decline with repetition? One reasonable conjecture is that subjects learn something during the experiment that induces them to adopt the dominant strategy of free riding. Perhaps the subjects did not understand the game in the first trial and only learned that free riding was dominant over time. This possibility, however, appears unlikely in light of other experimental evidence. For example, the usual cooperation rates of roughly 50 percent are observed on trial one even for *experienced* subjects, that is, subjects who have participated in other multiple trial public goods experiments (e.g., Isaac and Walker, forthcoming). Also, Andreoni (1987a) has investigated the learning hypothesis directly, using the simple procedure of restarting the experiment. Subjects were told they would play a ten-period public goods game. When the ten periods were completed, the subjects were told they would play again for another ten rounds with the same other players. In the first ten trials Andreoni replicated the decaying contribution rate found by previous investigators, but upon restarting the game contributions went back up to virtually the same contribution rates observed on the initial trial in the first game (44 percent on trial one of the second game vs. 48 percent in the first). Such results seem to rule out any explanation of cooperation based on subjects' misunderstanding the task.⁵

Reciprocal Altruism

One currently popular explanation of why we observe so much cooperation in and outside of the laboratory invokes reciprocal altruism as the mechanism. This explanation, most explicitly developed by Axelrod (1984), is based on the observation that people tend to reciprocate—kindness with kindness, cooperation with cooperation, hostility with hostility and defection with defection. Thus, being a free rider may actually be a less fruitful strategy when the chooser takes account of the probable future response of others to his or her cooperation or defection. A cooperative act itself—or a reputation for being a cooperative person—may with high probability be reciprocated with cooperation, to the ultimate benefit of the cooperator.

The most systematic strategy based on the principle of reciprocal altruism is a TIT-FOR-TAT one first suggested by Anatol Rapoport, in which a player begins by cooperating and then chooses on trial t the same response the other player has made on trial $t - 1$. The real strength of this explanation lies in demonstrating, both analytically and by computer tournaments of interacting players (programs) in iterated social dilemmas, that any person or small group of people practicing such reciprocal altruism will have a statistical tendency to receive higher payoffs “in the long run” than those who don't practice it. In fact, TIT-FOR-TAT “won” two computer tournaments Axelrod conducted in which game theorists proposed various strategies that were compared against each other in pairwise encounters with repeated plays. Because evolution is concerned with such long run probabilistic phenomena, it can be inferred that reciprocating people have greater “inclusive fitness” than do

⁵A similar conclusion is reached by Goetze and Orbell (forthcoming).

non-reciprocating ones. Hence, to the degree to which such a tendency has some genetic basis, it should evolve as an adaptation to the social world.

An implication of reciprocal altruism is that individuals will be uncooperative in dilemma situations when there is no possibility of future reciprocity from others, as in situations of anonymity or interacting with people on a “one-shot” basis. Yet we observe 50 percent cooperation rates even in single trial experiments, so reciprocal altruism cannot be used directly to explain the experimental results described so far. Also, of course, it is very difficult to play TIT-FOR-TAT, or any other strategy based on reciprocal altruism, when more than 2 people are involved in the repeated dilemma situation. If some members of a group cooperate on trial t while others defect, what should a player attempting to implement a TIT-FOR-TAT type strategy choose on the subsequent trial?

A related hypothesis that appears consistent with the decaying contribution rates observed in the multiple trial experiments is suggested by the theoretical work of Kreps, Milgrom, Roberts, and Wilson (1982). They investigate the optimal strategy in a repeated prisoner’s dilemma game with a finite number of trials. If both players are rational, then the dominant strategy for both is to defect on every trial. While TIT-FOR-TAT has been shown to be effective in infinitely repeated prisoner’s dilemma games (or equivalently, games with a constant small probability of ending after any given trial), games with a known end point are different. In any finite game both players know that they should defect on the last trial, so there is no point in cooperating on the penultimate trial, and by backward induction, it is never in one’s best interest to cooperate. What Kreps et al. show is that if you are playing against an opponent whom you think may be irrational (i.e., might play TIT-FOR-TAT even in a game with finite trials), then it may be rational to cooperate early in the game (to induce your irrational opponent to cooperate too). Since the public goods games have a similar structure, it could be argued that players are behaving rationally in the Kreps et al. sense. Once again, however, the data rule out this explanation. Cooperation never falls to zero, even in one-trial games or in the last period of multi-trial games when it can never be selfishly rational to cooperate.

Additional evidence against the reciprocity hypothesis comes from another experiment designed by Andreoni. One group of 15 subjects played repeated trials in groups of 5 as described above. Another group of 20 subjects played the same game in groups of 5, but the composition of the group varied on each trial. Moreover, the subjects did not know which 4 of the other 19 subjects would constitute their group in any given round of the game. In this condition, there can be no strategic advantage to cooperation, since the players in the next round will be, in essence, strangers. If cooperation in early rounds of these experiments is observed, strategic cooperation can be ruled out. Indeed, Andreoni found that cooperation was actually a bit higher in the stranger condition than in a comparable condition where the groups remained intact. (This effect was statistically significant, though slight.)

One conclusion which emerges from these experiments is that people have a tendency to cooperate until experience shows that those with whom they are interacting are taking advantage of them. This “norm of cooperation” will resemble

reciprocal altruism in infinitely repeated games; but the behavior, as we have seen, is also observed in cases when reciprocal altruism would be inappropriate. One explanation for this type of behavior is offered by Robert Frank (1987). Frank argues that people who adopt a norm of cooperation will do well by eliciting cooperation from others, and attracting interaction with other cooperators. The key to Frank's argument is that one cannot successfully fake being cooperative for an extended period of time—just as one cannot be successful getting people to believe too many lies.⁶ Furthermore, because cooperators are, by assumption, able to identify one another, they are able to interact selectively and exclude defectors.

Altruism

There are other explanations of why people cooperate both in the lab and the field. One is that people are motivated by “taking pleasure in others' pleasure.” Termed *pure altruism*,⁷ this motive has been eloquently stated by Adam Smith, in *The Theory of Moral Sentiments* (1759; 1976): “how selfish soever man may be supposed to be, there are evidently some principles in his nature, which interest him in the fate of others, and render their happiness necessary to him, though he derive nothing from it, except the pleasure of seeing it.” While the pleasure involved in seeing it may be considered “selfish” (following the sophomoric argument that altruism is by definition impossible, because people do what they “want” to do), the passage captures the idea that people are motivated by positive payoffs for others as well as for themselves. Consequently, they may be motivated to produce such results through a cooperative act. One problem with postulating such pure altruism as a reason for contributing to public goods is that such contributions cannot be explained purely in terms of their effects. If they could, for example, then governmental contributions to the same goal should “crowd out” private contributions on a dollar-for-dollar basis, since the results are identical no matter where the funding comes from. Such crowding out does not appear to be nearly complete. In fact, econometric studies indicate that an increase in governmental contribution to such activities is associated with a decrease in private contribution of only 5 to 28 percent (Abrams and Schmitz, 1978, 1984; Clotfelter, 1985).

Another type of altruism that has been postulated to explain cooperation is that involved in the act of cooperating itself, as opposed to its results. “Doing the right (good, honorable, . . .) thing” is clearly a motive for many people. Sometimes termed *impure altruism*, it generally is described as satisfaction of conscience, or of noninstrumental ethical mandates.

The roles of pure and impure altruism and other causes for cooperation (or the lack thereof) have been investigated over the last decade by the team of Robyn

⁶As the late Senator Sam Ervin said: “The problem with lying is that you have to have a perfect memory for what you said.” None of us do. It's easier to remember what actually happened, although that is not easy either.

⁷The terms pure and impure altruism are introduced by Andreoni (1987b).

Dawes, John Orbell and Alphons van de Kragt. In one set of experiments (Dawes et al., 1986), they examined the motives for free riding. The game used for these experiments had the following rules. Seven strangers were given \$5 each. If enough people contributed their stake to the public good (either 3 or 5 depending on the experiment), then every person in the group would receive a \$10 bonus whether or not they contributed. Thus, if enough subjects contribute, each contributor would leave with \$10 and each non-contributor would leave with \$15. If too few contributed, then non-contributors would keep their \$5 while contributors would leave with nothing. Subjects were not permitted to talk to one another (though this was modified in subsequent experiments). In this context two reasons for not contributing can be identified. First, subjects may be afraid that they will contribute but not enough others will, so their contribution will be futile. This motive for defecting was termed “fear.” Second, subjects may hope that enough others will contribute and hope to receive \$15 instead of \$10. This motive was called “greed.” The relative importance of fear and greed was examined by manipulating the rules of the game. In the “no greed” condition, payoffs were changed so that all subjects would receive \$10 if the number of contributors was sufficient (rather than \$10 for contributors and \$15 for free riders). In the “no fear” condition contributors were given a “money back guarantee:” if a subject contributed and not enough others did, the subject would receive the money back. (However, in this condition if the public good was provided, contributors would receive only \$10 while free riders would get \$15.) The results suggested that greed was more important than fear in causing free riding. In the standard game contribution rates averaged 51 percent. In the no fear (money back) game contributions rose to 58 percent, but in the no greed game contributions were 87 percent.⁸

Another possible interpretation is that the no greed condition can produce a stable equilibrium, while the no fear cannot. If subjects in the no greed condition believe that the mechanism of truncating payoffs works to motivate others to contribute, their motive will be enhanced as well, because the only negative result of contributing occurs if enough others *don't* contribute. In contrast, subjects in the no fear condition who conclude that the conditions will encourage others to contribute will be tempted to free ride themselves, leading to the conclusion that others will be tempted as well, leading to the conclusion that they should themselves contribute, etc.—an infinite loop.

One of the most powerful methods for inducing cooperation in these games is to permit the subjects to talk to one another. Twelve groups were run with the same payoffs described earlier, but under conditions in which discussions were allowed. The effect of this discussion was remarkable (van de Kragt, et al., 1983). Every group used the discussion period to specify a group of people who were designated to cooperate. The most common means of making the distributional decision was by lottery, though

⁸Notice that contributing could be selfishly rational if a subject thought that the probability his or her contribution would be critical (that is, exactly $M - 1$ others will contribute) was greater than one-half. However, subjects who contributed did not generally believe that their contribution was necessary. Virtually no contributors believed they were critical to obtaining the public good with a probability greater than 0.50. In fact, pooling across all conditions, 67 percent of the contributors believed so many others would contribute that their own contributions would be redundant.

volunteering was also observed. One group attempted interpersonal utility comparisons to determine relative “need.” Whatever methods the groups used, they worked. All 12 groups provided the public good, and in 3 of the groups more than the required number of subjects contributed. These results are consistent with the earlier ones. Subjects designated as contributors cannot greedily expect more from free riding, because their contributions are (believed to be) crucial for their obtaining the bonus (and were in all but 3 groups). Moreover, belief that others in the designated set of contributors will be motivated to contribute by the designated contributor mechanism will enhance—rather than diminish—each designated contributor’s motive to contribute.

One possible explanation for the value of discussion is that it “triggers” ethical concerns that yield a utility for doing the “right” thing (that is, impure altruism). Elster (1986), for example, has argued that group discussions in such situations yield arguments for group-regarding behavior (it is hard to argue for selfishness), and that such arguments have an effect not only on the listener but on the person making them as well. To test this hypothesis, a new set of experiments was conducted (Orbell, van de Kragt and Dawes, forthcoming). In this set of experiments all 7 subjects were given \$6 each. They could either keep the money or contribute it to the public good in which case it would be worth \$12 to the other 6 members of the group. In this case, keeping the \$6 is a dominant strategy because the person who does so receives both that \$6 and \$2 from each of the other 6 group members who gave away the money.

Subjects first met in groups of 14 in a waiting room in which they were not allowed to talk; they were then divided into the two groups on a clearly random basis. Half of these subgroups were allowed to talk about the decision, half not. The experimenters told half of the groups that the \$12 given away would go to the other six people in their own group, while the other half of the groups were told that the money would go to six people in the other group. There are thus four conditions—discussion or no discussion crossed with money goes to own group or other group. If discussion simply makes individuals’ egoistic payoffs clear, then it should not increase cooperation rate in any of these conditions since free riding is dominant. If, however, discussion increases utility for the act of cooperation per se, then discussion should be equally effective whether the money given away goes to members of their own group or to the other group—which consists, after all, of very similar people who were indistinguishable prior to the random drawing (usually college students or poorer members of the community).

The results were clear. In the absence of discussion, only about 30 percent of the subjects gave away the money, and those who did so indicated that their motive was to “do the right thing” irrespective of the financial payoffs.⁹ Discussion raises the

⁹In a similar—but simulated—one-shot experiment Hofstadter (1983) had discovered a roughly identical cooperation rate among his eminent friends. Most defect, but some cooperate, and for reasons of impure altruism. As one cooperator, Professor Daniel C. Dennett of Tufts, put it: “I would rather be the person who bought the Brooklyn bridge than the person who sold it. Similarly, I feel better spending \$3 gained by cooperation than \$10 gained by defection.” (Hofstadter terms that a “wrong reason” for cooperating in a dilemma situation; yet it is the one often given by the subjects who cooperate without discussion in the experiments described above, and similar ones.)

cooperation rate to 70 percent, but only when the subjects believe the money is going to members of their own groups; otherwise, it is usually less than 30 percent. Indeed, in such groups it was common to hear comments that the “best” possible outcome would be for all group members to keep their money while those in the other group gave it away (again, people from whom the subjects have been randomly separated about 10 minutes earlier).

Thus, group identity appears to be a crucial factor in eschewing the dominating strategy. That result is compatible with previous social-psychological research on the “minimal group” paradigm (e.g. Tajfel and Turner, 1979; the papers contained in Turner and Giles, 1981), which has repeatedly demonstrated that allocative decisions can be sharply altered by manipulations substantially weaker than 10 minutes of discussion. For example, a “common fate” group identity—where groups received differing levels of payoffs depending on a coin toss—led subjects to attempt to “compensate” for non-cooperators in their own group by increasing cooperation rates, while simultaneously decreasing cooperation when the non-cooperators were believed to be in the other group, even when the identities of the people involved were unknown (Kramer and Brewer, 1986).

In the groups in which discussion was permitted it was very common for people to make promises to contribute. In a second series of experiments, Orbell, van de Kragt and Dawes investigated whether these promises were important in generating cooperation. Perhaps people feel bound by their promises—or believe they will receive a “satisfactory” payoff if they give away the money when others promise to do so because others will be bound by such promises. The main result was that promise making was related to cooperation only when every member of the group promised to cooperate. In such groups with universal promising, the rate of cooperation was substantially higher than in other groups. In groups in which promising was not universal there was no relationship between each subject’s choice to cooperate or defect and (1) whether or not a subject made a promise to cooperate, or (2) the number of other people who promised to cooperate. Consequently, the number of promises made in the entire group and the group cooperation rate were unrelated. These data are consistent with the importance of group identity if (as seems reasonable) universal promising creates—or reflects—group identity.

Commentary

In the rural areas around Ithaca it is common for farmers to put some fresh produce on a table by the road. There is a cash box on the table, and customers are expected to put money in the box in return for the vegetables they take. The box has just a small slit, so money can only be put in, not taken out. Also, the box is attached to the table, so no one can (easily) make off with the money. We think that the farmers who use this system have just about the right model of human nature. They feel that enough people will volunteer to pay for the fresh corn to make it worthwhile to put it out there. The farmers also know that if it were easy enough to take the money, someone would do so.

In contrast to these farmers, economists either avoid judgments of human nature, or make assumptions that appear excessively harsh. It is certainly true that there is a “free rider problem.” Not all people can be expected to contribute voluntarily to a good cause, and any voluntary system is likely to produce too little of the public good (or too much of the public bad in the case of externalities). On the other hand, the strong free rider prediction is clearly wrong—not everyone free rides all of the time.

There is a big territory between universal free riding and universal contributing at the optimal rate. To understand the problems presented by public goods and other dilemmas it is important to begin to explore some issues that are normally ignored in economics. For example, what factors determine the rate of cooperation? It is encouraging to note that cooperation is positively related to the investment return on the public good. The more the group has to gain through cooperation, the more cooperation is observed—the supply of cooperation is upward sloping. The results involving the role of discussion and the establishment of group identity are, however, more difficult to incorporate into traditional economic analyses. (One economist attempting to do so proposed that group discussion simply confuses subjects to the point that they no longer understand it is in their best interests to be defectors.)

More generally, the role of selfish rationality in economic models needs careful scrutiny. Amartya Sen (1977) has described people who are always selfishly rational as “rational fools,” because mutual choices based only on egoistic payoffs consistently lead to suboptimal outcomes for all involved. Perhaps we need to give more attention to “sensible cooperators.”

■ *We wish to thank James Andreoni, Linnda Caporael, Mark Isaac, and John Orbell for helpful comments on an earlier draft.*

References

- Abrams, Burtran A., and Mark A. Schmitz, "The Crowding Out Effect of Government Transfers on Private Charitable Contributions," *Public Choice*, 1978, 33, 29–39.
- Abrams, Burtran A., and Mark A. Schmitz, "The Crowding Out Effect of Government Transfers on Private Charitable Contributions: Cross Sectional Evidence," *National Tax Journal*, 1984, 37, 563–68.
- Andreoni, James, "Why Free Ride? Strategies and Learning in Public Goods Experiments," unpublished, University of Wisconsin, Department of Economics, 1987a.
- Andreoni, James, "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," unpublished, University of Wisconsin, Department of Economics, 1987b.
- Axelrod, Robert, *The Evolution of Cooperation*, New York: Basic Books, 1984.
- Bergstrom, Theodore, Lawrence E. Blume, and Hal Varian, "On the Private Provision of Public Goods," *Journal of Public Economics*, 1986, 29, 25–49.
- Clotfelter, Charles T., *Federal Tax Policy and Charitable Giving*, Chicago: The University of Chicago Press, 1985.
- Dawes, Robyn M., John M. Orbell, Randy T. Simmons, and Alphons J. C. van de Kragt, "Organizing Groups for Collective Action," *American Political Science Review*, 1986, 80, 1171–1185.
- Elster, Jon, "The Market and the Forum: Three Varieties of Political Theory." In Jon Elster and Aanund Hylland, eds., *Foundations of Social Choice Theory: Studies in Rationality and Social Change*. Cambridge: Cambridge University Press, 1986, 103–132.
- Frank, Robert, "If *Homo Economicus* Could Choose his own Utility Function, Would He Want One with a Conscience?" *American Economic Review*, September 1987, 77, 593–605.
- Goetze, David, and John M. Orbell, "Understanding and Cooperation," *Public Choice*, forthcoming.
- Hirshleifer, Jack, "The Expanding Domain of Economics," *American Economic Review*, December 1985, 75, Number 6, 53–70.
- Hofstadter, Douglas, "Metamagical Themas," *Scientific American*, 1983, 248, 14–28.
- Isaac, R. Mark, Kenneth F. McCue, and Charles Plott, "Public Goods Provision in an Experimental Environment," *Journal of Public Economics*, 1985, 26, 51–74.
- Isaac, R. Mark, James M. Walker, and Susan H. Thomas, "Divergent Evidence on Free Riding: An Experimental Examination of Possible Explanations," *Public Choice*, 1984, 43, 113–149.
- Isaac, R. Mark, and James M. Walker, "Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism," *Quarterly Journal of Economics*, forthcoming.
- Kim, Oliver, and Mark Walker, "The Free Rider Problem: Experimental Evidence," *Public Choice*, 1984, 43, 3–24.
- Kramer, R. M., and Marilyn Brewer, "Social Group Identity and the Emergence of Cooperation in Resource Conservation Dilemmas." In H. Wilke, D. Messick, and C. Rutte, eds., *Psychology of Decision and Conflict. Vol. 3. Experimental Social Dilemmas*, Frankfurt Am Main: Verlag Peter Lang, 1986, pp. 205–230.
- Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson, "Rational Cooperation in Finitely Repeated Prisoners' Dilemmas," *Journal of Economic Theory*, 1982, 27, 245–252.
- Marwell, Gerald and Ruth Ames, "Economists Free Ride, Does Anyone Else?" *Journal of Public Economics*, 1981, 15, 295–310.
- Orbell, John M., Robyn M. Dawes, and Alphons J. C. van de Kragt, "Explaining Discussion Induced Cooperation," *Journal of Personality and Social Psychology*, forthcoming.
- Sen, Amartya K., "Rational Fools: A Critique of the Behavioral Foundations of Economic Theory," *Journal of Philosophy and Public Affairs*, 1977, 6, 317–344.
- Smith, Adam, *The Theory of Moral Sentiments*. Oxford: Clarendon Press, 1976. (Originally published in 1759.)
- Tajfel, Henri and John C. Turner, "An Integrative Theory of Intergroup Conflict." In W. Austin and S. Worchel, eds., *The Social Psychology of Intergroup Relations*, Monterey, CA: Brooks/Cole, 1979, pp. 33–47.
- Turner, John C., and Howard Giles, *Intergroup Behavior*, Chicago: University of Chicago Press, 1981.
- van de Kragt, Alphons J. C., John M. Orbell, and Robyn M. Dawes, "The Minimal Contributing Set as a Solution to Public Goods Problems," *American Political Science Review*, 1983, 77, 112–22.