

Harnessing Reciprocity to Promote Cooperation and the Provisioning of Public Goods

David G. Rand¹, Erez Yoeli^{2,3}, and Moshe Hoffman^{2,3}

Policy Insights from the Behavioral and Brain Sciences
1–7
© Federation of Associations in Behavioral & Brain Sciences 2014
DOI: 10.1177/2372732214548426
bbs.sagepub.com



Abstract

How can we maximize the common good? This is a central organizing question of public policy design, across political parties and ideologies. The answer typically involves the provisioning of public goods such as fresh air, national defense, and knowledge. Public goods are costly to produce but benefit everyone, thus creating a social dilemma: Individual and collective interests are in tension. Although individuals may want a public good to be produced, they typically would prefer not to be the ones who have to pay for it. Understanding how to motivate individuals to pay these costs is therefore of great importance for policy makers. Research provides advice on how to promote this type of “cooperative” behavior. Synthesizing a large body of research demonstrates the power of “reciprocity” for inducing cooperation: When others know that you have helped them, or acted to benefit the greater good, they are often more likely to reciprocate and help you in turn. Several conclusions stem from this line of thinking: People will be more likely to do their part when their actions are observable by others; people will pay more attention to how effective those actions are when efficacy is also observable; people will try to avoid situations where they could help, but often *will* help if asked directly; people are more likely to cooperate if they think others are also cooperating; and people can develop habits of cooperation that shape their default inclinations.

Keywords

cooperation, reciprocity, game theory, behavioral economics, public goods, reputation

Tweet

To harness reciprocity concerns and increase contributions to public goods, improve transparency, and make people’s contributions observable.

Summary

- Public goods benefit everyone but are costly (in terms of money, time, effort) to produce.
- People are typically more likely to help you if you are seen as a cooperative person who does his or her part (i.e., “reciprocity” drives human cooperative behavior).
- Therefore, making people’s contributions to public goods observable will increase contribution levels (and making efficacy observable will make people more likely to contribute to efficient causes).
- Preventing people from avoiding situations where they are asked to contribute will also increase contributions.
- So too will giving the impression that many others are also contributing and creating habits of cooperation.

Introduction

A central goal of public policy is to improve the welfare of the state’s citizens. Although policy makers often disagree about which specific policies best achieve this goal, wide agreement centers on the overall aim of benefiting the greater good. This article provides a suite of tools that can help to motivate individuals to act for the greater good, even when doing so is personally costly (e.g., effortful, time-consuming, expensive, uncomfortable, or dangerous). The suggestions here are not unique to any one particular issue, and thus are inherently non-partisan. We aim to provide a toolkit that can be applied to a wide range of issues, as desired by any given policy maker.

To this end, we leverage a large body of interdisciplinary research that uses game theory and psychology experiments to study human cooperation (for a more detailed overview, see Rand & Nowak, 2013). In this literature, cooperative (or

¹Yale University, New Haven, CT, USA

²Harvard University, Cambridge, MA, USA

³Massachusetts Institute of Technology, Cambridge, USA

Corresponding Author:

David G. Rand, Department of Psychology, Department of Economics, School of Management, Yale University, 2 Hillhouse Ave., New Haven, CT 06511, USA.

Email: david.rand@yale.edu

“prosocial”) behavior is defined as paying a personal cost to provide a benefit to one or more others. Costs are defined broadly, and may include time, effort, resources, money, discomfort, and physical harm. For example, it is cooperative to pay the full amount due on your taxes although the likelihood of being caught for underreporting income is low (Spicer & Thomas, 1982), to vote although your individual ballot will not change the election’s outcome (Gelman, Silver, & Edlin, 2012), to join the military and put yourself at risk to protect your country, and to volunteer your time at (and donate your money to) charitable organizations. At a more personal level, it is also cooperative to lend money to a friend, to help your neighbors move into their house, to do your fair share on a group project in the workplace, and to collaborate in good faith with someone, although you have different beliefs or opinions.

What all of these examples have in common is that they describe behaviors that (a) increase social welfare, which is to say that societies in which people do engage in these behaviors are better off than societies in which people do not, but (b) are individually costly, such that purely selfish and self-interested people will not choose to engage in them. Thus, a tension exists between what is best for the individual and what is best for the society as a whole. Situations that involve this tension are called “social dilemmas,” and it is the resolution of such dilemmas that is the focus here.

Framework

Understanding Social Dilemmas Using Game Theory

To study decision making, game theorists create simple games in which each player has a set of choices, and receives a payoff based on his or her own choice and the choices of the other players. For thinking about cooperation in social dilemmas, the simplest and most common of such games is the famous “Prisoner’s Dilemma” (PD). In the PD, each of two players simultaneously chooses whether to pay a cost c to give the other person a benefit b , with $b > c > 0$ (i.e., to “cooperate”), or to do nothing (i.e., to “defect”). If both players cooperate, they each earn payoffs of $b - c$, which is greater than 0. If both players defect, neither earns anything (because they both do nothing). So a pair of cooperators does better than a pair of defectors, illustrating the collective benefit of cooperation. Yet, regardless of what the other player chooses, defecting always earns you more than cooperating (because you avoid paying the cost of cooperation). So a selfish individual will always defect; and as a result, the collective benefits of cooperation will be missed. The PD therefore illustrates the tension between individual and collective interests that is at the heart of social dilemmas such as tax compliance, voter participation, environmental conservation, and cooperation across political or ideological lines.

The PD can expand to consider cooperation among groups using the “Public Goods Game” (PGG). In the PGG, each of n group members chooses how much money to keep, and how much to contribute to a “public good.” All contributions are multiplied by r and split evenly among all n group members, with $r < n$. If everyone contributes fully, everyone’s money is multiplied by r , and group welfare is substantially increased. However, because contributions are multiplied by r and then split n ways, each individual receives back only r/n (< 1) dollar for each dollar contributed; therefore, players individually lose money on contributing and maximize their payoff by free riding off the contributions of others. Although a group of full contributors earns more than a group of non-contributors, the best possible payoff comes from being in a group where everyone *else* contributes while you do not. The PGG represents all kinds of situations involving public goods where the benefits created are shared equally by all (regardless of their personal contribution level), such as clean air, national security, knowledge creation, websites such as Wikipedia, non-toll roads, and cultures of trust.

Reciprocity: The Key to Promoting Cooperation in Social Dilemmas

The question, then, is how to get people to act cooperatively and help to provision these public goods, although contributing is personally costly. Many (and perhaps most) people will cooperate, at least to some extent, just based on an intrinsic desire to help others and to do the right thing. Such intrinsic prosociality, however, will only get us so far as a society. Selfish actors will exploit these more altruistic people and undermine their prosocial attempts to provision public goods. Most research on human cooperation has therefore focused on identifying elements that can be added to social dilemmas to *resolve* the dilemma, by aligning the interests of the individual and the group.

The most effective solutions for human social dilemmas typically involve the principle of “reciprocity” (Axelrod, 1984; Fudenberg & Maskin, 1986; Nowak & Sigmund, 1993, 2005; Trivers, 1971). Reciprocity takes into account the fact that most important interactions involve repeated interactions and reputational concerns, and thus, most actions have future consequences. Imagine what would happen if you played a series of PDs with the same partner, rather than just a single one-shot interaction. Now, defecting no longer necessarily maximizes individual payoffs: Suppose your partner will cooperate with you in the next interaction, but only if you cooperate in this interaction. With such a partner, a self-interested player will choose to cooperate to gain the benefit of the partner’s reciprocal cooperation in the future. That is, in the context of long-term relationships, it is short-sighted to defect: Although you might be able to take advantage of a friend or colleague today, he or she will probably find out, and you will then pay a price for your bad action later. Repeated

interactions allow cooperative people to avoid being exploited by defectors and thus stabilize cooperation.

This principle of reciprocity also extends to group interactions and public goods. To see how, remember that group interactions do not occur in a vacuum, but rather are superimposed on a network of pairwise personal relationships. If your partner is only willing to cooperate with you in the next interaction if you both (a) cooperate in your personal relationship in this interaction, *and* (b) contribute to the public good, then you have an incentive to do both. For this to work, however, your decisions in the public good context must be observable by your personal interaction partners (Ellingsen, Herrmann, Nowak, Rand, & Tarnita, 2012; Rand, Dreber, Ellingsen, Fudenberg, & Nowak, 2009). You may be less likely to help neighbors in need if they underreported their income to the Internal Revenue Service, or did not bother to vote in the last election, or have not upgraded their extremely inefficient old furnace—but only if you are aware of these pieces of information.

Practical Lessons

The remainder of this article outlines practical lessons for policy makers gleaned from research on reciprocity.

Make Behavior Observable

To be rewarded for contributing to public goods, we must be observed contributing. For instance, consider participants playing a repeated PGG in the laboratory for real money, along with pairwise PD with each group member. When they were informed of the public goods contributions of each group member, they preferentially cooperated as individuals with high contributors. As a result, people learned to contribute over time. When the PGG was not linked to the PDs (i.e., public goods contributions were not observable), conversely, cooperation failed, and they learned to free ride (Ellingsen et al., 2012; Rand et al., 2009). Similarly, participants contributed substantially more in a PGG where their photos were shown to other players, compared with a fully anonymous game (Andreoni & Petrie, 2004).

Ample evidence of the power of observability also comes from real-world field experiments, which find large effects of making behavior observable in settings as diverse as blood donation (Lacetera & Macis, 2010), blackout prevention (Yoeli, Hoffman, Rand, & Nowak, 2013), and support for national parks (Alpizar, Carlsson, & Johansson-Stenman, 2008). In Switzerland, voting rates fell in small communities when voters were given the option to vote by mail (Funk, 2010), which makes it harder to tell who did not vote. As one would expect from the game theory models, increases in contributions are largest when the observers are more important for future interactions. For example, in the blackout-prevention study, making

participation in the program observable increased participation 4.7 times more among homeowners than renters (renters move frequently and thus are less invested in relationships with their neighbors). Subtle cues of being observed also have dramatic effects on contributions, both in the lab (Haley & Fessler, 2005) and in the field. For example, displaying stylized “eyespot,” which induce an unconscious sense of being watched, increased the money donated to an honor-system payment box in a cafeteria by threefold (Bateson, Nettle, & Roberts, 2006).

Thus, policy makers and practitioners should make contributions to public goods observable, particularly to people who are important future interaction partners. The field experiments provide some guidance on how to implement this suggestion. A newsletter with lists of donors increased blood donations (Lacetera & Macis, 2010), and one could easily imagine using newsletters—print or online—to increase observability in other settings. Enrolling customers in a utility’s blackout-prevention program using publicly visible sign-up sheets, instead of asking customers to register by phone, led to substantially higher participation (Yoeli et al., 2013). Whenever those you are attempting to reach live in dense housing with shared common space, this solution is feasible and cost-effective.¹ Simply placing the donation bin for a national park within view of a park ranger, again, can often be costlessly implemented (Alpizar et al., 2008). Other examples provide potentially useful models as well. The Prius is thought to have been successful in large part because of its instantly recognizable design (Kahn, 2007), voter participation has successfully been encouraged with “I voted” stickers, as well as Facebook badges that informed members’ friends when members clicked on an “I voted” button (Bond et al., 2012). Facebook has similarly encouraged organ donations (Cameron et al., 2013). Social media offers great promise for making public goods contributions observable; researchers and practitioners have only scratched the surface. Such observability efforts can be particularly effective if they are targeted at peers and neighbors (important future or potential interaction partners). Again, social media offers a fruitful platform, due to the ease of identifying contributors’ most-frequent interaction partners.

To be sure, contributions cannot always be observed by others, but some of these cases could still use subtle cues, including abstract shapes that vaguely resemble faces or eyes (similar to those used by Haley & Fessler, 2005). We are inspired by a group of our students who, when some of their dorm-mates ate more than their fair share of free pizza that was left out in the dorm’s common area, placed a life-sized cardboard cutout of a person near the pizza boxes. An image of a park ranger or park-goer near an unstaffed donation box at a national park might similarly reduce the number of visitors who passed by without giving.

Although observability offers great promise for increasing contributions to public goods, there are caveats.

First, although people are motivated to contribute by reputational benefits, they also are less cooperative when their cooperation can be construed as selfishly motivated for the explicit purpose of achieving a good reputation (Benabou & Tirole, 2006; Hoffman, Yoeli, & Nowak, 2014). Making reputational benefits explicit can therefore have perverse effects on contributions (Frey & Oberholzer-Gee, 1997; Gneezy & Rustichini, 2000). Thus, in making contributions observable, policy makers simultaneously must ensure that it is not explicit that the goal is to create observability. Sign-up sheets, relocated donation boxes, and eyespots naturally achieve this aim, as their purpose is not obvious; newsletters should perhaps be sent under other auspices. When observability is not subtle, as with Facebook badges, we recommend stating that observability has some purpose other than garnering reputational benefits, for example, helping to convince others to contribute.

Second, making contributions to make public goods observable should work only when observers recognize that public good as desirable. In one experiment (Ariely, Bracha, & Meier, 2009), observability *reduced* donations to a conservative charity (the National Rifle Association) that was unpopular on the liberal college campus where they ran their experiments. Charities seeking donations in areas where their causes are unpopular might do better by maintaining donor anonymity.

Focus on Effectiveness

Reciprocity typically rewards people for their actions more than for the consequences of those actions (which are often less easily observable). Thus, people tend to be undermotivated to pay attention to the efficacy of their contributions. In one study, researchers explicitly asked donors whether they considered effectiveness when giving to charities, and only 6% responded affirmatively (Baltussen, Sylla, Frick, & Mariotti, 2005). In another study, researchers varied the size of a matching donation (and thus the efficacy of donating) and offered subjects the opportunity to pay a small amount to reveal the size of the match. Subjects typically failed to act on this opportunity (Null, 2011). Complementing this, little change in giving behavior occurred when a matching grant was changed from 2 to 3 times the donor's gift (Karlan & List, 2007). Furthermore, a large literature on "scope neglect" finds no difference in subjects' willingness to pay to save 2,000 versus 200,000 birds, or 1 versus 40 clean ponds (Desvousges et al., 1992). Finally, to the extent that people do sometimes try to evaluate different charities, they typically focus on minimizing overhead ratios (i.e., administrative expenses) rather than maximizing cost-effectiveness (i.e., number of saved lives per dollar), because the former attribute is easier to evaluate than the latter (Caviola, Faulmüller, Everett, Savulescu, & Kahane, 2014).

We propose the following solution: Make effectiveness of contributions known not only to the contributor but also to observers. For example, when it comes to charitable donations, we echo the recommendation of prominent figures such as Peter Singer (2010), who have been urging donors to evaluate the efficiency of the charities they donate to, using resources such as www.CharityNavigator.com and GiveWell. These resources facilitate identification of the more effective charities. Furthermore, make such ratings observable by third parties as much as possible, for example, by being featured in acknowledgements of donations, and otherwise incorporated into the public discourse on donations. This will give donors a stronger incentive to consider efficacy when making their contribution choices. Preliminary work supports this suggestion: In an experiment, donors gave the greatest share to the most effective charities when they knew the effectiveness of their contributions would be known not only by them but also by observers (Hauser, Hoffman, & Yoeli, 2014). Similarly, consider efforts by agencies and utilities to encourage homeowners to invest in energy-efficiency improvements. The Department of Energy recently released a tool called Home Energy Score (HEScore) that does for home improvements what CharityNavigator does for charities, making it easier for homeowners to identify the most effective improvements. Such a tool is additionally useful when advertising homeowners' improvements on websites managed by utilities or in realtors' Multiple Listing Service. When implementing these solutions, potential contributors must know that effectiveness will be public at the time when they make their contributions.

Do Not Let People "Avoid the Ask"

Often, people have the chance to avoid putting themselves in a position where they would have the chance to contribute. For example, supermarket shoppers often avoid doors at which Salvation Army volunteers are stationed (Andreoni, Rao, & Trachtman, 2011), and experimental subjects who told experimenters they would donate money to their partner in a dictator game, then chose to pay a small fee to keep the remaining money and not inform the potential recipient that a dictator game was supposed to take place (Dana, Cain, & Dawes, 2006). Crossing the street to avoid a panhandler is an example from our everyday life. By "avoiding the ask" in this way, people can choose not to contribute while incurring less reputational risk—they avoid having to obviously say "no," by hiding the fact that they had the opportunity to contribute. This then makes it less costly to free ride, and it reduces contributions.

Therefore, when possible, policy makers should not give people the opportunity to "avoid the ask." Outreach should be designed to prevent potential contributors from plausibly claiming to have been unaware of the solicitation, or to not allow them to avoid being seen walking away from a

contribution opportunity. Those of you who have encountered fundraisers on the sidewalk know that many of them have already incorporated this takeaway: They personally greet all passers-by, “Hi, do you have a minute to help save the rainforest?” Similarly, when possible, make contribution the default, such that choosing to not contribute requires an explicit action. Organ donation defaults provide a potent example (Johnson & Goldstein, 2003). In European countries that make organ donation the default, more than 97% of drivers register as donors. In countries where drivers must opt in to donating, an average of only 15% of drivers register. The New York Metropolitan Museum of Art provides another helpful example: Admission is free, but one is charged the suggested donation of US\$25 by default and must request to change this amount if one wishes to pay less.

Project a “Norm” of Cooperation

Whether or not one gets punished or rewarded for one’s contribution to a particular public good depends on whether contributing is the norm. In contexts where most people contribute, and withhold cooperation from non-contributors, it pays to contribute if your behavior is observable. However, if most people do not contribute, and people will happily cooperate with non-contributors, it does *not* pay to contribute, even when others can observe your choice (Ellingsen et al., 2012). This suggests that people should attend to the contribution behavior of others (i.e., the norms in a given context) when deciding whether to contribute.

Indeed, in a wide range of studies, most people are “conditionally cooperative”—they are willing to cooperate more when they believe others contribute more. For example, students asked to donate to a university charity gave 2.3 percentage points more when told that others had given at a rate of 64% than when they were told giving rates were 46% (Frey & Meier, 2004). Hotel patrons were 26% more likely to reuse their towels when informed others had done the same (Goldstein, Cialdini, & Griskevicius, 2008). Households have been shown to meaningfully reduce electricity consumption when told neighbors are consuming less, both in the United States (Ayres, Raseman, & Shih, 2012) and in India (Sudarshan, 2014).

Thus, let potential contributors know when cooperation is high, by informing them of contribution rates. As the examples above illustrate, this is typically straightforward: “Our guests don’t litter. You shouldn’t either.” is an effective prescription against littering; “65% of people buy smaller cars to help reduce emissions. Think big, buy small” will help direct car buyers toward smaller cars; and “Most of our members contribute. You should, too.” can drive donation rates higher (see Cialdini, 1993, for a detailed discussion). Some have suggested that charity seed money helps encourage contributions because it indicates high contribution rates. Partially filled, see-through public donation boxes may work for the same reason. One would

do well to never completely empty the box, following the lead of panhandlers and baristas, who always leave a few coins in their cup.

Importantly, one should also avoid providing such information when contribution rates are low, especially explicitly, but also implicitly. A TV commercial urging viewers to eschew littering, in which litter was ubiquitous, could very likely have given viewers the impression that littering was acceptable (Cialdini, 1993).

Create Habits of Cooperation

Choosing to contribute can pay off in the long run because of reciprocity effects, but only in the context of norms that endorse contributing. As a large body of evidence from social and cognitive psychology suggests, we internalize behaviors that are typically successful, and we adopt them as our defaults (Gigerenzer & Goldstein, 1996; Kahneman, 2003). This suggests that individuals who generally interact in environments where cooperation is advantageous will be predisposed toward contributing, even in atypical contexts when it does not actually pay off, whereas those for whom defection is typically successful will internalize non-contribution (Rand et al., 2014). In support of such automaticity, participants in public-goods experiments contribute less when they are forced to carefully consider their decisions (Rand, Greene, & Nowak, 2012), but only if they generally trust their daily interaction partners (Rand & Kraft-Todd, in press).

These findings emphasize the importance of organization structures and institutions aimed at giving people more opportunities to experience the payoff benefits of cooperation. For example, cross-nation correlations suggest that people living in countries with less corruption and stronger rule of law (factors that incentive cooperative behavior) are less likely to behave antisocially (Ellingsen et al., 2012; Gächter & Herrmann, 2009; Herrmann, Thoni, & Gächter, 2008). And across a range of societies from hunter-gatherers to industrialized nations, greater market integration was associated with greater generosity in lab experiments, with the argument that markets are institutions that facilitate productive (cooperative) interactions with strangers (Henrich et al., 2010).

Direct experimental evidence also supports the effect of social environment on cooperation. In one set of experiments, subjects were randomly assigned to play a series of repeated PDs either under a set of “good” rules that favored cooperation (high likelihood of future consequences for present action) or a set of “bad” rules that favored defection (low likelihood of future consequences for present action); afterward, all subjects played an identical battery of one-shot anonymous games (Peysakhovich & Rand, 2013). Subjects assigned to “good” environments contributed more to public goods and were more altruistic, trusting, and trustworthy in the subsequent one-shot games, as well as more inclined to punish

selfishness. Furthermore, this effect was especially strong among individuals who relied more on their intuitive, default responses.

Thus, great care should be paid to the kind of prosocial or selfish behavior that is rewarded, as such behavior often becomes internalized. The manner in which rewards are given is also essential: When rewards are explicit (i.e., you get a bigger bonus for good behavior or incur a fine for bad behavior), the consequences can be disastrous. Explicit incentives can “crowd out” intrinsic motivation to help, and actually reduce prosociality. Thus, incentivizing good behavior must use more subtle tools, such as repeated interactions and reputational concerns.

Conclusion

Research on cooperation teaches us simple lessons that can have a big impact on contributions to public goods, such as making contributions observable, making effectiveness public, eliminating opportunities to avoid the ask, advertising others’ high contribution rates, giving people experience in cooperative situations, and letting people choose their peers. In many cases, existing studies demonstrate the feasibility, cost-effectiveness, and potential impact of these lessons. We hope that policy makers and practitioners will find these suggestions helpful in their quest to solve society’s public goods problems, however small or large.

Acknowledgment

We thank Jillian Jordan and Jim Everett for helpful comments and suggestions.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: We gratefully acknowledge financial support from the John Templeton Foundation.

Note

1. It is interesting and potentially useful to note that in this experiment, when participation was anonymous, framing participation as a contribution to the public good actually reduced participation. We thus caution against using public goods messaging when contributions are not observable.

References

- Alpizar, F., Carlsson, F., & Johansson-Stenman, O. (2008). Anonymity, reciprocity, and conformity: Evidence from voluntary contributions to a national park in Costa Rica. *Journal of Public Economics*, 92, 1047-1060.
- Andreoni, J., & Petrie, R. (2004). Public goods experiments without confidentiality: A glimpse into fund-raising. *Journal of Public Economics*, 88, 1605-1623.
- Andreoni, J., Rao, J. M., & Trachtman, H. (2011). *Avoiding the ask: A field experiment on altruism, empathy, and charitable giving* (No. w17648). National Bureau of Economic Research.
- Ariely, D., Bracha, A., & Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99, 544-555.
- Axelrod, R. (1984). *The evolution of cooperation*. New York, NY: Basic Books.
- Ayres, I., Raseman, S., & Shih, A. (2012). Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *Journal of Law, Economics, & Organization*. Advance online publication.
- Baltussen, R. M., Sylla, M., Frick, K. D., & Mariotti, S. P. (2005). Cost-effectiveness of trachoma control in seven world regions. *Ophthalmic Epidemiology*, 12, 91-101.
- Bateson, M., Nettle, D., & Roberts, G. (2006). Cues of being watched enhance cooperation in a real-world setting. *Biology Letters*, 2, 412-414.
- Bénabou, R., & Tirole, J. (2005). *Incentives and prosocial behavior* (No. w11535). National Bureau of Economic Research.
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489, 295-298.
- Cameron, A. M., Massie, A. B., Alexander, C. E., Stewart, B., Montgomery, R. A., Benavides, N. R., . . . Segev, D. L. (2013). Social media and organ donor registration: The Facebook effect. *American Journal of Transplantation*, 13, 2059-2065.
- Caviola, L., Faulmüller, N., Everett, J. A. C., Savulescu, J., & Kahane, G. (2014). The evaluability bias in charitable giving: Saving administration costs or saving lives? *Journal of Judgment and Decision Making*, 9, 303-315.
- Cialdini, R. B. (1993). *Influence: The psychology of persuasion*.
- Dana, J., Cain, D. M., & Dawes, R. M. (2006). What you don’t know won’t hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes*, 100, 193-201.
- Desvousges, W. H., Johnson, F. R., Dunford, R. W., Boyle, K. J., Hudson, S. P., & Wilson, K. N. (1992). *Measuring nonuse damages using contingent valuation: An experimental evaluation of accuracy*. Research Triangle Institute.
- Ellingsen, T., Herrmann, B., Nowak, M. A., Rand, D. G., & Tarnita, C. E. (2012). *Civic capital in two cultures: The nature of cooperation in Romania and USA*. Retrieved from <http://ssrn.com/abstract=2179575>
- Frey, B. S., & Meier, S. (2004). Social comparisons and pro-social behavior: Testing “conditional cooperation” in a field experiment. *American Economic Review*, 94, 1717-1722.
- Frey, B. S., & Oberholzer-Gee, F. (1997). The cost of price incentives: An empirical analysis of motivation crowding-out. *American Economic Review*, 87, 746-755.

- Fudenberg, D., & Maskin, E. S. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, *54*, 533-554.
- Funk, P. (2010). Social incentives and voter turnout: Evidence from the Swiss mail ballot system. *Journal of the European Economic Association*, *8*, 1077-1103.
- Gächter, S., & Herrmann, B. (2009). Reciprocity, culture, and human cooperation: Previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*, 791-806.
- Gelman, A., Silver, N., & Edlin, A. (2012). What is the probability your vote will make a difference? *Economic Inquiry*, *50*, 321-326.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650-669.
- Gneezy, U., & Rustichini, A. (2000). Pay enough or don't pay at all. *Quarterly Journal of Economics*, *115*, 791-810.
- Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, *35*, 472-482.
- Haley, K. J., & Fessler, D. M. (2005). Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evolution & Human Behavior*, *26*, 245-256.
- Hauser, O., Hoffman, M., & Yoeli, E. (2014). *Increasing the efficiency of charitable giving* (Harvard University Working Paper).
- Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., . . . Ziker, J. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science*, *327*, 1480-1484.
- Herrmann, B., Thoni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science*, *319*, 1362-1367.
- Hoffman, M., Yoeli, E., & Nowak, M. (2014). *Cooperating without looking* (Harvard University Working Paper).
- Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*, *302*, 1338-1339.
- Kahn, M. E. (2007). Do greens drive Hummers or hybrids? Environmental ideology as a determinant of consumer choice. *Journal of Environmental Economics and Management*, *54*, 129-145.
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, *58*, 697-720.
- Karlan, D., & List, J. (2006). *Does price matter in charitable giving? Evidence from a large-scale natural field experiment* (No. w12338). National Bureau of Economic Research.
- Lacetera, N., & Macis, M. (2010). Social image concerns and prosocial behavior: Field evidence from a nonlinear incentive scheme. *Journal of Economic Behavior & Organization*, *76*, 225-237.
- List, J. A., Berrens, R. P., Bohara, A. K., & Kerkvliet, J. (2004). Examining the role of social isolation on stated preferences. *American Economic Review*, *94*, 741-752.
- Nowak, M. A., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, *364*, 56-58.
- Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, *437*, 1291-1298.
- Null, C. (2011). Warm glow, information, and inefficient charitable giving. *Journal of Public Economics*, *95*, 455-465.
- Peysakhovich, A., & Rand, D. G. (2013). *Habits of virtue: Creating norms of cooperation and defection in the laboratory*. Retrieved from <http://ssrn.com/abstract=2294242>
- Rand, D. G., Dreber, A., Ellingsen, T., Fudenberg, D., & Nowak, M. A. (2009). Positive interactions promote public cooperation. *Science*, *325*, 1272-1275.
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, *489*, 427-430.
- Rand, D. G., & Kraft-Todd, G. T. (in press). Reflection does not undermine self-interested prosociality. *Frontiers in Behavioral Neuroscience*.
- Rand, D. G., & Nowak, M. A. (2013). Human cooperation. *Trends in Cognitive Sciences*, *17*, 413-425.
- Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., & Green, J. D. (2014). Social heuristics shape intuitive cooperation. *Nature Communications*, *5*, Article 3677.
- Singer, P. (2010). *The life you can save: How to do your part to end world poverty*. Random House.
- Spicer, M. W., & Thomas, J. E. (1982). Audit probabilities and the tax evasion decision: An experimental approach. *Journal of Economic Psychology*, *2*, 241-245.
- Sudarshan, A. (2014). *Nudges in the marketplace: Using peer comparisons and incentives to reduce household electricity consumption* (Harvard University Working Paper).
- Trivers, R. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, *46*, 35-57.