

Powering up with indirect reciprocity in a large-scale field experiment

Erez Yoeli^a, Moshe Hoffman^{b,c}, David G. Rand^{c,d}, and Martin A. Nowak^{c,e,f,1}

^aFederal Trade Commission, Washington, DC 20580; ^bDepartment of Computer Science and Engineering, University of California at San Diego, La Jolla, CA 92093; ^cProgram for Evolutionary Dynamics, Harvard University, Cambridge, MA 02138; ^dDepartment of Psychology, Yale University, New Haven, CT 06520; and Departments of ^eMathematics and ^fOrganismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138

Edited by John C. Avise, University of California, Irvine, CA, and approved April 3, 2013 (received for review February 14, 2013)

A defining aspect of human cooperation is the use of sophisticated indirect reciprocity. We observe others, talk about others, and act accordingly. We help those who help others, and we cooperate expecting that others will cooperate in return. Indirect reciprocity is based on reputation, which spreads by communication. A crucial aspect of indirect reciprocity is observability: reputation effects can support cooperation as long as peoples' actions can be observed by others. In evolutionary models of indirect reciprocity, natural selection favors cooperation when observability is sufficiently high. Complimenting this theoretical work are experiments where observability promotes cooperation among small groups playing games in the laboratory. Until now, however, there has been little evidence of observability's power to promote large-scale cooperation in real world settings. Here we provide such evidence using a field study involving 2413 subjects. We collaborated with a utility company to study participation in a program designed to prevent blackouts. We show that observability triples participation in this public goods game. The effect is over four times larger than offering a \$25 monetary incentive, the company's previous policy. Furthermore, as predicted by indirect reciprocity, we provide evidence that reputational concerns are driving our observability effect. In sum, we show how indirect reciprocity can be harnessed to increase cooperation in a relevant, real-world public goods game.

evolutionary game theory | experimental economics

Cooperation occurs when we take on costs to benefit the greater good. By this definition, everyone is better off when everyone cooperates, but self-interest undermines cooperation and leads to free-riding. Promoting cooperation is a central challenge for human societies, both today and over our evolutionary history (1–10). There are five mechanisms for the evolution of cooperation (11): direct and indirect reciprocity, spatial selection, group selection, and kin selection. Each of these mechanisms is an interaction structure that can lead cooperators to outperform noncooperators, and therefore be favored by selection.

Direct and indirect reciprocity involve repeated interactions, creating future consequences for one's actions: it can pay to cooperate today to receive cooperation from others tomorrow. Spatial selection occurs when players' interactions are structured rather than occurring at random. As a result, cooperators may be more likely to interact with other cooperators and thus preferentially receive the benefits of cooperation. Spatial selection operates when cooperators cluster in physical space, on social networks, in sets, or in phenotype space (12). Group selection (or multilevel selection) occurs when competition and reproduction happen at multiple levels: not only do players compete with others in their group, but groups compete with each other. If cooperative groups outcompete noncooperative groups, then group-level selection can favor the evolution of cooperation. Finally, kin selection may be defined as preferring to cooperate with those who are closely related. Kin recognition can allow players to cooperate with close genetic relatives and defect otherwise.

Most of the literature on the evolution of cooperation uses the Prisoner's Dilemma and related frameworks: players can pay a cost to give a greater benefit to one or more others. Thus, within the context of these games, cooperation is good for everyone. However, cooperation need not be good for everyone more generally (6, 13). There are situations in which cooperating may give a benefit to some, but impose costs on others. For example, in intergroup conflict and war, people cooperate with members of their own group in an attempt to harm members of other groups (14). Or in the context of markets, companies may collude to keep prices high, benefiting each other but harming consumers (15). The five mechanisms for the evolution of cooperation may promote both total welfare-enhancing cooperation, as well as these more pernicious forms of cooperation.

All of these mechanisms are relevant for the evolution of human cooperation, but direct reciprocity and indirect reciprocity occupy a central place: most of our key interactions are repeated and reputation is usually at stake. Direct reciprocity is based on repeated encounters between the same two individuals: my behavior toward you depends on what you have done to me. Indirect reciprocity is based on repeated encounters in a group of individuals: my behavior toward you also depends on what you have done to others (Fig. 1). We take a keen interest in who does what to whom and why, which requires sophisticated social intelligence. We talk to each other about others. As David Haig said: "For direct reciprocity you need a face, for indirect reciprocity you need a name" (4). The evolution of indirect reciprocity is linked to the evolution of human language. Supported by human language, reputation systems allow us to track the good and bad behavior of others and to use this information to incentivize cooperation. Whatever is specifically human about our mental machinery is derivative of human language, social intelligence, and thus indirect reciprocity (4, 16).

The evolution of cooperation via indirect reciprocity has been a topic of great interest in recent years. Mathematical models and computer simulations have demonstrated the power of indirect reciprocity for promoting cooperative behavior (17–40). In these models, players typically engage in a series of one-shot interactions with others selected at random from the population. In some of those interactions, players' previous decisions are observable by their partners. Observability allows players to use conditional strategies that base their actions on the partner's behavior in

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, "In the Light of Evolution VII: The Human Mental Machinery," held January 10–12, 2013, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, CA. The complete program and audio files of most presentations are available on the NAS Web site at www.nasonline.org/evolution_vii.

Author contributions: E.Y. designed research; E.Y. performed research; E.Y., M.H., D.G.R., and M.A.N. analyzed data; and E.Y., M.H., D.G.R., and M.A.N. wrote the paper.

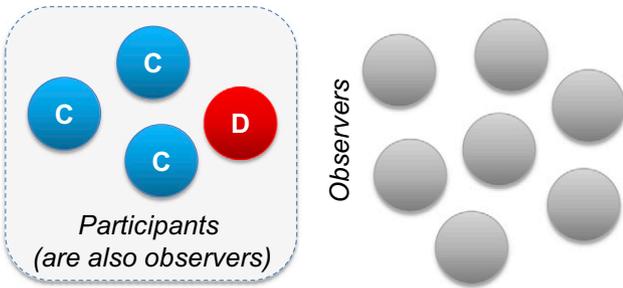
The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: martin_nowak@harvard.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1301210110/-DCSupplemental.

1. Public goods game



2. Gossip, Communication, Evaluation

3. Conditional response:

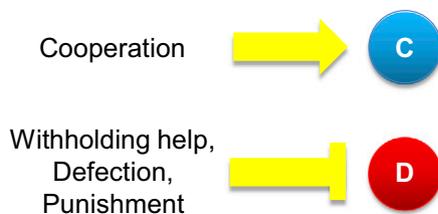


Fig. 1. Indirect reciprocity can support contributions to public goods, as depicted here. In indirect reciprocity, my behavior toward you depends on how you have behaved toward both myself and others. This process occurs in three stages. First, people engage in a public goods game (cooperators in blue and defectors in red). Peoples' behavior is observed, both by other players and third parties. Second, information about this behavior can spread from person to person. Based on the information received, peoples' opinions about the players are updated. Third, as a result, public goods cooperators will receive cooperation in future interactions whereas defectors will be denied cooperation, defected upon, or punished. Thus, indirect reciprocity creates an incentive to contribute to public goods and can promote the evolution of cooperation.

the past. When past actions are sufficiently observable, natural selection can favor strategies that cooperate as long as the partner has behaved well in the past. What constitutes “good” behavior worthy of receiving cooperation depends on the social norm. A simple social norm is called “image scoring” and prescribes cooperating with those who have cooperated in sufficiently many previous interactions (20). More complicated norms also take into account the behavior of the partner's previous partners. For example, under the “standing” norm, players can maintain their good reputation by defecting against those with bad reputation (24). The many models of indirect reciprocity differ in their details, yet across a wide range of assumptions, making previous decisions observable allows cooperators to selectively target their cooperation at other cooperators and withhold cooperation from defectors. Thus, free-riders are at a disadvantage, and cooperation can spread.

This body of theoretical work is supported by behavioral experiments where subjects play economic games in the laboratory. People are substantially more cooperative when their decisions are observable and when others can respond accordingly (41–60). Subjects understand that having a good reputation is valuable in these settings (49) and so are willing to pay the cost of cooperation. Observability particularly increases cooperation when the prosocial nature of the cooperative choice is made salient (55, 61). Moreover, experimental evidence indicates that indirect reciprocity is deeply entrenched in human psychology: subtle cues of

observability have large effects on cooperation levels (62–65), and our initial impulse to cooperate in one-shot anonymous settings (66–69) is likely the result of adaptation in a world dominated by reputational concerns (66, 68).

These laboratory experiments are extremely valuable. They generate powerful insights into human psychology and provide clear evidence for the importance of indirect reciprocity. To do so, however, they typically use abstract economic games and involve the interaction of only a handful of subjects. Thus, the question of whether observability affects large-scale cooperation in real world settings outside of the laboratory remains largely unexplored (exceptions include refs. 70–72). The extent to which findings from theory and the laboratory generalize to natural field settings is of great importance, both for scientific understanding and for public policy (73).

Here, we address this question by running a large-scale field experiment on the effect of observability in a public goods game (PGG). We collaborated with a major electric utility company to enroll consumers in a “demand response” program. This program is designed to help prevent blackouts by reducing excessive use of air conditioning during periods of high electricity demand. The cost of electricity production can spike hundredsfold during demand peaks. However, the price consumers pay is typically constant across time. Thus, during peak periods there is a dramatic mismatch between price and actual cost, leading to excessive energy use. This mismatch reduces grid reliability, drives up energy costs, increases the risk of blackouts, and harms the environment. In recent years, reducing excessive peak energy use has become a target of regulatory efforts to increase efficiency in the electricity industry.

Encouraging participation in demand response programs such as the one used in this study is the primary policy tool available for reducing peak energy use (74). Demand response programs are voluntary programs in which people allow their utility to remotely restrict their energy consumption during peak hours. To do so, the utility usually installs a remote switch in-line with the circuitry of an appliance such as a hot water heater or air conditioner. Estimates suggest that these voluntary programs could reduce the need to invest in additional generation capacity by at least 38% over the next two decades, generating cost savings of at least \$129 billion (75). Voluntary energy efficiency and demand response programs have been widely available for years, but participation is frustratingly low (76). Demand response programs exemplify the public goods dilemma: participation helps reduce on-peak demand, benefitting all energy grid users, but energy consumers find participating inconvenient. Participation is socially optimal because the inconvenience is minimal for most individuals relative to the societal costs of a black out.

To explore the effect of observability on this real-world public goods problem, we solicited residents of 15 homeowners associations (HOAs) to participate in a demand response program. Residents who volunteered for this program allowed the utility to install a device that remotely curbs their central air conditioners when necessary: on days with unusually high demand or in the case of an unexpected plant or transmission failure. Residents who volunteered, therefore, contributed to a public good by improving the stability of the electrical grid in all of California, at the cost of some personal inconvenience. We solicited volunteers by delivering mailers to residents and asking them to participate. Sign-up sheets were posted in a communal area near their home, usually by a shared mailbox kiosk. In our primary manipulation, we varied whether residents' neighbors could tell who had signed up for the program. We did so by varying whether the publicly posted sheets required residents to print their name and unit number (observable treatment) or only a code that does not reveal their identity (anonymous treatment).

Results

We found that residents in the observable treatment are nearly three times as likely to participate in the demand response program as residents in the anonymous treatment (fraction of residents participating: anonymous = 0.030, observable = 0.088, $P < 0.01$, $n = 1408$; Fig. 2). All statistics presented are from probit regressions including various controls, with SEs clustered at the HOA level; for details and regression tables, see *Supporting Information*.

The effect of the observable treatment was over seven times that of offering a \$25 incentive (the estimated effect of the incentive is 0.009; a Wald test rejects that the coefficients on observability and the \$25 incentive are identical, $P = 0.024$). This incentive was what the utility had used before the experiment, and they had previously argued the incentive would be far more effective than observability. In fact, this incentive appears to have been too small to be effective, and such small financial incentives are known to sometimes backfire (77). For the sake of comparison, we followed convention and estimated how large the financial incentive would have to be to achieve the same results if its effect is linear (78). We found that the utility would have had to offer an incentive of \$174 to increase participation as much as our observable treatment.

We now explore the mechanism through which observability functions to increase participation. Indirect reciprocity theory is based on reputational concerns: when groups of people interact repeatedly and actions are observable, it becomes advantageous to be seen contributing to public goods. Based on this account, we predict that observability will have a greater effect among populations where ongoing relationships and reputations are expected to play a larger role. We evaluate this prediction in two ways.

First, we test whether the effect of the observable treatment was greater in apartment buildings compared with row houses and individual homes. In apartment buildings, residents are more likely to interact with their neighbors in public spaces, and sign-up sheets were typically posted in especially conspicuous locations. Thus, indirect reciprocity theory predicts that observability will have a larger effect in apartment buildings. As shown in Fig. 3A, the results confirm this prediction: observability increased participation among those living in apartment buildings (fraction of residents participating: anonymous = 0.048, observable = 0.114, $P < 0.01$, $n = 582$) whereas it had little effect on the inhabitants of row houses or individual homes (fraction of residents participating:

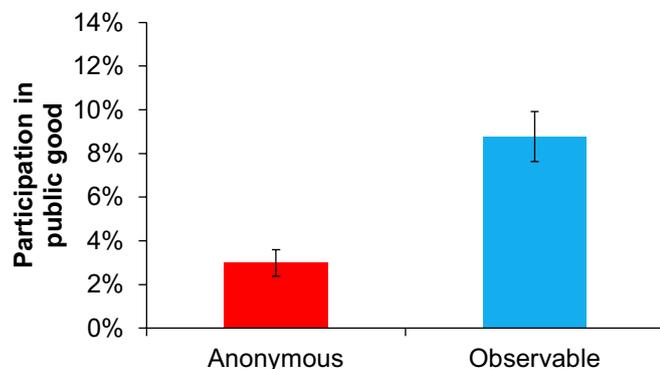


Fig. 2. We solicited 1,408 customers of a major electric utility for participation in a program designed to prevent blackouts. Residents signed up for the program on sheets posted in a communal area near their home, usually by a shared mailbox kiosk. We varied whether residents' neighbors could tell who signed up for the program: publicly posted sheets required residents to print their name and unit number (observable treatment) or only a code that does not reveal their identity (anonymous treatment). Observability tripled participation in the program.

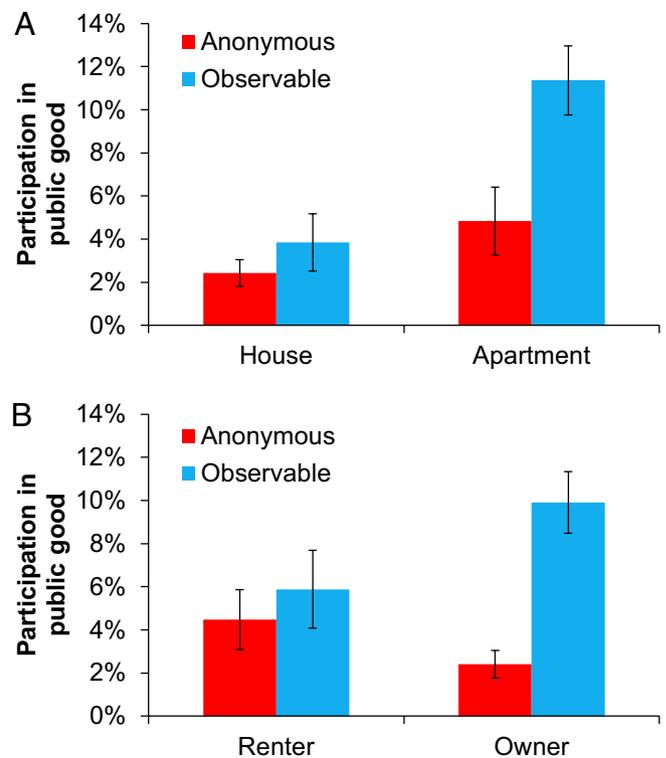


Fig. 3. Observability increased participation more in settings where reputational concerns matter more. (A) Observability increased participation more in apartment buildings where residents are more likely to interact with their neighbors in public spaces and sign-up sheets were typically posted in especially conspicuous locations, compared with row houses or individual homes, where neighbors are less likely to interact and sign-up sheets were less easily visible by others. (B) Similarly, observability increased participation more among those who own their homes/apartments relative to those who rent because renters are more transient and therefore likely to be less invested in long-term relationships with their neighbors.

anonymous = 0.024, observable = 0.038, not significant, $n = 826$; yielding an estimated interaction of 0.052, $P = 0.04$).

Second, we test whether observability had a larger effect among those who own their homes/apartments relative to those who rent. Renters are more transient and therefore likely to be less invested in relationships with their neighbors. Thus, indirect reciprocity theory predicts that observability will have a larger effect among owners. As shown in Fig. 3B, the results were again consistent with this prediction: observability dramatically increased participation among owners (fraction of residents participating: anonymous = 0.024, observable = 0.099, $P < 0.01$, $n = 1015$) but had little effect on renters (fraction of residents participating: anonymous = 0.045, observable = 0.059, not significant, $n = 393$; yielding an estimated interaction of 0.046, $P < 0.01$).

Residents of apartment buildings and individual homes differ on any number of dimensions, as do those who own versus rent their homes. Thus, although the results are consistent with the predictions of the indirect reciprocity framework, alternative explanations of these results are possible. To partially address this issue, we ran the same analysis with additional controls such as Spanish language preference, ethnicity, missed payments, rebate use, and historical electrical use, and found even stronger results (*Supporting Information*). Therefore, differences on these dimensions do not account for the differential effects of observability seen in Figs. 2 and 3.

Finally, we provide evidence that the effect of observability is unique to public goods. Not participating in the program should carry the threat of social sanctions only if participation is considered

to be a public good. Therefore, indirect reciprocity theory predicts that observability should not increase participation among subjects who do not think of participation as a public good. To test this prediction, an additional 1,005 subjects received exactly the same treatment as described above, except that the mailers they received were stripped of any language that framed the demand response program as a public good. Consistent with our hypothesis, Fig. 4 shows that the effect of observability was reduced in this cohort (fraction of residents participating: anonymous = 0.061, observable = 0.086, not significant, $n = 1005$; estimated interaction between observability and the public good message in a pooled regression is 0.035, $P = 0.098$).

Discussion

We have shown that indirect reciprocity promotes cooperation in a real-world public goods game affecting thousands of people. Making participation in the public good observable substantially increased sign-ups and did so significantly more than offering a cash incentive. Moreover, the effect of observability was larger in settings where individuals were more likely to have future interactions with those who observed them, and when participation was framed as a public good. These results provide evidence that reputational concerns were the driving force behind the effect of observability in our study.

Our study is part of a nascent literature exploring reputation and prosociality using field experiments. Consistent with our findings in the domain of energy efficiency, there is evidence that publicizing the names of donors increases the frequency of blood donation (70) as well as the level of giving to a college charity (71). Nonfinancial incentives involving reputation have also been shown to outperform monetary incentives in motivating the sale of condoms on behalf of a health organization in Namibia (72). Our work adds to these studies by directly manipulating observability, allowing a comparison with monetary incentives while avoiding other potential confounds present in previous experiments. We also test specific hypotheses generated by indirect reciprocity theory regarding when observability will and will not increase cooperation. Taken together, this body of work provides clear evidence that reputational incentives can be a powerful force for increasing cooperation in the field. Our paper in particular adds to efforts aimed at promoting energy conservation via nonfinancial incentives, such as providing people information about their own energy use and how it compares with the energy use of their neighbors (79–82).

A question arising from our study is the extent to which our subjects were conscious of their indirect reciprocity motives. One

possibility is that they explicitly considered the reputational costs of not participating in the observable public goods treatment. Alternatively, they may have learned or evolved sensitivity to subtle cues that subconsciously increased their desire to participate when their decisions were observable, as has been shown in other settings (62, 63, 65, 83). Perhaps the degree of “warm glow” they feel is sensitive to the degree of observability in their environment and the likelihood of interacting with observers in the future. Subsequent studies should further investigate this issue.

A related issue is the universality of reputation concerns. Observability can promote cooperation, but only in populations where the proper social norms are in place. For example, in a laboratory experiment in the United States, making public goods contributions observable by linking the PGG to a set of pairwise Prisoner’s Dilemma games led to high contributions (53). However, when the same experiment was run using students in Romania, no such positive effects were observed because the Romanians did not sanction bad behavior in the PGG (84). Similarly, providing feedback on how one’s energy use compares with one’s neighbors had reduced consumption among American liberals but may have had the opposite effect among conservatives (80). Studying the interaction between norms and institutional policies is an important direction for future research.

In our experiment, the observability mechanism was designed so that participation was automatically displayed to all: because sign-up sheets were posted in public areas, no special effort was required by individuals to spread reputational information. Most indirect reciprocity models, however, rely on individuals communicating information about the observed actions of others (21). Fortunately, we are more than happy to talk about how others have behaved: gossip is a central element of human communication (85, 86). However, why did we come to have this predilection for gossiping about the previous behavior of others? Why spend time and effort on evaluating others, and why give honest evaluations of competitors? Indirect reciprocity itself offers a potential answer: providing honest information or not is another game of cooperation and defection, which is also linked to reputation. Your reputation can be damaged not just by defection in the primary public goods game, but also by the distribution of incomplete or false information. Another important question involves large-scale reputation systems such as those used by the online market eBay (87) or the business rating website Yelp (88): to what extent does our intrinsic desire to gossip extend to these more distributed settings? Why do people bother to leave evaluations, and how can secondary reputation systems be designed to encourage honest feedback? Exploring these issues is an important direction for further study.

Indirect reciprocity offers a powerful tool for promoting cooperation in contexts of great societal importance. Here, we offer quantitative evidence for one example: curbing electricity use during periods of high demand. However, this is just one of many such opportunities (70–72, 89). For example, people might be induced to drive more efficient cars if all vehicles bore a visible indication of fuel efficiency, perhaps via mandated color coding of license plates for the most efficient and most wasteful vehicles. Or home energy use might be reduced if utility companies made individuals’ power use statistics publicly available. One might even apply this logic to scientific discovery: a measure of “scientific carbon efficiency” could be calculated by dividing an author’s number of citations (or h-index) by the number of miles flown to attend conferences. Of course, privacy is an important issue that must be balanced against the benefits of reputational pressure. However, there are also indirect reciprocity applications that do not infringe on the privacy rights of individuals. For example, businesses might reduce their environmental impact if they were required to disclose the overall carbon footprint of their operations. Reputational concerns might discourage financial institutions from taking excessive risk because of changes in

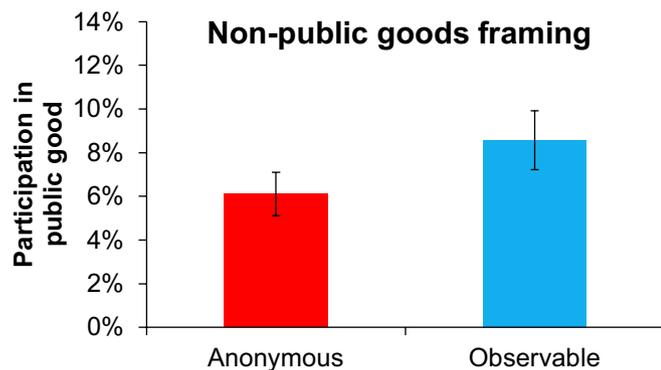


Fig. 4. We solicited an additional 1,005 customers with exactly the same treatment as described above, except that the informational materials they received were stripped of any language that framed the blackout prevention program as a public good. The effect of observability was dramatically reduced among subjects who did not receive the public good framing.

the price at which they have to borrow. Or government agencies might reduce waste if the relevant statistics were readily accessible to the public. Developing interventions that harness indirect reciprocity is a promising direction for future public policy initiatives.

Methods

We administered the field experiment with the collaboration of the Pacific Gas and Electric Company (PG&E), a regulated utility that operates the majority of Northern California's retail residential electricity market. The experiment was incorporated into a routine marketing effort for a demand response program called SmartAC, which is designed to help prevent or shorten power interruptions by curbing demand from central air conditioners on days with unusually high demand, or in the case of an unexpected plant or transmission failure. The program is voluntary; subjects who participate contribute to a public good by contributing to the stability of the electrical grid in all of California, at the cost of some personal inconvenience and possibly some discomfort. The SmartAC switch is installed free-of-charge. At the time of the program, participants received a \$25 check for signing up. The SmartAC program is a typical demand-side management, direct load control, or load shedding program.

Subjects in the field experiment were residential customers living in homeowners associations (HOAs) and one rental complex in Santa Clara County. We focused on tenants of HOAs because it was necessary to choose residences with public spaces where sign-up sheets could be posted. We focused on Santa Clara County because PG&E had not marketed in this area before the field experiment. Furthermore, Santa Clara County is hot enough that customers there were likely to have air conditioners, and dense enough to have a sufficient number of HOAs. Finally, we restricted the analysis to HOAs where all residents were known to have central air conditioning because central air conditioning was required to participate in the SmartAC program.

We invited subjects to participate in the program by sliding marketing materials under subjects' doors, placing them on their doorstep, or mailing

materials to subjects. The materials included an informational letter describing the program and an instruction card that directed subjects to sign up for the program on sign-up sheets posted next to their mailboxes or in another central location. We left the sign-up sheets up for 3 to 10 d, depending on managers' preferences, the weather, and other conditions. After distributing the marketing materials, we removed the sign-up sheets, noted subjects' participation decisions, and provided the list of participants to PG&E's contractor for processing and installation. Note that subjects were not aware that they were participating in an experiment. This study is therefore classified as a natural field experiment (90).

In the experiment's main treatment, we varied observability by varying the design of the sign-up sheets on which subjects register for the program: some sheets were designed so that subjects' identities were easily revealed to others who observed the sign-up sheet whereas others were designed to conceal subjects' identities. In the latter "anonymous" design, the fields for subjects' names and apartment numbers were omitted from the sign-up sheet. Instead, subjects were identified only by their randomly generated personal code.

Simultaneously, we varied the design of the marketing materials along two dimensions. First, we varied whether the materials framed the decision to sign up as a contribution to a public good that would benefit others, or just as a new feature being offered by PG&E. Second, we varied whether subjects were offered a \$25 incentive for signing up for the program. See [Supporting Information](#) for further details of the experimental design.

ACKNOWLEDGMENTS. This research was partially funded by National Science Foundation Proposal 0752818, Russell Sage Foundation Award 980803, the Stigler Center for the Study of the Economy and the State's Energy Initiative, the University of Chicago Graduate School of Business, and the John Templeton Foundation. Any opinions expressed in this article are those of the authors and not of the Federal Trade Commission or any individual Commissioner.

- Hardin G (1968) The tragedy of the commons. The population problem has no technical solution; it requires a fundamental extension in morality. *Science* 162(3859):1243–1248.
- Ostrom E (1990) *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge Univ Press, Cambridge, UK).
- Axelrod R (1984) *The Evolution of Cooperation* (Basic Books, New York).
- Nowak MA, Highfield R (2011) *SuperCooperators: Altruism, Evolution, and Why We Need Each Other to Succeed* (Free Press, New York).
- Churchland PS (2012) *Braintrust: What Neuroscience Tells Us About Morality* (Princeton Univ Press, Princeton).
- Oakley B, Knaflo A, Madhavan G, Wilson DS, eds (2011) *Pathological Altruism* (Oxford Univ Press, Oxford).
- Levin SA (2006) Learning to live in a global commons: Socioeconomic challenges for a sustainable environment. *Ecol Res* 21(3):328–333.
- Walker B, et al. (2009) Environment. Looming global-scale failures and missing institutions. *Science* 325(5946):1345–1346.
- Helbing D, Yu W (2009) The outbreak of cooperation among success-driven individuals under noisy conditions. *Proc Natl Acad Sci USA* 106(10):3680–3685.
- Sigmund K (2010) *The Calculus of Selfishness* (Princeton Univ Press, Princeton).
- Nowak MA (2006) Five rules for the evolution of cooperation. *Science* 314(5805):1560–1563.
- Nowak MA, Tarnita CE, Antal T (2010) Evolutionary dynamics in structured populations. *Philos Trans R Soc Lond B Biol Sci* 365(1537):19–30.
- Dasgupta P (2012) Dark matters: Exploitation as cooperation. *Journal of Theoretical Biology* 299:180–187.
- Choi JK, Bowles S (2007) The coevolution of parochial altruism and war. *Science* 318(5850):636–640.
- Carlton DW, Perloff JM (1994) *Modern Industrial Organization* (HarperCollins, New York).
- Cela-Conde CJ, Ayala FJ (2007) *Human Evolution: Trails from the Past* (Oxford University Press, Oxford).
- Fu F, Hauert C, Nowak MA, Wang L (2008) Reputation-based partner choice promotes cooperation in social networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 78(2 Pt 2):026117.
- Manapat ML, Nowak MA, Rand DG (2012) Information, irrationality and the evolution of trust. *J Econ Behav Organ*, in press.
- Nowak MA, Page KM, Sigmund K (2000) Fairness versus reason in the ultimatum game. *Science* 289(5485):1773–1775.
- Nowak MA, Sigmund K (1998) Evolution of indirect reciprocity by image scoring. *Nature* 393(6685):573–577.
- Nowak MA, Sigmund K (2005) Evolution of indirect reciprocity. *Nature* 437(7063):1291–1298.
- Ohtsuki H, Iwasa Y, Nowak MA (2009) Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. *Nature* 457(7225):79–82.
- Masuda N, Ohtsuki H (2007) Tag-based indirect reciprocity by incomplete social information. *Proc Biol Sci* 274(1610):689–695.
- Ohtsuki H, Iwasa Y (2006) The leading eight: Social norms that can maintain cooperation by indirect reciprocity. *J Theor Biol* 239(4):435–444.
- Panchanathan K, Boyd R (2004) Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature* 432(7016):499–502.
- Kandori M (1992) Social norms and community enforcement. *Rev Econ Stud* 59:63–80.
- Brandt H, Sigmund K (2006) The good, the bad and the discriminator—errors in direct and indirect reciprocity. *J Theor Biol* 239(2):183–194.
- Pacheco JM, Santos FC, Chalub FACC (2006) Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity. *PLoS Comput Biol* 2(12):e178.
- Ohtsuki H, Iwasa Y (2007) Global analyses of evolutionary dynamics and exhaustive search for social norms that maintain cooperation by reputation. *J Theor Biol* 244(3):518–531.
- Leimar O, Hammerstein P (2001) Evolution of cooperation through indirect reciprocity. *Proc Biol Sci* 268(1468):745–753.
- Panchanathan K, Boyd R (2003) A tale of two defectors: The importance of standing for evolution of indirect reciprocity. *J Theor Biol* 224(1):115–126.
- Nowak MA, Sigmund K (1998) The dynamics of indirect reciprocity. *J Theor Biol* 194(4):561–574.
- Suzuki S, Akiyama E (2007) Evolution of indirect reciprocity in groups of various sizes and comparison with direct reciprocity. *J Theor Biol* 245(3):539–552.
- Ohtsuki H, Iwasa Y (2004) How should we define goodness?—reputation dynamics in indirect reciprocity. *J Theor Biol* 231(1):107–120.
- Manapat ML & Rand DG (2012) Delayed and inconsistent information and the evolution of trust. *Dyn Games Appl* 2:401–410.
- Uchida S, Sigmund K (2010) The competition of assessment rules for indirect reciprocity. *J Theor Biol* 263(1):13–19.
- Nakamura M, Masuda N (2011) Indirect reciprocity under incomplete observation. *PLoS Comput Biol* 7(7):e1002113.
- Suzuki S, Akiyama E (2007) Three-person game facilitates indirect reciprocity under image scoring. *J Theor Biol* 249(1):93–100.
- Berger U (2011) Learning to cooperate via indirect reciprocity. *Games Econ Behav* 72(1):30–37.
- Mani A, Rahwan I, Pentland A (2013) Inducing Peer Pressure to Promote Cooperation. *Scientific Reports* 3:1735.
- Seinen I, Schram A (2006) Social status and group norms: Indirect reciprocity in a repeated helping experiment. *Eur Econ Rev* 50(3):581–602.
- Ule A, Schram A, Riedl A, Cason TN (2009) Indirect punishment and generosity toward strangers. *Science* 326(5960):1701–1704.
- Milinski M, Semmann D, Bakker TCM, Krambeck HJ (2001) Cooperation through indirect reciprocity: Image scoring or standing strategy? *Proc Biol Sci* 268(1484):2495–2501.
- Milinski M, Semmann D, Krambeck HJ (2002) Donors to charity gain in both indirect reciprocity and political reputation. *Proc Biol Sci* 269(1494):881–883.
- Milinski M, Semmann D, Krambeck HJ (2002) Reputation helps solve the 'tragedy of the commons'. *Nature* 415(6870):424–426.

46. Rockenbach B, Milinski M (2006) The efficient interaction of indirect reciprocity and costly punishment. *Nature* 444(7120):718–723.
47. Semmann D, Krambeck H-J, Milinski M (2005) Reputation is valuable within and outside one's own social group. *Behav Ecol Sociobiol* 57(6):611–616.
48. Wedekind C, Milinski M (2000) Cooperation through image scoring in humans. *Science* 288(5467):850–852.
49. Pfeiffer T, Tran L, Krumme C, Rand DG (2012) The value of reputation. *J R Soc Interface* 9(76):2791–2797.
50. Wedekind C, Braithwaite VA (2002) The long-term benefits of human generosity in indirect reciprocity. *Curr Biol* 12(12):1012–1015.
51. Bolton GE, Katok E, Ockenfels A (2005) Cooperation among strangers with limited information about reputation. *J Public Econ* 89(8):1457–1468.
52. Jacquet J, Hauert C, Traulsen A, Milinski M (2011) Shame and honour drive cooperation. *Biol Lett* 7(6):899–901.
53. Rand DG, Dreber A, Ellingsen T, Fudenberg D, Nowak MA (2009) Positive interactions promote public cooperation. *Science* 325(5945):1272–1275.
54. Choi J-K, Ahn TK (2013) Strategic reward and altruistic punishment support cooperation in a public goods game experiment. *J Econ Psychol*, 35:17–30.
55. Rege M, Telle K (2004) The impact of social approval and framing on cooperation in public good situations. *J Public Econ* 88(7):1625–1644.
56. Andreoni J, Petrie R (2004) Public goods experiments without confidentiality: A glimpse into fund-raising. *J Public Econ* 88(7):1605–1623.
57. Tadelis S (2007) The power of shame and the rationality of trust. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1006169. Accessed April 20, 2013.
58. Andreoni J, Bernheim BD (2009) Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica* 77(5):1607–1636.
59. Linardi S, McConnell MA (2008) Volunteering and Image Concerns. *California Institute of Technology Working Paper*.
60. List JA, Berrens RP, Bohara AK, Kerkvliet J (2004) Examining the role of social isolation on stated preferences. *Am Econ Rev* 94(3):741–752.
61. Kossmeier S, Ariely D, Bracha A (2009) Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *Am Econ Rev* 99(1):544–555.
62. Haley KJ, Fessler DMT (2005) Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evol Hum Behav* 26:245–256.
63. Burnham T, Hare B (2007) Engineering Human Cooperation. *Hum Nat* 18(2):88–108.
64. Powell KL, Roberts G, Nettle D (2012) Eye Images Increase Charitable Donations: Evidence From an Opportunistic Field Experiment in a Supermarket. *Ethology* 118(11):1096–1101.
65. Ernest-Jones M, Nettle D, Bateson M (2011) Effects of eye images on everyday cooperative behavior: a field experiment. *Evol Hum Behav* 32(3):172–178.
66. Rand DG, Greene JD, Nowak MA (2012) Spontaneous giving and calculated greed. *Nature* 489(7416):427–430.
67. Schulz JF, Fischbacher U, Thöni C, & Utikal V (In press) Affect and fairness: Dictator games under cognitive load. *Journal of Economic Psychology* (0).
68. Rand DG, et al. (2013) Intuitive Cooperation and the Social Heuristics Hypothesis: Evidence from 15 Time Constraint Studies. Available at SSRN: <http://ssrn.com/abstract=2222683>.
69. Cornelissen G, Dewitte S, Warlop L (2011) Are social value orientations expressed automatically? Decision making in the dictator game. *Pers Soc Psychol Bull* 37(8):1080–1090.
70. Lacetera N, Macis M (2010) Social image concerns and prosocial behavior: Field evidence from a nonlinear incentive scheme. *J Econ Behav Organ* 76(2):225–237.
71. Karlan D, McConnell MA (2012) *Hey Look at Me: The Effect of Giving Circles on Giving* (National Bureau of Economic Research, Cambridge, MA).
72. Ashraf N, Bandiera O, Jack K (2012) No margin, no mission? A field experiment on incentives for pro-social tasks. *Harvard Business School Working Papers*. Available at www.hbs.edu/faculty/Pages/download.aspx?name=12-008.pdf. Accessed April 20, 2013.
73. Levitt SD, List JA (2007) What do laboratory experiments measuring social preferences reveal about the real world? *J Econ Perspect* 21:153–174.
74. US Committee on Technology (2011) *A Policy Framework for the 21st Century Grid: Enabling Our Secure Energy Future* (Executive Office of the President: National Science and Technology Council, Washington).
75. Chupka MW, Earle R, Fox-Penner P, Hledik R (2008) *Transforming America's Power Industry: The Investment Challenge 2010–2030* (The Edison Foundation, Washington).
76. Kathan D, et al. (2011) *Assessment of Demand Response and Advanced Metering: Staff Report* (Federal Energy Regulatory Commission, Washington).
77. Gneezy U, Rustichini A (2000) Pay Enough or Don't Pay At All. *Q J Econ* 115(3):791–810.
78. Bertrand M, Karlan D, Mullainathan S, Shafir E, Zinman J (2010) What's advertising content worth? Evidence from a consumer credit marketing field experiment. *Q J Econ* 125(1):263–306.
79. Ayres I, Raseman S, Shih A (2009) *Evidence from Two Large Field Experiments That Peer Comparison Feedback Can Reduce Residential Energy Usage* (National Bureau of Economic Research, Cambridge, MA).
80. Costa DL, Kahn ME (2010) *Energy Conservation "Nudges" and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment* (National Bureau of Economic Research, Cambridge, MA).
81. Thompson C (2007) Clive Thompson thinks: Desktop orb could reform energy hogs. *Wired Magazine* 15(8).
82. Minosi A, et al. (2003) Intelligent, low-power and low-cost measurement system for energy consumption. *IEEE International Symposium on Virtual Environments, Human-Computer Interfaces and Measurement Systems (IEEE, Washington)*, pp 125–130.
83. Bateson M, Nettle D, Roberts G (2006) Cues of being watched enhance cooperation in a real-world setting. *Biol Lett* 2(3):412–414.
84. Ellingsen T, Herrmann B, Nowak MA, Rand DG, Tarnita CE (2012) Civic Capital in Two Cultures: The Nature of Cooperation in Romania and USA. Available at <http://ssrn.com/abstract=2179575>. Accessed April 20, 2013.
85. Dunbar RIM, Marriott A, Duncan NDC (1997) Human conversational behavior. *Hum Nat* 8(3):231–246.
86. Sommerfeld RD, Krambeck H-J, Semmann D, Milinski M (2007) Gossip as an alternative for direct observation in games of indirect reciprocity. *Proc Natl Acad Sci USA* 104(44):17435–17440.
87. Resnick P, Zeckhauser R, Swanson J, Lockwood K (2006) The value of reputation on eBay: A controlled experiment. *Exp Econ* 9(2):79–101.
88. Luca M (2011) Reviews, reputation, and revenue: The case of Yelp.com. Available at <http://hbswk.hbs.edu/item/6833.html>. Accessed April 20, 2013.
89. Rand DG, Nowak MA (2009) Name and Shame. *New Sci* 204(2734):28–29.
90. Harrison GW, List JA (2004) Field experiments. *J Econ Lit* 42(4):1009–1055.

Supporting Information

Yoeli et al. 10.1073/pnas.1301210110

SI Text

Experimental Methods

Setting. We administered the field experiment with the collaboration of Pacific Gas and Electric Company (PG&E), a regulated utility that operates the majority of Northern California's retail residential electricity market. The experiment was incorporated into a routine marketing effort for a program called SmartAC, which helps to prevent or shorten power interruptions by curbing demand from central air conditioners on days with unusually high demand, or in the case of an unexpected plant or transmission failure. The program is voluntary; subjects who participate contribute to a public good by contributing to the stability of the electrical grid in all of California, at the cost of some personal inconvenience and possibly some discomfort. PG&E installs a small radio-operated switch in series with participants' central air conditioners' thermostats. During system emergencies, PG&E remotely cycles the switch on-and-off, reducing the frequency with which the central air conditioner's compressor operates. When engaged, the switches raise the temperature in participants' homes by no more than four degrees Fahrenheit, while providing a substantial aggregate reduction in demand proportional to the number of participants. The SmartAC switch is installed free-of-charge. At the time of the program, participants traditionally received a \$25 check for signing up.

The SmartAC program is a typical demand-side management (DSM), direct load control, or load shedding program. Broadly speaking, DSM programs are energy efficiency programs in which end users grant a utility or system operator the right to remotely curb a designated part of their energy consumption when costs are exorbitant or there is a risk of a blackout.

Sample. Subjects in the field experiment were residential customers living in homeowners associations (HOAs) and one rental complex in Santa Clara County. We restricted the sample to tenants of HOAs because it was necessary to choose residences with public spaces where sign-up sheets could be posted. We restricted the sample to Santa Clara County because PG&E had not marketed in this area before the field experiment. Furthermore, Santa Clara County is hot enough that customers there were likely to have air conditioners, and dense enough to have a sufficient number of HOAs. Finally, we restricted the analysis to HOAs where all residents were known to have central air conditioning because central air conditioning was required to participate in the SmartAC program.

Procedure. We began by obtaining permission from managers and boards to solicit residents in their HOA for the program. If either did not approve, the HOA was removed from the sample. If we obtained approval, we invited subjects to participate in the program by sliding marketing materials under subjects' doors or placing them on their doorstep. When managers did not permit this, we mailed materials to subjects. The materials included an informational letter describing the program and an instruction card that directed subjects to sign up for the program on sign-up sheets posted next to their mailboxes or in another central location. Both the letter and the instruction card were stuffed into a SmartAC envelope. We left the sign-up sheets up for 3–10 d, depending on managers' preferences, the weather, and other conditions. After distributing the marketing materials, we removed the sign-up sheets, noted subjects' participation decisions, and provided the list of participants to PG&E's contractor

for processing and installation. Note that subjects were not aware that they were participating in an experiment. This study is therefore classified as a natural field experiment (1).

Individual subjects were dropped from the analysis when we could identify them as having had extensive interaction with us (the researchers), for example if they attended a board meeting where we described the program in full. Entire HOAs were dropped from the analysis if circumstances prevented subjects from signing up, e.g., the sign-up sheets were repeatedly removed or vandalized (one HOA was dropped for this reason), or if HOAs were located in areas that had been mailed by PG&E before the experiment (one HOA was dropped for this reason). The following describes the variables used in the analysis.

Treatments. In the experiment's main treatment, we varied observability by varying the design of the sign-up sheets on which subjects register for the program: some sheets were designed so that subjects' identities were easily revealed to others who observed the sign-up sheet whereas others were designed to conceal subjects' identities.

Simultaneously, we varied the design of the marketing materials along two dimensions. First, we varied whether the materials framed the decision to sign up as a contribution to a public good. Second, we varied whether subjects were offered a \$25 incentive for signing up for the program. There were also three additional treatments which are presented for completeness.

Varying observability. We varied observability by manipulating the design of the sign-up sheets on which subjects signed up for the SmartAC program. We designed the treatment group's sign-up sheets to reveal subjects' identities to others who observed the sign-up sheet. This "observable" design included fields for subjects' names and apartment numbers. Also, when we generated subjects' personal codes, which uniquely identified them within their building, we did so using their apartment numbers. We designed the control group's sign-up sheets to conceal subjects' identities from others who observed the sign-up sheet. This "anonymous" design omitted the fields for subjects' names and apartment numbers. Subjects were identified only by their personal code, which we generated randomly. Reproductions of the two versions of the sign-up sheets are displayed in Fig. S2. We randomized subjects into treatment groups at the sign-up sheet level: neighbors who shared a mailbox location were necessarily randomized into treatment groups together.

This treatment is similar to existing experimental tests of social approval, all of which test for social approval by comparing individuals' contributions to a public good under varying degrees of peer or experimenter scrutiny.

Varying public good message. We manipulated whether subjects considered their sign-up decision in the context of a contribution to a public good, namely, reducing the risk of power interruptions in California. Subjects in the treatment group were informed of the SmartAC's role in preventing power interruptions in the printed marketing materials that they receive. The materials' design made the difference between the treatment groups striking: in one, the public good message is greatly emphasized and is accompanied with a butterfly graphic, whereas in the other, both the image and the message are missing. We used this distinction in the informational flier (displayed in Fig. S3), as well as in the instruction card and the envelope in which both were enclosed. This treatment is similar to those used by laboratory studies which also vary framing simultaneously with observability (2,3).

Varying financial incentive. We manipulated whether subjects were offered a \$25 incentive for signing up for the SmartAC program, again by varying the content of the printed marketing materials. We followed common practice (4) and used this treatment as a benchmark against which other experimental effects were compared.

Additional treatments. The experiment included three additional treatments. First, we informed one group that a relatively small number of subjects had been offered the \$25 sign up incentive and another that a larger number of subjects had been offered the incentive.

Second, we attempted to exogenously vary sign-up rates of subjects' peers by varying whether residents in a sign-up sheet group were eligible for the sign-up incentive. In some of the groups, subjects were offered the incentive with probability 0.8, whereas in other groups, they were never offered the incentive. Although we randomized subjects into treatment groups at the sign-up sheet level, some subjects in eligible groups did not receive the incentive because marketing materials were still randomized at the individual level as described above.

Finally, we varied whether subjects had the opportunity to observe their peers' decision before signing up for the SmartAC. We did this by delaying delivery of the marketing materials for a small group of residents in a building and observing whether their response to the mailer in the first few days was different from the remaining residents' response in the first few days. We randomized the timing of mailer delivery at the individual level.

These treatments were designed with the expectation that the \$25 incentive would have a large impact on sign-up rates. Although this expectation was based on information from industry experts, it did not bear fruit in this setting. Consequently, we repressed the bulk of the tests associated with these treatments, and the associated discussion. For completeness, we present the results of a regression with controls for these treatments in column 3 of Table S1.

Randomization Procedures. The following describes the randomization procedures for all treatments.

Sign-up sheets. Subjects were always assigned to a revealing sign-up sheet with probability 0.5. This assignment was independent of the remaining treatments.

Incentives. Subjects' sign-up sheet group was designated as eligible for incentives with probability 0.8. Conditional on being eligible for the incentive, subjects received the incentive with probability 0.8. The unconditional probability of receiving an incentive was thus 0.64.

Public good message. Subjects were assigned the public good message with probability 0.45 if their sign-up sheet group was eligible for incentives and 0.87 if it was not. The unconditional probability of receiving the public good message was thus 0.49.

Crowding out message. Conditional on receiving the public good message and the incentive, subjects received the high and low crowding out message each with probability 0.42.

Late mailers. Subjects were assigned to receive their mailer late with probability 0.2. This assignment was independent of the remaining treatments.

Additional Data. The experimental treatments were augmented with individual characteristics collected through visual inspection or from PG&E's customer files. They included the following: whether subjects lived in an apartment building, whether subjects were homeowners, subjects' ethnicity inferred from their names, subjects' Spanish language preference, whether subjects had missed payments, whether subjects had ever requested rebates, and subjects' electrical use history in kilowatt hours.

Statistical Methods and Regression Results

In the manuscript, we present sample means across treatments accompanied by P values, which indicate the likelihood that the sample came from a population in which there was no difference across treatments. The P values presented are from probit regressions of whether subjects signed up for the SmartAC program on experimental manipulations and subject characteristics. We use variants of the following regression specification:

$$s_{ijk} = \begin{cases} 1 & \text{if } \alpha + \beta_0 o_{ij} + \beta_1 g_{ijk} + \beta_2 o_{ij} g_{ijk} + \beta_3 t_{ijk} + \beta_4 Z_{ijk} > \epsilon_{ijk} \\ 0 & \text{otherwise} \end{cases} \quad [1]$$

where s_{ijk} indicates whether subject k of sign-up sheet group j in HOA i participated in the program, o_{ij} is a dummy indicating whether her sign-up sheet was observable, g_{ijk} indicates whether she received the public goods message, t_{ijk} indicates whether k received a \$25 sign-up incentive, Z_{ijk} are demographics controls (see *Additional Data*) and interactions included in some specifications, and ϵ_{ijk} is a normally distributed idiosyncratic error term that is assumed to be clustered at the HOA level. For example, in describing our main result, we compare the mean sign-up rate for subjects assigned to the observable treatment to the mean sign-up rate for subjects assigned to the anonymous treatment and report the P value associated with the average marginal effect of observability from our estimates of the regression in Eq. 1 (see column 1 in Table S1).

We present average marginal effects, SEs, and P values from regressions that include all subjects for whom data are available, evaluated at parameters of interest. For example, P values for subjects who received the public good framing are from regressions on all subjects, but for the average marginal effect of observability evaluated at public good = 1. Similarly, in describing the different effect of observability across groups, we present the estimated average marginal effects of observability from a single regression that includes subjects in both groups, evaluated separately for each group. For example, in comparing the effect of observability for homeowners and renters, we present the estimated average marginal effect of observability from a single regression that includes both homeowners and renters, evaluated separately at homeowner = 0 and homeowner = 1. We present the P value for the coefficient on the interaction of observability with a dummy variable indicating whether subjects were residents of apartment buildings and homeowners (see columns 4 and 5 in Table S1). In Table S1, we present both the number of subjects in the regression, as well as the number of subjects in the group of interest. For example, there were 582 subjects who lived in apartment buildings and received the public good framing; however, we present results for these subjects from regressions that include all 2,413 subjects (see column 2 in Table S1). For each result presented, we have confirmed that results from separate regressions (e.g., separate regressions for subjects who received the public good framing and those that did not) yield similar results, as expected with balanced treatment groups.

In all, there were 2,413 subjects in 15 HOAs. The average sign-up rate was 6.2%. Table S1 presents the probit regression estimates for all regressions referenced in the manuscript. Column 1 presents our main test of the effect of observability when subjects receive the public good framing. Column 2 adds controls for additional treatments not included in this study. Column 3 examines whether the effect of observability is stronger for subjects who live in apartment buildings. Column 4 does the same, but adds demographic controls. Column 5 examines whether the effect of observability is stronger for homeowners. Column 6 does the same, but adds demographic controls. Finally, column 7 presents the effect of observability when subjects do not receive the public good framing.

- Harrison GW, List JA (2004) Field experiments. *J Econ Lit* 42:1009–1055.
- Rege M, Telle K (2004) The impact of social approval and framing on cooperation in public good situations. *J Public Econ* 88:1625–1644.
- Ariely D, Bracha A, Meier S (2007) *Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially* (Federal Reserve Bank of Boston, Boston).
- Bertrand M, Karlan D, Mullainathan S, Shafir E, Zinman J (2010) What's advertising content worth? Evidence from a consumer credit marketing field experiment. *Q J Econ* 125:263–306.



Fig. S1. Example of posted sign-up sheet. Subjects were instructed to sign up for the program on sign-up sheets posted in public spaces, usually near mailboxes.

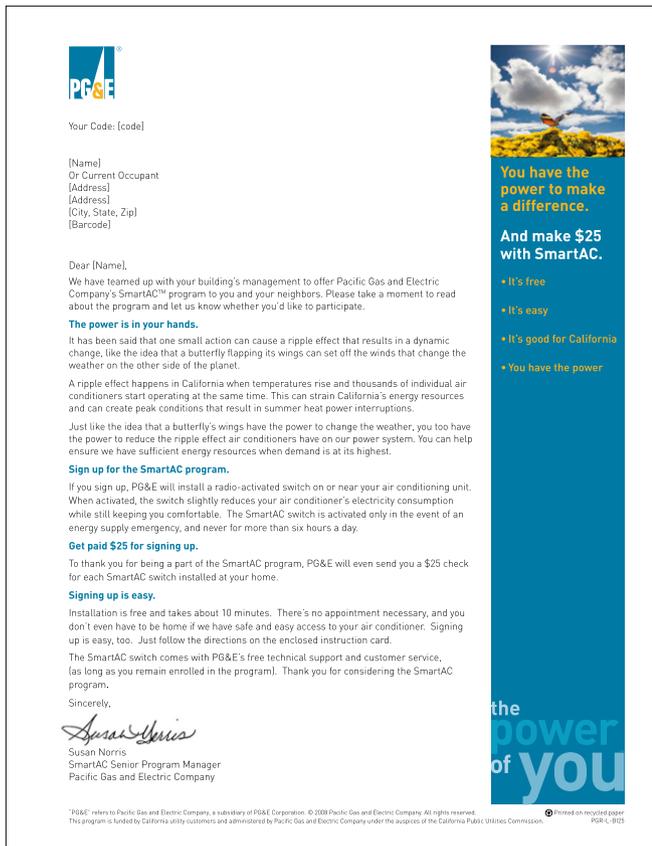
Pacific Gas and Electric Company		SmartAC [®] SIGN-UP SHEET			
Sign Me Up!	Personal Code (From your letter)	Date (MM/DD)	Apt. # (Please Print)	First Name (Please Print)	Last Name (Please Print)
<input type="checkbox"/> Yes	_____	____/____/____	_____	_____	_____
<input type="checkbox"/> No	_____	____/____/____	_____	_____	_____
<input type="checkbox"/> Yes	_____	____/____/____	_____	_____	_____
<input type="checkbox"/> No	_____	____/____/____	_____	_____	_____
<input type="checkbox"/> Yes	_____	____/____/____	_____	_____	_____
<input type="checkbox"/> No	_____	____/____/____	_____	_____	_____

(a) Observable

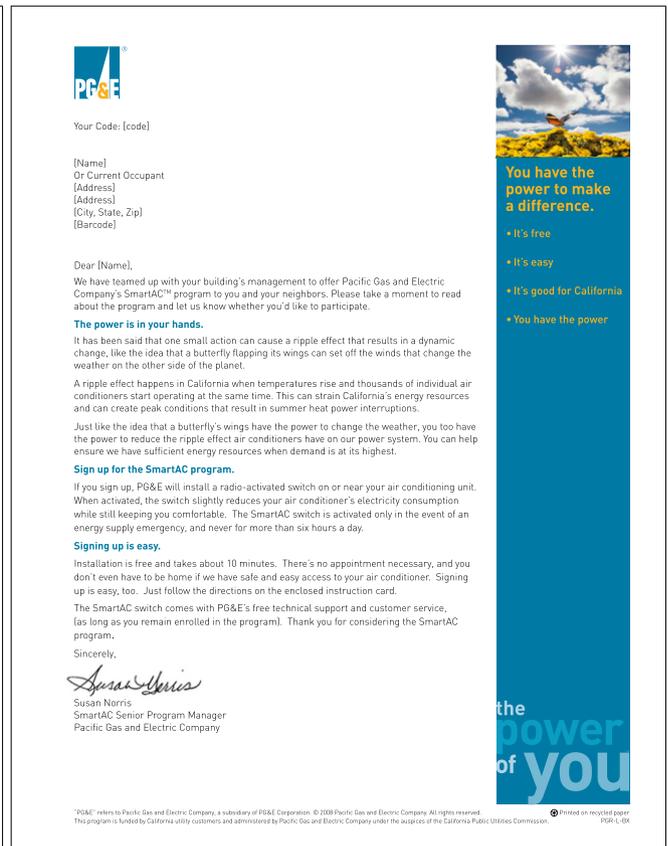
Pacific Gas and Electric Company		SmartAC [®] SIGN-UP SHEET			
Sign Me Up!	Personal Code (From your letter)	Date (MM/DD)	[Blacked out area]		
<input type="checkbox"/> Yes	_____	____/____/____	[Blacked out area]		
<input type="checkbox"/> No	_____	____/____/____	[Blacked out area]		
<input type="checkbox"/> Yes	_____	____/____/____	[Blacked out area]		
<input type="checkbox"/> No	_____	____/____/____	[Blacked out area]		
<input type="checkbox"/> Yes	_____	____/____/____	[Blacked out area]		
<input type="checkbox"/> No	_____	____/____/____	[Blacked out area]		

(b) Anonymous

Fig. S2. Varying observability using sign-up sheets. In the observable treatment, subjects were instructed to sign up by writing in their name and apartment number on the sheet. In the anonymous treatment, subjects were instructed to sign up using an anonymous code.



(e) Public Good and Monetary Incentive



(f) Public Good Only

Fig. S4. Varying monetary incentive using marketing materials. Participants in the program traditionally received \$25 for signing up. We varied whether participants were offered this \$25; subjects in the monetary incentive treatment were offered \$25 for participating whereas subjects in the control treatment were not.

Table S1. Effect of observability on participation in PG&E's SmartAC program

Variable	Public good framing						No public good framing
	1	2	3	4	5	6	7
Observable	0.059 (0.019)***	0.057 (0.019)***	0.036 (0.012)***	0.048 (0.011)***	0.058 (0.02)***	0.069 (0.018)***	0.024 (0.025)
Incentive	0.009 (0.014)	0.007 (0.025)	0.01 (0.012)	0.013 (0.013)	0.008 (0.014)	0.011 (0.014)	0.0103 (0.017)
Apartment building	—	—	0.048 (0.014)***	0.04 (0.011)***	—	—	—
Homeowner	—	—	—	—	0.006 (0.011)	0.002 (0.011)	—
Observable × apartment building	—	—	0.052 (0.026)**	0.067 (0.019)***	—	—	—
Observable × homeowner	—	—	—	—	0.046 (0.016)**	0.057 (0.022)**	—
Controls	—	OT	—	SD	—	SD	—
Observations in group	1,408	1,408	582	478	1,015	855	1,005
Observations in regression	2,413	2,413	2,413	2,073	2,413	2,073	2,413

Probit estimation of the effect of randomized interventions and subject characteristics on participation in the SmartAC program. Dependent variable is whether a subject signed up for the SmartAC program. The mean of the dependent variable is 0.062. Coefficients displayed are average marginal effects. Columns 1–6 display average marginal effects evaluated at public good = 0. Column 7 displays average marginal effects evaluated at public good = 1. Column 2 includes controls for treatments other than experimental treatments (OT, other treatments). Columns 4 and 6 include controls for Spanish language preference, rebate requests, missed payments, electricity use, and ethnicity (SD, subject demographics). SEs are displayed in parentheses and are clustered at the HOA level. —, omitted from the regression. **Statistical significance at the 5% level; ***statistical significance at the 1% level.