

Dynamic social networks promote cooperation in experiments with humans

David G. Rand^{a,b,1}, Samuel Arbesman^{c,d,1}, and Nicholas A. Christakis^{c,d,e,f,2}

^aProgram for Evolutionary Dynamics, ^bDepartment of Psychology, ^cInstitute for Quantitative Social Science, and ^eDepartment of Sociology, Harvard University, Cambridge, MA 02138; and ^dDepartment of Health Care Policy and ^fDepartment of Medicine, Harvard Medical School, Boston, MA 02115

Edited by Douglas S. Massey, Princeton University, Princeton, NJ, and approved October 18, 2011 (received for review May 23, 2011)

Human populations are both highly cooperative and highly organized. Human interactions are not random but rather are structured in social networks. Importantly, ties in these networks often are dynamic, changing in response to the behavior of one's social partners. This dynamic structure permits an important form of conditional action that has been explored theoretically but has received little empirical attention: People can respond to the cooperation and defection of those around them by making or breaking network links. Here, we present experimental evidence of the power of using strategic link formation and dissolution, and the network modification it entails, to stabilize cooperation in sizable groups. Our experiments explore large-scale cooperation, where subjects' cooperative actions are equally beneficial to all those with whom they interact. Consistent with previous research, we find that cooperation decays over time when social networks are shuffled randomly every round or are fixed across all rounds. We also find that, when networks are dynamic but are updated only infrequently, cooperation again fails. However, when subjects can update their network connections frequently, we see a qualitatively different outcome: Cooperation is maintained at a high level through network rewiring. Subjects preferentially break links with defectors and form new links with cooperators, creating an incentive to cooperate and leading to substantial changes in network structure. Our experiments confirm the predictions of a set of evolutionary game theoretic models and demonstrate the important role that dynamic social networks can play in supporting large-scale human cooperation.

collective action | economic games | evolutionary game theory | homophily | reciprocity

Cooperation is central to the success of human societies and is widespread (1–5). However, cooperation poses a challenge in both the social and biological sciences: How can this high level of cooperation be maintained in the face of possible exploitation? One answer involves networked interactions and population structure. Accounting for the fact that individuals are embedded in a social network and interact only with others in their neighborhood can lead natural selection to support even unconditional cooperation in evolutionary game theoretic models (6–11). The reason is that these local, nonrandom interactions can lead to the clustering of strategy types, so that cooperators are more likely to interact with other cooperators and therefore earn higher payoffs. However, empirical investigations using behavioral experiments have found little effect of network structure on promoting cooperation (12–15), despite evidence that cooperation and defection (as well as other, related behaviors) can spread among experimental subjects (15–17).

A key element missing from most prior network experiments is that real social networks typically are dynamic (18, 19). People often have control over whom they interact with, and interaction patterns change over time. This possibility of rewiring ties fundamentally alters the role of the network: Dynamic networks not only afford the opportunity for the clustering of strategy types but also make it possible for population structure to vary in response to cooperation. This variability creates a new form of

conditional action, one that occurs via changes in network structure rather than via changes in cooperation behavior.

Behavioral reciprocity is a central mechanism for the evolution of cooperation (1, 20, 21). In evolutionary game theory, reciprocity is defined as occurring when my actions toward you depend on your actions in the past. Reciprocity traditionally has been conceptualized in two-player game theory as the emergence of concordant behaviors within dyads. For example, the “tit-for-tat” strategy engages in reciprocity by cooperating only if the opponent cooperated in the previous round. Reciprocity creates future consequences for one's choices and has been shown experimentally to promote cooperation in repeated two-player interactions (22–25). However, reciprocity is problematic in group interactions involving more than two players: If the only way to sanction defectors is to defect, this action also harms the other cooperators in one's group (26).

Strategic tie formation and dissolution in dynamic networks offer a solution to this problem by providing players with an additional method of responding to the past actions of others. Players can reciprocate not only by changing their cooperation behaviors but also by creating or dissolving ties. Thus, cooperators need not switch to defection to punish defectors in their group; instead they can establish and maintain links with cooperators but sever connections with defectors, engaging in what we call “link reciprocity.” (Note that this reciprocity is different from the use of the term in social network analysis, where reciprocity refers to the existence of tie concordance in directed graphs—that is, if ego nominates alter, alter also nominates ego, and a mutually reciprocated tie exists.)

In recent years, a number of evolutionary game theory models have demonstrated the ability of link reciprocity to promote the evolution of cooperation in group interactions (27–32). Although these articles differ in the details of their methods and assumptions, several qualitative results emerge consistently across dynamic network models (see ref. 33 for a review). Most importantly, these models predict that rapid rewiring of the network supports cooperation. If the network updates too slowly, the threat of severed links cannot be carried out often enough to make defection maladaptive. In addition, several other predictions arise regularly across models: Rapidly updating networks are predicted to have more variation across individuals in the number of connections (i.e., a greater degree heterogeneity) than static or slowly updating networks; connections between two cooperators are predicted to be more stable than connections involving defectors in rapidly updating networks; and

Author contributions: D.G.R., S.A., and N.A.C. designed research; D.G.R. and S.A. performed research; D.G.R., S.A., and N.A.C. analyzed data; and D.G.R., S.A., and N.A.C. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹D.G.R. and S.A. contributed equally to this work.

²To whom correspondence should be addressed: E-mail: christak@hcp.med.harvard.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1108243108/-DCSupplemental.

cooperators are predicted to acquire more connections than defectors in rapidly updating networks.

Despite the considerable body of theoretical work exploring the ability of dynamic networks to promote multilateral cooperation, this issue thus far has received relatively little attention empirically. Here we evaluate these theoretical predictions by leveraging tools for running economic games online (12, 34–36) to conduct a series of large-scale behavioral cooperation experiments using dynamic networks. We randomly assigned 785 participants to one of four conditions in a series of 40 realizations of our network experiments (average network size = 19.6; SD, 6.4). In all conditions, subjects play a repeated cooperative dilemma with other subjects in an artificial social network created in the virtual laboratory.

We aim to capture the essential elements of the family of evolutionary game theory models exploring dynamic networks, and to do so using the simplest possible experimental design. To that end, our subjects lie on a nonweighted graph, and each subject interacts with her neighbors (determined as described below). As in most theoretical models, every subject chooses a single action simultaneously toward all neighbors, either cooperation (C) or defection (D). In our experiment, cooperation entails paying 50 units for each neighbor and results in each neighbor gaining 100 units; defection involves paying no costs and generating no benefits. Before making each decision, subjects are reminded of their number of neighbors and the neighbors' previous decisions. At the end of each turn, subjects are informed about the decisions of their neighbors, along with their own payoff. Subjects also are informed that, after every round, the probability that another round will occur is 0.8.

At the beginning of the experiment, the social network is initialized with 20% of possible links being formed at random. We examine three kinds of network conditions: random link updating, fixed links, and strategic link updating. In the random-link condition, the social network is regenerated randomly after every round, creating a well-mixed population. In the fixed-link condition, the network is static and remains in its initial conformation for the duration of the experiment.

In the strategic link updating conditions, each cooperation round is followed by a rewiring round in which subjects choose whether to alter their network connections. Previous theoretical models have investigated a wide range of strategic updating rules. In our experiments, we aim to implement the simplest rewiring process: In each round, a percentage k of subject pairs are picked at random to have their connections updated. If a connection already exists between the pair of subjects, one of the two (picked at random) is offered the chance to break the connection. If no connection already exists, one of the two (again picked at random) is offered the chance to form a new connection. In both cases, before choosing to break or form a connection, the deciding subject is informed of the other's action in the preceding round. We inform subjects about the previous play of potential new partners to investigate whether different conditional strategies affect the making versus breaking of connections. We do not inform subjects about the structure of the network or how many of their neighbors are connected to the player they currently are evaluating. At the end of every rewiring round, each subject is told the number of others who chose to break links with her and the number of others who formed new links with her. Note that a particular subject may be part of multiple selected subject pairs and thus have the chance to update multiple links in a given round.

We examine two strategic updating scenarios. In the viscous condition, the network updates relatively infrequently, with $k = 10\%$ of pairings potentially changing each round. In the fluid condition, the network updates relatively frequently, with $k = 30\%$ of pairings potentially changing each round.

Results

We begin by evaluating the central prediction of the theoretical models, that rapid network updating can support cooperation. To do so, we examine how cooperation varies across our four conditions (Fig. 1A). Unless otherwise noted, all statistical analyses use logistic regression with robust SEs clustered on subject and session (see *Methods* for details). We see that cooperation declines steeply over time in the random network condition (coeff = -0.11 , $P < 0.001$), recreating the classic tragedy of the commons (37–39). In line with previous experimental results on static networks (12–15), we find the same pattern in the fixed network condition (coeff = -0.19 , $P < 0.001$). This result is also consistent with theoretical predictions (11), given that the average number of neighbors exceeds the benefit-to-cost ratio of cooperation. Static interaction networks do not facilitate cooperation in our experiments.

What, then, is the effect of allowing subjects to alter their interaction structure? In the viscous dynamic network condition, where 10% of connections update each round, we again see that cooperation decreases over time (coeff = -0.22 , $P = 0.013$). Subjects cannot update their partnerships quickly enough to incentivize cooperation. However, in the fluid dynamic network condition, where 30% of connections update each round, we see a qualitatively different outcome: Cooperation is robust and stable (coeff = -0.04 , $P = 0.386$).

These differences in cooperation across conditions emerge over time. Initially, as expected, there is no difference in cooperation between the fluid network condition and the other three conditions ($P > 0.30$ for all comparisons). However, over subsequent rounds of play, cooperation persists in the fluid condition and hence becomes increasingly more common relative

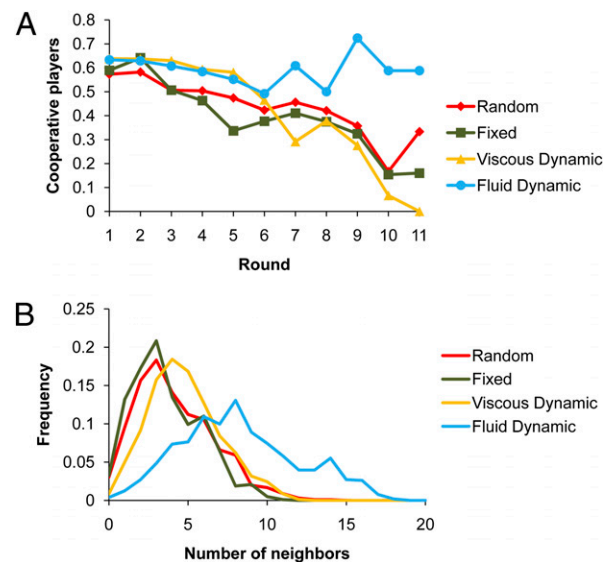


Fig. 1. (A) Dynamic social networks prevent the tragedy of the commons. The fraction of players choosing to cooperate is stable in fluid dynamic networks (blue) but declines over time in random networks (red), fixed networks (green), and viscous dynamic networks (yellow). Game length is stochastic and varies across sessions, with a constant 80% chance of a subsequent round. Although one might expect to see more cooperation in the viscous condition than in the static condition and more cooperation in the static condition than in the random condition, any differences in cooperation across these conditions are far from statistical significance (considering either all rounds or only rounds 7–11; $P > 0.45$ for all comparisons). (B) As predicted, there is greater variation in the number of connections in the rapidly updating fluid dynamic condition than in the other conditions. The fraction of individuals having each possible number of connections is shown by condition, across all sessions and rounds.

to the other conditions. Thus, we find significantly more cooperation in the fluid condition than in the other conditions in the later rounds (rounds 7–11, $P = 0.006$). This result is demonstrated further by a significant positive interaction between round number and a variable indicating membership in the fluid condition when predicting cooperation (coeff = 0.135, $P = 0.006$). This interaction indicates that cooperation in the fluid condition increases, relative to the other conditions, as round number increases. Thus, we find strong support for the prediction that rapid network updating promotes cooperation.

We now evaluate the three additional theoretical predictions (33), all of which are supported by our data. As predicted, we find greater variation in the number of connections (i.e., a greater degree of heterogeneity in the network) in the fluid condition than in the other conditions (Fig. 1*B*). For each session, we calculate the variance of number of connections across all subjects and rounds. Comparing these session-level variances shows significantly higher average variance in the fluid condition (Wilcoxon rank-sum; fluid vs. random, $P = 0.0002$; fluid vs. static, $P = 0.0002$; fluid vs. viscous, $P = 0.0003$). We also find that links between two cooperators are much more stable (i.e., less likely to be broken) in the fluid condition than are links between a cooperator and a defector or links between two defectors ($P < 0.001$ for both comparisons; Fig. 2*A*). Furthermore, we find that cooperators have more connections than defectors on average in the fluid condition ($P = 0.021$; Fig. 2*B*).

Having evaluated the theoretical predictions, we turn next to exploring the behavioral mechanisms underlying the success of cooperation in the fluid dynamic networks. We begin by examining how subjects change their cooperation decisions in response to the cooperation and defection of the individuals with whom they have interacted in the prior round. Across all conditions, we see the same intuitive pattern: The greater the percentage of a subject's partners that defected in the previous period, the less likely the subject is to cooperate (random, coeff = -0.44 , $P = 0.001$; fixed, coeff = -1.8 , $P < 0.001$; viscous, coeff = -3.62 , $P < 0.001$; fluid, coeff = -2.01 , $P < 0.001$; see *SI Appendix* for further analysis). This behavior represents classical reciprocation via changing one's cooperative action in response to the actions of one's interaction partners. However, despite the

presence of this form of reciprocation in all conditions, cooperation fails in the random, fixed, and viscous dynamic conditions and succeeds only in the fluid dynamic condition. Thus, reciprocation through changing one's cooperation action is not enough to stabilize cooperation in our experiments, most likely because of the large number of simultaneously interacting partners.

Instead, what differentiates the fluid condition from the others is the ability to reciprocate effectively by forming or breaking social ties. Consistent with the assumptions of many of the theoretical models, we find that subjects seek connections with cooperators and shun defectors (Fig. 3*A*). Subjects are more likely to form new links with those who cooperated in the previous round than with those who defected ($P < 0.001$), and they are more likely to break existing links with those who defected than with those who cooperated ($P < 0.001$). Interestingly, subjects are more discriminating when breaking links than when making links: the likelihood of making a link with a defector is significantly higher than the likelihood of breaking a link with a cooperator ($P < 0.001$). That is, although subjects almost never break links with cooperators, they sometimes make new links with defectors. This observation suggests that subjects are willing to give potential new partners the benefit of the doubt, which is reminiscent of a form of forgiveness or leniency (20, 25).

How subjects respond to alterations in the network is also important for cooperation. We find that a defector's probability of switching to cooperation increases with the number of players that broke links with her following her decision to defect, but is unaffected by the number of players that formed new links with her following her defection (logistic regression including both number of links broken, coeff = 0.19, $P < 0.001$, and number of new links formed, coeff = -0.03 , $P = 0.82$; difference between the two coefficients, $P = 0.016$). Thus, breaking links serves to discipline defectors, but forming new links with defectors does not encourage cooperation (Fig. 3*B*). Unlike defectors, cooperators are not affected by either the breaking or forming of new links (logistic regression including both number of links broken,

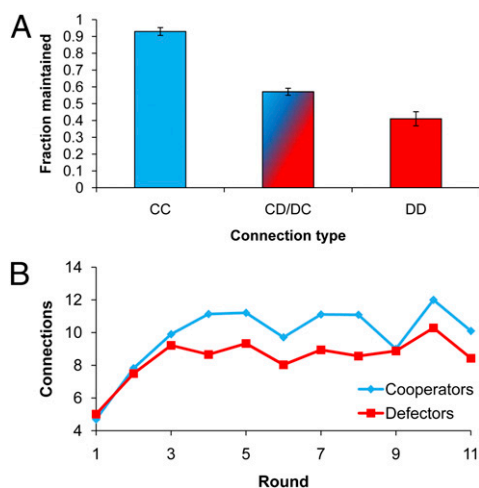


Fig. 2. (A) In the fluid dynamic condition, connections between two cooperators (CC links) are much more likely to be maintained during rewiring than connections between a cooperator and a defector (CD/DC links) or between two defectors (DD links). Interestingly, CD/DC links also are significantly more stable than DD links ($P = 0.009$). Error bars indicate SEMs, clustered on subject and session. (B) In the fluid dynamic condition, cooperators come to have more connections on average than defectors.

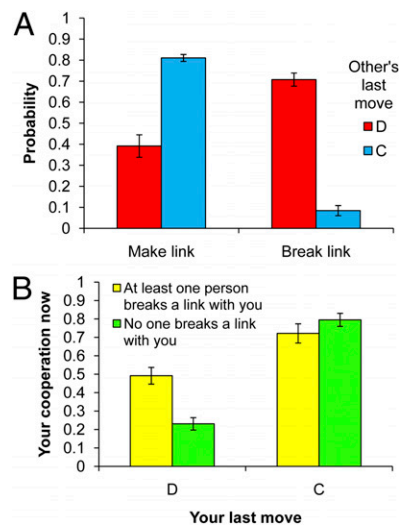


Fig. 3. (A) Subjects are more likely to make new links with players who cooperated in the previous round than with those who defected, and are more likely to break existing links with partners who defected than with partners who cooperated. (B) Subjects who defected in the previous last round are more likely to switch to cooperation if other players break links with them. Conversely, the breaking of links does not affect the behavior of cooperators. Interestingly, subjects' cooperation is not affected by whether others made new links with them. Error bars indicate SEMs, clustered on subject and session. Data from the fluid dynamic condition is shown. C, cooperation; D, defection. See *SI Appendix* for further analysis.

coeff = -0.14 , $P = 0.25$, and number of new links formed, coeff = -0.41 , $P = 0.26$). See *SI Appendix* for further analysis.

Thus, we find in our fluid dynamic network condition that subjects preferentially break links with defectors and form links with cooperators; that defectors switch to cooperation after links are broken; and that as a result stable cooperation emerges. Note that allowing frequent network updates does not necessarily mean that the network structure actually changes frequently but only that the opportunity for change often exists. Indeed, we find that the probability of a player altering the network when given the option decreases significantly over time (coeff = -0.146 , $P < 0.001$) as the network structure equilibrates. To give an intuitive feel for how the development of behavior and network structure are related in our experiments, we show a series of snapshots from fluid and static networks in Fig. 4 (see *SI Appendix* for further analysis).

Discussion

We have shown that allowing subjects to update their social network connections dynamically can stabilize cooperation in groups, where cooperation otherwise is difficult to maintain solely through traditional reciprocation via changes in cooperative behavior. Despite having to choose the same action toward an average of 8.2 other players, subjects in our fluid condition maintained a high level of cooperation by strategically making and breaking social ties. As predicted by evolutionary game theory models, our experiments show that rapid network updating promotes cooperation and leads to greater degree heterogeneity; that connections between two cooperators are longer lived than connections involving defectors; and that cooperators acquire more connections than defectors.

Our fixed-condition results are consistent with several recent experimental studies on static networks (12–15), which find no effect on cooperation. Although theoretical models suggest that

static networks can promote cooperation (6–11), these models often consider networks that are substantially larger than those typical of laboratory experiments. Exploring cooperation on very large static networks experimentally is an important direction for future study, particularly given the large scale of real-world social networks (18). The failure of static networks to promote cooperation experimentally also may be the result of subjects engaging in high rates of experimentation (or “mutation”) (15). Clustering is key to the success of cooperation in static networks, and the random variation introduced by mutation breaks up clusters. In dynamic networks, however, clustering may be maintained in the face of mutation by players rewiring the network. Should a well-connected cooperator mutate into a defector, her neighbors will sever their ties, and she will be excluded from the cooperative cluster. Exploring the effect of high mutation rates in dynamic networks is a promising area for future theoretical and empirical research.

Our results showing the ability of dynamic networks to promote cooperation in group interactions are complemented by a recent experimental study exploring pairwise interactions (24). There, it is shown that allowing subjects to choose a separate action toward each of three partners produces a substantial amount of cooperation, and that allowing subjects to break ties further improves cooperation. Similar results also are found in theoretical models of repeated two-player games on dynamic networks (40, 41). Our results also are consistent with the possibility that humans may have evolved to manipulate aspects of their social network structure to maximize their fitness (42, 43).

In an effort to use the simplest possible experimental set-up, we do not impose an explicit cost on forming, maintaining, or breaking ties in our experiment (although having more ties makes cooperating more costly, because one pays a cost for each neighbor to provide her with a benefit); and we do not limit the number of connections a subject can have. Exploring the effect of such costs and limits is an interesting direction for future empirical research. We also do not normalize payoffs across subjects with different numbers of connections, creating an incentive to increase the number of cooperative partners (the most efficient outcome is a fully connected network of cooperators). Cooperators, however, still have an incentive to avoid defecting partners in our experiments, because cooperation is costly, and the cost increases with each additional connection (unlike settings in which the lowest possible payoff is 0; e.g., ref. 8). Furthermore, although the average number of connections is higher in the fluid condition than in the other conditions, we find that cooperators accrue more connections than defectors in the fluid condition even when we restrict analysis to subjects whose degree is in the range of those observed in the other conditions (*SI Appendix*). Experimental investigation of interaction protocols with payoff normalization or asynchronous decisions is a promising direction for future research.

Following the convention of most dynamic network models from evolutionary game theory, network updating in our experiment is dyadic. When deciding whether to make or break a link with another player, subjects do not know how many neighbors they have in common with the other player. Therefore, a subject cannot be influenced explicitly by the behaviors her partners exhibit toward each other, a topic that has received considerable attention outside evolutionary game theory (44, 45). To the extent that cooperation and defection can be taken as forms of positive and negative ties, respectively, our experimental set-up can allow future exploration of the role of positive and negative valence as related to balance theory (46). Integrating information about relationships among partners in our experimental framework is a promising direction for future research.

Breaking links with defectors in our dynamic networks can be seen as a form of punishment. Most previous studies of cooperation and punishment have focused on costly punishment, in which

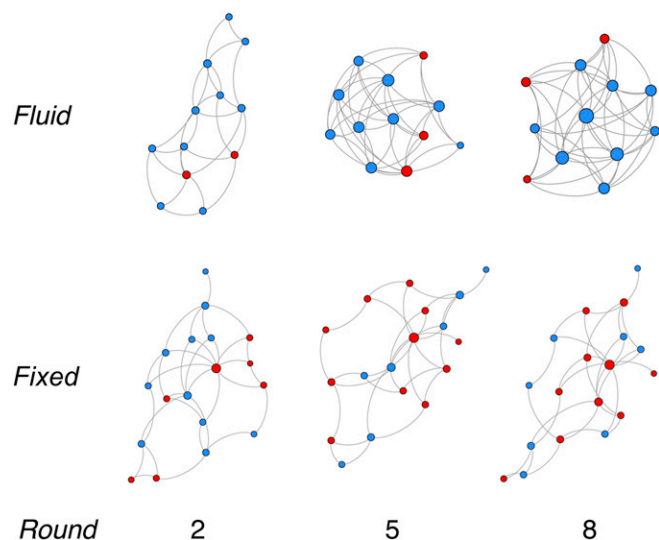


Fig. 4. Structure and strategy snapshots over time in two representative experimental sessions. Blue nodes represent cooperating subjects; red nodes represent defecting subjects. Individual connections are shown as gray lines. The network is arranged using a force-based algorithm where the edges act like springs, so that nodes in a more highly connected network are drawn more closely together. In addition, the nodes are sized according to their number of connections, and nodes with no connections are omitted. In the fluid dynamic condition, cooperation is stable, the network evolves from being relatively sparse to being quite dense, and cooperators come to have more connections than defectors. In the fixed condition, conversely, cooperation declines, and subjects with many connections are mostly defectors. Note that the connections do not change in the static network although the visualization algorithm alters the position of the nodes.

players can pay a cost to cause others to incur a cost. Costly punishment can promote cooperation (37, 38, 47–50) but reduces the payoffs for both parties and can be used against cooperators as well as defectors, often as part of retaliatory vendettas (23, 51–55). Breaking links, on the other hand, is not costly to either party and cannot be used by defectors to harm cooperators.

Another form of punishment which has been shown to promote cooperation effectively is ostracism (56, 57). In these studies, players can choose to eject each other from the group using various voting schemes, with ejected players being universally excluded from the benefits of any cooperation undertaken by nonejected group members. Breaking links in our experiments represents a form of decentralized ostracism, where each subject makes her own decisions about whom to exclude from the benefits of her individual cooperation decisions. The success of cooperation in our experiments shows that ostracism need not be coordinated to be effective.

Opting out in voluntary social dilemmas is also related to the breaking of links. In voluntary games, players can choose to not interact with their partner or group, typically earning a fixed loner's payoff instead. It has been shown both experimentally (58–60) and theoretically (61, 62) that allowing opting out promotes cooperation. However, although voluntary games are characterized by a continual cycling between cooperators, defectors, and loners, dynamic networks can lead to stable high levels of cooperation.

In the same way that breaking links can be construed as punishment, forming new links is similar to costly rewarding (but only if the player forming the new link is a cooperator). Previous experimental studies have typically found that costly rewards have mixed effects in one-shot games or in the final period of finite-length games but effectively promote cooperation as long as future interactions are possible (39, 47, 48, 63, 64; see ref. 65 for further discussion). Whereas costly rewarding typically involves a separate stage following the cooperation game, forming new links allows the rewarding of good behavior within the single framework of the cooperation game, without the addition of a second, positively non-zero-sum interaction.

In summary, we provide empirical evidence regarding cooperation in dynamic networks, confirming the predictions of a

family of evolutionary game theoretic models. Our experiments demonstrate that dynamically updating social networks can support cooperation in large groups. When social ties are fluid, people need not abandon cooperation to punish free-riders. Instead we can shun them, excluding them from the benefits of future cooperation and disincentivizing defection. It pays to cooperate today, lest you find yourself alone tomorrow.

Methods

A total of 785 subjects participated in our incentivized economic game experiments. Subjects were recruited using the online labor market Amazon Mechanical Turk (12, 34, 36) and interacted anonymously over the internet using custom software playable in a browser window. Subjects were not allowed to participate in more than one session of the experiment. In all, 40 sessions were conducted in March 2010. Each session lasted approximately 1 h on average. In each session, the subjects were paid a \$3 show-up fee. Each subject's final score summed over all rounds was converted into dollars at an exchange rate of \$1 = 1,000 points. For further discussion of the validity of experiments conducted using Mechanical Turk, as well as the details of our experiment setup, see *SI Appendix*. This research was approved by the Harvard University Committee on the Use of Human Subjects.

Unless otherwise noted, all statistical analyses are conducted at the level of the individual decision (cooperate/defect, break/maintain link, or create/do not create link) using logistic regression. Because multiple observations from the same subject are not independent, and observations from multiple individuals within the same session are not independent, we cluster the SEs in our regressions on both subject and session, as per ref. 66. Differences in cooperation between any two conditions are assessed by setting one of the conditions as the baseline, including a binary ("dummy") variable to indicate decisions from the other condition, and examining the *P* value associated with the binary variable.

ACKNOWLEDGMENTS. We are grateful for extensive input and suggestions from Feng Fu and helpful comments from Benjamin Allen, Anna Dreber, James Fowler, John Horton, and Erez Lieberman Aiden. Mark McKnight provided expert programming needed for the online experiments, and Rob Kane provided invaluable assistance running the experiments. This work was supported by Grant P01-AG031093 from the National Institute on Aging, a grant from the Pioneer Portfolio of the Robert Wood Johnson Foundation, and by the Science of Generosity Initiative of the University of Notre Dame, supported by The John Templeton Foundation. D.G.R. is funded by The John Templeton Foundation's Foundational Questions in Evolutionary Biology Prize Fellowship.

- Axelrod R, Hamilton WD (1981) The evolution of cooperation. *Science* 211:1390–1396.
- Helbing D, Yu W (2009) The outbreak of cooperation among success-driven individuals under noisy conditions. *Proc Natl Acad Sci USA* 106:3680–3685.
- Henrich J, et al. (2010) Markets, religion, community size, and the evolution of fairness and punishment. *Science* 327:1480–1484.
- Cressman R (2003) *Evolutionary Dynamics and Extensive Form Games* (The MIT Press, Cambridge, MA).
- Sigmund K (2010) *The Calculus of Selfishness* (Princeton Univ Press, Princeton).
- Nowak MA, May RM (1992) Evolutionary games and spatial chaos. *Nature* 359:826–829.
- Abramson G, Kuperman M (2001) Social games in a social network. *Phys Rev E Stat Nonlin Soft Matter Phys* 63:030901.
- Santos FC, Pacheco JM (2005) Scale-free networks provide a unifying framework for the emergence of cooperation. *Phys Rev Lett* 95:098104.
- Szabo G, Fath G (2007) Evolutionary games on graphs. *Phys Rep* 446:97–216.
- Nowak MA, Tarnita CE, Antal T (2010) Evolutionary dynamics in structured populations. *Philos Trans R Soc Lond B Biol Sci* 365:19–30.
- Ohtsuki H, Hauert C, Lieberman E, Nowak MA (2006) A simple rule for the evolution of cooperation on graphs and social networks. *Nature* 441:502–505.
- Suri S, Watts DJ (2011) Cooperation and contagion in web-based, networked public goods experiments. *PLoS ONE* 6:e16836.
- Grujić J, Fosco C, Araujo L, Cuesta JA, Sánchez A (2010) Social experiments in the mesoscale: Humans playing a spatial prisoner's dilemma. *PLoS ONE* 5:e13749.
- Cassar A (2007) Coordination and cooperation in local, random and small world networks: Experimental evidence. *Games Econ Behav* 58:209–230.
- Traulsen A, Semmann D, Sommerfeld RD, Krambeck H-J, Milinski M (2010) Human strategy updating in evolutionary games. *Proc Natl Acad Sci USA* 107:2962–2966.
- Kearns M, Suri S, Montfort N (2006) An experimental study of the coloring problem on human subject networks. *Science* 313:824–827.
- Fowler JH, Christakis NA (2010) Cooperative behavior cascades in human social networks. *Proc Natl Acad Sci USA* 107:5334–5338.
- Palla G, Barabási A-L, Vicsek T (2007) Quantifying social group evolution. *Nature* 446:664–667.
- Christakis NA, Fowler JH (2007) The spread of obesity in a large social network over 32 years. *N Engl J Med* 357:370–379.
- Nowak MA, Sigmund K (1992) Tit for tat in heterogeneous populations. *Nature* 355:250–253.
- Fudenberg D, Maskin E (1986) The Folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54:533–554.
- Dal Bó P (2005) Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *Am Econ Rev* 95:1591–1604.
- Dreber A, Rand DG, Fudenberg D, Nowak MA (2008) Winners don't punish. *Nature* 452:348–351.
- Fehl K, van der Post DJ, Semmann D (2011) Co-evolution of behaviour and social network structure promotes human cooperation. *Ecol Lett* 14:546–551 10.1111/j.1461-0248.2011.01615.x.
- Fudenberg D, Rand DG, Dreber A Slow to anger and fast to forgive: Cooperation in an uncertain world. *Am Econ Rev*, in press.
- Boyd R, Richerson PJ (1988) The evolution of reciprocity in sizable groups. *J Theor Biol* 132:337–356.
- Skyrms B, Pemantle R (2000) A dynamic model of social network formation. *Proc Natl Acad Sci USA* 97:9340–9346.
- Ebel H, Bornholdt S (2002) Coevolutionary games on networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 66:056118.
- Zimmermann MG, Eguiluz VM, San Miguel M (2004) Coevolution of dynamical states and interactions in dynamic networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 69:065102.
- Santos FC, Pacheco JM, Lenaerts T (2006) Cooperation prevails when individuals adjust their social ties. *PLoS Comput Biol* 2:e140.
- Hanaki N, Peterhansl A, Dodds PS, Watts DJ (2007) Cooperation in evolving social networks. *Manage Sci* 53:1036–1050.
- 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* 78:026117.
- Perc M, Szolnoki A (2010) Coevolutionary games—a mini review. *Biosystems* 99:109–125.

34. Horton JJ, Rand DG, Zeckhauser RJ (2011) The Online Laboratory: Conducting experiments in a real labor market. *Exp Econ*, 10.1007/s10683-10011-19273-10689.
35. Centola D (2010) The spread of behavior in an online social network experiment. *Science* 329:1194–1197.
36. Rand DG (2011) The promise of mechanical turk: How online labor markets can help theorists run behavioral experiments. *J Theor Biol* 10.1016/j.jtbi.2011.1003.1004.
37. Fehr E, Gächter S (2002) Altruistic punishment in humans. *Nature* 415:137–140.
38. Ostrom E, Walker J, Gardner R (1992) Covenants with and without a sword: Self-governance is possible. *Am Polit Sci Rev* 86:404–417.
39. Milinski M, Semmann D, Krambeck HJ (2002) Reputation helps solve the 'tragedy of the commons'. *Nature* 415:424–426.
40. Pacheco JM, Traulsen A, Ohtsuki H, Nowak MA (2008) Repeated games and direct reciprocity under active linking. *J Theor Biol* 250:723–731.
41. Wadriil L, de Silva JKL (2009) Adoption of simultaneous different strategies against different opponents enhances cooperation. *Europhys Lett* 86:38001.
42. Fowler JH, Dawes CT, Christakis NA (2009) Model of genetic variation in human social networks. *Proc Natl Acad Sci USA* 106:1720–1724.
43. Fowler JH, Settle JE, Christakis NA (2011) Correlated genotypes in friendship networks. *Proc Natl Acad Sci USA* 108:1993–1997.
44. Rapoport A (1957) Contribution to the theory of random and biased nets. *Bull Math Biol* 19:257–277.
45. Davis JA (1970) Clustering and hierarchy in interpersonal relations: Testing two graph theoretical models on 742 sociomatrices. *Am Sociol Rev* 35:843–851.
46. Marvel SA, Kleinberg J, Kleinberg RD, Strogatz SH (2011) Continuous-time model of structural balance. *Proc Natl Acad Sci USA* 108:1771–1776.
47. Rand DG, Dreber A, Ellingsen T, Fudenberg D, Nowak MA (2009) Positive interactions promote public cooperation. *Science* 325:1272–1275.
48. Rockenbach B, Milinski M (2006) The efficient interaction of indirect reciprocity and costly punishment. *Nature* 444:718–723.
49. Hauert C, Traulsen A, Brandt H, Nowak MA, Sigmund K (2007) Via freedom to coercion: The emergence of costly punishment. *Science* 316:1905–1907.
50. Boyd R, Gintis H, Bowles S, Richerson PJ (2003) The evolution of altruistic punishment. *Proc Natl Acad Sci USA* 100:3531–3535.
51. Nikiforakis N (2008) Punishment and counter-punishment in public goods games: Can we still govern ourselves? *J Public Econ* 92:91–112.
52. Cinyabuguma M, Page T, Putterman L (2006) Can second-order punishment deter perverse punishment? *Exp Econ* 9:265–279.
53. Herrmann B, Thöni C, Gächter S (2008) Antisocial punishment across societies. *Science* 319:1362–1367.
54. Rand DG, Armao JJ, Nakamaru M, Ohtsuki H (2010) Anti-social punishment can prevent the co-evolution of punishment and cooperation. *J Theor Biol* 265:624–632.
55. Rand DG, Nowak MA (2011) The evolution of antisocial punishment in optional public goods games. *Nat Commun* 2:434.
56. Masclet D (2003) Ostracism in work teams: A public good experiment. *Int J Manpow* 24:867–887.
57. Cinyabuguma M, Page T, Putterman L (2005) Cooperation under the threat of expulsion in a public goods experiment. *J Public Econ* 89:1421–1435.
58. Semmann D, Krambeck H-J, Milinski M (2003) Volunteering leads to rock-paper-scissors dynamics in a public goods game. *Nature* 425:390–393.
59. Miller RR (1967) No play: A means of conflict resolution. *J Pers Soc Psychol* 6:150–156.
60. Orbell JM, Dawes RM (1993) Social welfare, cooperators' advantage, and the option of not playing the game. *Am Sociol Rev* 58:787–800.
61. Hauert C, De Monte S, Hofbauer J, Sigmund K (2002) Replicator dynamics for optional public good games. *J Theor Biol* 218:187–194.
62. Hauert C, De Monte S, Hofbauer J, Sigmund K (2002) Volunteering as Red Queen mechanism for cooperation in public goods games. *Science* 296:1129–1132.
63. Sefton M, Shupp RS, Walker JM (2006) The effect of rewards and sanctions in provision of public goods. *Econ Inq* 45:671–690.
64. Sutter M, Haigner S, Kocher MG (2010) Choosing the Stick or the Carrot? Endogenous Institutional Choice in Social Dilemma Situations. *Rev Econ Stud* 77:1540–1566.
65. Almenberg J, Dreber A, Apicella CL, Rand DG (2011) Third Party Reward and Punishment: Group Size, Efficiency and Public Goods. *Psychology and Punishment* (Nova Science Publishers, Hauppauge, NY).
66. Cameron AC, Gelbach JB, Miller DL (2011) Robust inference with multiway clustering. *J Bus Econ Stat* 29:238–249.