Cooperation and assortativity with dynamic partner updating

Jing Wang[a,1], Siddharth Suri[b,1], and Duncan J. Watts[b,1]

[a]Department of Information, Operations and Management Sciences, Leonard N. Stern School of Business, New York University, New York, NY 10012; and [b]Microsoft Research New York City, New York, NY 10104

Edited by Matthew O. Jackson, Stanford University, Stanford, CA, and accepted by the Editorial Board July 10, 2012 (received for review December 19, 2011)

The natural tendency for humans to make and break relationships is thought to facilitate the emergence of cooperation. In particular, allowing conditional cooperators to choose with whom they interact is believed to reinforce the rewards accruing to mutual cooperation while simultaneously excluding defectors. Here we report on a series of human subjects experiments in which groups of 24 participants played an iterated prisoner's dilemma game where, critically, they were also allowed to propose and delete links to players of their own choosing at some variable rate. Over a wide variety of parameter settings and initial conditions, we found that dynamic partner updating significantly increased the level of cooperation, the average payoffs to players, and the assortativity between cooperators. Even relatively slow update rates were sufficient to produce large effects, while subsequent increases to the update rate had progressively smaller, but still positive, effects. For standard prisoner's dilemma payoffs, we also found that assortativity resulted predominantly from cooperators avoiding defectors, not by severing ties with defecting partners, and that cooperation correspondingly suffered. Finally, by modifying the payoffs to satisfy two novel conditions, we found that cooperators did punish defectors by severing ties, leading to higher levels of cooperation that persisted for longer.

network science | dynamic networks | web-based experiments

Why, and under what circumstances, presumptively selfish individuals cooperate is a question of longstanding interest to social science (1, 2) and one that has been studied extensively in laboratory experiments, often as some variant of a public goods game (also called a voluntary contribution mechanism) or a prisoner's dilemma (see refs. 2 and 3 for surveys). One broadly replicated result from this literature is that in unmodified, finitely repeated games cooperation levels typically start around 50–60% and steadily decline to around 10% in the final few rounds (2). The explanation for this decline is that when conditional cooperators, who are thought to constitute as much as 50% of the population (4), are forced to interact with a heterogeneous mix of other types, in particular free riders, the conditional cooperators tend to reduce their contributions over time (3).

Previous work has shown that explicit enforcement mechanisms such as punishment (5) and reward (6) can alleviate the observed decline in cooperation in fixed groups. Here we investigate a related but distinct mechanism that exploits a well-known feature of human social interactions—namely that they change over time, as individuals form new relationships or sever existing ones (7). By allowing participants engaged in a repeated game of cooperation to update their interaction partners dynamically, cooperation levels might be enhanced in two ways. First, conditional cooperators could implicitly punish defectors by denying them a partner, thereby encouraging potential defectors to cooperate. Second, conditional cooperators could benefit from assortative mixing (1) by avoiding defectors and seeking cooperators, thus sustaining their own cooperative tendencies.

Two recent studies have argued that dynamic partner updating is capable of promoting cooperation (8, 9); however, the studies, in fact, reached somewhat different conclusions. In particular, whereas one study (8) found that allowing individuals to update one partner every round led to a significant increase in cooperation, the other study (9) found no significant increase at that rate. The latter result has been interpreted as supporting prior theoretical claims that dynamic partner updating enhances cooperation only when it exceeds a critical threshold rate (10). Because the former study considered only one rate, however, and the latter study considered only two rates, and because in both cases the effect sizes were small, the existence or otherwise of a threshold rate remains ambiguous. Moreover, both studies used an "active linking" (10) design in which players were offered opportunities to make and break links with randomly chosen partners, but were not allowed to choose with whom they wished to make and break links; moreover, players were unable to refuse new links proposed by others. Although appealingly simple, active linking is somewhat unrealistic. In real-world social networks individuals can generally select among a multiplicity of potential partners (11) and hence can choose new partners selectively. Moreover, social networks are typified by high levels of reciprocity (12–15), implying that mutual acceptance of new links is the social norm.

Our study builds upon this work in three ways. First, our design is fully endogenous, allowing individuals to decide with whom they will make and break ties. As we explain below, the resulting effect sizes are much larger than in previous studies of dynamic networks (8, 9), reaching close to 100% cooperation in some cases. Second, we consider an extremely wide range of update rates, affording us a much clearer understanding of the importance of varying rates. We find no evidence of the hypothesized threshold effect (9, 10), instead finding significant and positive increases in cooperation at rates well below those previously reported. Finally, and in contrast to both previous studies that considered only one set of payoffs, we manipulate the payoff structure itself, effectively varying the attractiveness of the "outside option" (16), meaning roughly the payoff associated with choosing not to interact with a potential partner. We find that only in the presence of an attractive outside option do conditional cooperators punish defectors (by proactively deleting ties with them). By contrast, when the outside option is less attractive, we find that cooperators tolerate defecting partners, eventually leading them to defect themselves.

Our work is also related more generally to a number of recent experiments that have investigated various aspects of the relationship between cooperation and partner selection, such as unilateral vs. bilateral choice (17, 18), the effect of introducing an outside option of varying attractiveness (16), and the attributes of the individuals (age, sex, race, etc.) as predictors of selection and cooperation (19, 20). Although our treatment of the outside option is consistent with previous work (16), it is distinct in that it extends it to the case of a dynamic network. Finally,

Author contributions: J.W., S.S., and D.J.W. designed research; J.W. and S.S. performed research; J.W., S.S., and D.J.W. analyzed data; and S.S. and D.J.W. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. M.O.J. is a guest editor invited by the Editorial Board.

Freely available online through the PNAS open access option.

[1]To whom correspondence may be addressed. E-mail: jwang5@stern.nyu.edu, suri@microsoft.com, or duncan@microsoft.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1120867109/-/DCSupplemental.

SOCIAL SCIENCES

other related work (21, 22) has examined how individuals select groups or are excluded by them. Although at a high level these papers clearly resemble both the partner selection literature and dynamic updating studies such as ours, they differ substantially from both literatures in that the object of selection (21) or the actor (22) is the group, not the individual.

Experimental Setup

We conducted a series of online human subjects experiments in which groups of 24 participants played an iterated prisoner's dilemma (PD) game, where in addition to choosing their action each round—cooperate or defect—they also were given the opportunity to update their interaction partners at some specified rate, which was varied across experimental conditions. (See *SI Appendix*, Figs. S1 and S2 for details of the experimental platform and recruiting.) All games comprised 12 "strategy update" rounds during which players could update their strategy: cooperate (C) or defect (D). Consistent with standard PD conditions, a cooperator received four points when interacting with another cooperator, but lost one point when interacting with a defector. A defector received seven points when interacting with a cooperator and one point when interacting with another defector (see *SI Appendix* for details).

In addition, after every r strategy update rounds, players entered a "partner updating" turn in which they were permitted to make up to k partner updates. By adjusting r and k we were therefore able to explore an extremely wide range of relative updating rates, from one opportunity every several strategy update rounds to several opportunities every round. A single partner update comprised either severing a link with an existing partner or proposing a link to a new partner, where, importantly, players could choose the partner in question. Also of importance, our design specified that severing a link was a unilateral action, requiring no consent from the corresponding partner; however, proposing a link was a bilateral action that required acceptance for the edge to be formed. These requirements in turn necessitated that each partner-updating turn consist of two phases: a proposal phase, during which players submitted their link proposals and link deletions, and an approval phase during which they were required to accept or reject any outstanding link proposals. If a proposal was rejected, the proposing player could not reuse that action, and hence players had an incentive to make proposals they thought would be accepted. After each partner-updating turn was completed, the network of partners was updated to reflect severed and accepted links, and a new strategy update round commenced.

Players were shown the identities (anonymous player IDs) and strategy choices for up to the previous five rounds for all players. Players were also shown who they were and were not connected to, their current cumulative payoff, their payoff from the previous round, and the time remaining in the current round. Consistent with previous work (23, 24), players were not given explicit information about the structure of the network beyond their immediate network neighbors (see *SI Appendix*, Figs. S3–S5 for screenshots). Nevertheless, to test for the possibility that initial conditions could affect outcomes, players were randomly assigned to positions in one of two initial network topologies: "cliques" composed of four cliques of six players each; and "random" comprising a random regular graph, where in both initial graphs, each player had exactly five neighbors (i.e., partners). These topologies were chosen because they are as different as possible in terms of standard network metrics such as path length and clustering coefficient (25, 26) while still maintaining the same number of ties per person.

Results

Fig. 1 shows the average fraction of cooperators by round for k = 1, 3, 5 and r = 1, 3, 6 (*Top*, *Middle*, and *Bottom* rows, respectively) and for the cliques (*Left* column) and random (*Right* column) initial conditions, respectively. For r = 1 and r = 3, we observe three striking features of networks with dynamic partner updating: first, cooperation levels start significantly higher than

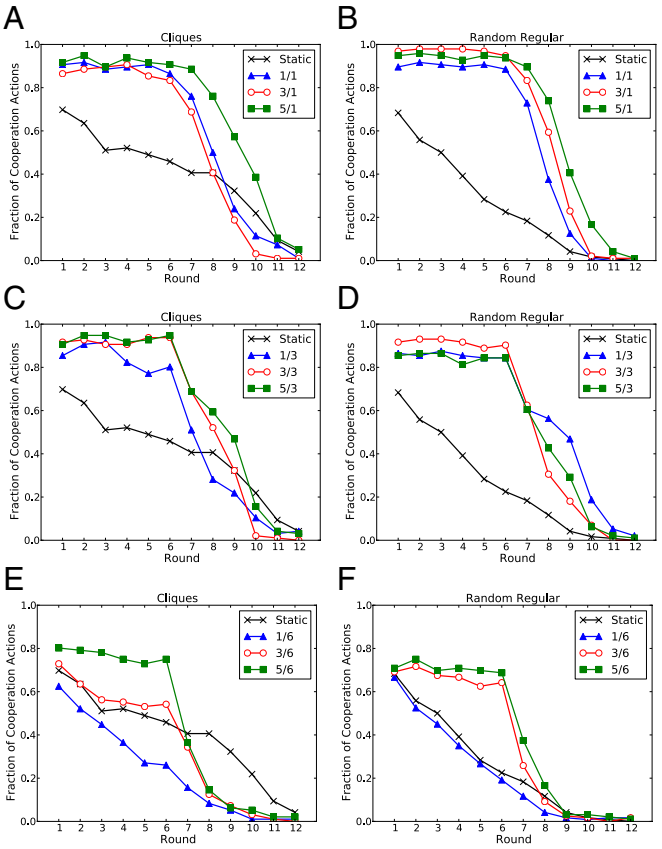

Fig. 1. Fraction of cooperation by round for the cliques (*A*, *C*, and *E*) and random (*B*, *D*, and *F*) initial conditions. (*Top* row) r = 1; (*Middle* row) r = 3; (*Bottom* row) r = 6. Symbols indicate different values of k (solid triangles, k = 1; open circles, k = 3; solid squares, k = 5).

in fixed networks; second, cooperation levels remain between 80% and 100% in the presence of updates even as they decline in fixed networks; and third, cooperation declines rapidly as the game nears its end, finishing as low as in the absence of partner updates. Taken as a whole this behavior is far from the Nash prediction of all players defecting on all turns (see *SI Appendix* for the theorems and proofs). We note, however, that for r = 6, the initial increase is largely absent, and the persistence effect is present only for the higher values of k = 3, 5. This lack of effect for the r = 6 case can be understood by noting that the players experienced only one partner-updating opportunity (because round 12 was the final round of the game); thus for the r = 6, k = 1 case, players were permitted to update just one partnership in the entire game. Because this treatment is only slightly different from the static case, it is unsurprising that its effect, if any, was small.

Next, Fig. 2*A* summarizes these findings for all values of r and k, showing the average rate of cooperation as a function of the total number of updates u per player over the course of a game [i.e., $u = k^*(12/r - 1)$]. Consistent with Fig. 1, Fig. 2*A* shows that increases in cooperation rates were relatively small for the very lowest (r = 6) rates of updating (i.e., compared with the variation between the two static cases). However, when r = 1, 3 the average cooperation rate was substantially higher than the static (i.e., no partner updating) case. Correspondingly, average payoffs also increased severalfold over the static case (see *SI Appendix*, Fig. S6*A* for details). To account for subject- and game-level variations, the treatment effects in Fig. 2*A* were estimated using a nonnested, multilevel model (27) with error terms for treatment, subject, and game as well as the experience level of a given subject in a given game (see *Materials and Methods* for more details). To test for significance, Fig. 2*B* shows the estimated

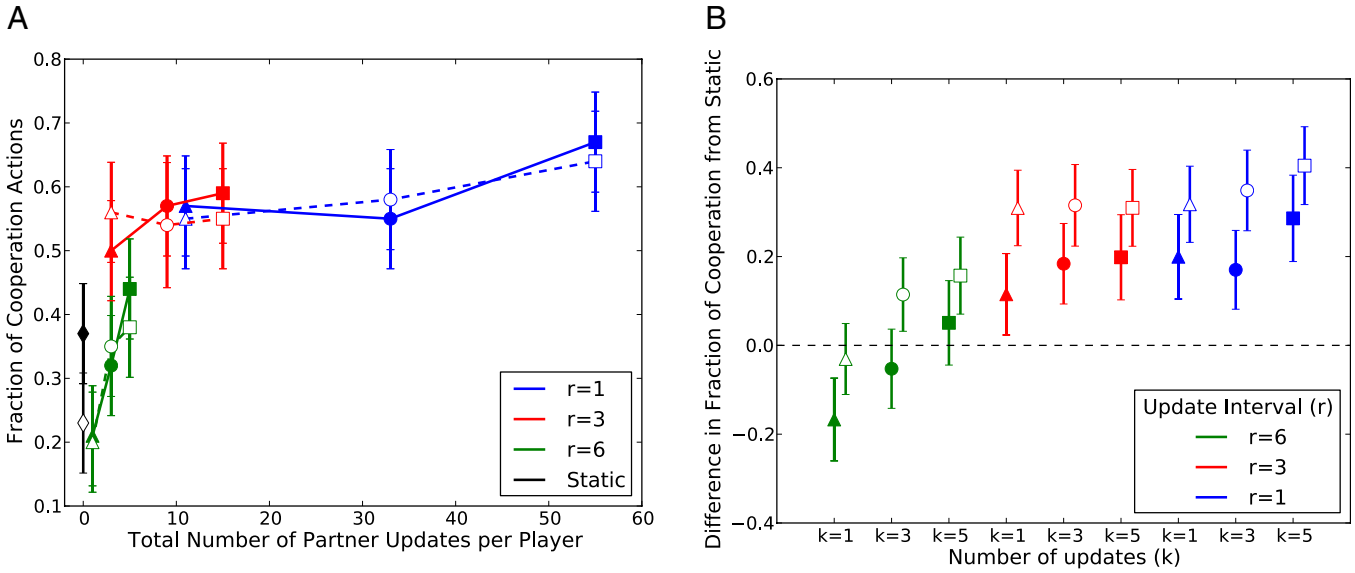

Fig. 2. Average fraction of cooperation as a function of partner update rate (*A*) and estimated difference in fraction of cooperation from the corresponding static cases as a function of k (*B*) for cliques (dashed lines) and random (solid lines) initial conditions, for $r = 1, 3, 6$ and $k = 0, 1, 3, 5$. Symbols indicate different values of k (triangles, $k = 1$; circles, $k = 3$; squares, $k = 5$). Error bars are 95% confidence intervals (see *Materials and Methods* for details).

difference in average cooperation levels between the various treatments and the corresponding static case, where error bars represent 95% confidence intervals. For the cliques initial condition all $r = 1$ and $r = 3$ treatments yield positive effects that are significant at the 5% level, and for the random regular initial condition the $r = 6$, $k = 3, 5$ conditions are also positive and significant. In general, regardless of initial condition, allowing as few as one update every three rounds was sufficient to significantly increase cooperation (see *SI Appendix*, Fig. S6*B* for a similar analysis of average payoff levels), a rate that is well below the previously reported threshold for a positive effect (9).

Next, Fig. 3 shows the relationship between assortativity and cooperation for $r = 1$ (see *SI Appendix*, Figs. S7 and S8 for equivalent $r = 3, 6$ results). To quantify assortativity, we first label each link in the network as CC, CD, or DD according to the states of the players who share the link. Then we define CC, CD, and DD assortativity as the difference between the observed fraction of CC, CD, and DD links and the corresponding fraction that would be expected under a random permutation of the player IDs. The difference is necessary to account for "baseline assortativity" (also called baseline homophily) (28), which varies dramatically over the course of the game as the overall fraction of cooperators changes. Fig. 3 *A* and *B* show that for both cliques and random initial conditions CC assortativity is positive for most of the game and CD assortativity is negative, whereas DD assortativity is essentially absent. In other words cooperators displayed a tendency to avoid defectors and gravitate to other cooperators, whereas defectors were neutral with respect to other defectors.

Finally, Fig. 3 *C* and *D* show that players continued to add links throughout the game (see also *SI Appendix*, Figs. S7 and S8 for $r = 3, 6$). Although the addition of links is superficially consistent with the Nash prediction (see *SI Appendix* for details), the equilibrium analysis also predicts that all players defect on all turns; hence players form links with each other on the grounds that the payoff to (D, D) exceeds their outside option (16) (a payoff of zero for having no links). In reality, however, it is not only defectors who accept and maintain ties with other defectors. For the $r = 1$ case, for example, cooperators also accepted proposals from defectors roughly 40% of the time and rarely deleted them, even though such ties were costly. Overall, deletions accounted only for 10% of updates (see *SI Appendix*, Table S1 and Figs. S9–S11 for more details). Moreover, defectors were also more than twice as likely to propose links to other defectors than to cooperators (0.24 for D → D vs. 0.1 for D → C). Together, these results suggest that observed assortativity derived less from cooperators "punishing" defectors by deleting ties and more from two related mechanisms: (*i*) cooperators avoiding defectors and (*ii*) defectors failing to propose links to cooperators in the first place. A striking illustration of the lack of punitive deletion can be seen in the cases of $r = 1$ and $k = 3, 5$ in Fig. 3 *C* and *D*. By round 4 the graph was close to, and in some trials precisely, a clique. Because in these cases there were almost no edges available to be added, deleting edges exerted no opportunity cost. Nevertheless when exposed to a small number of defectors in later rounds, cooperators chose to defect rather

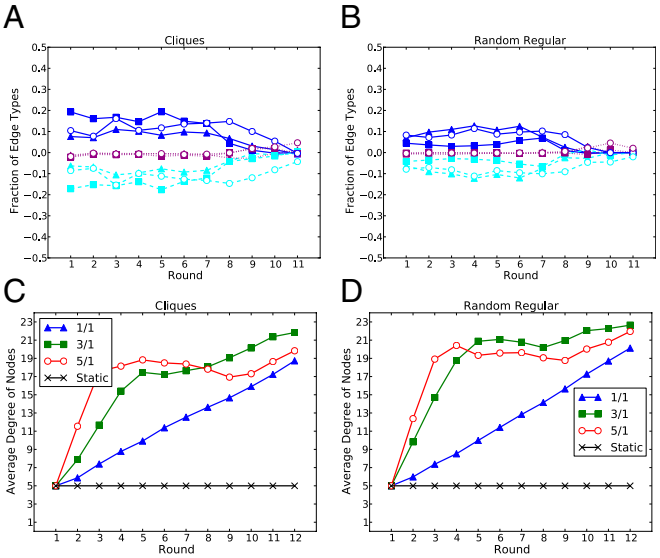

Fig. 3. Assortativity by round for cliques (*A*) and random (*B*) initial conditions for $r = 1$. The differences between observed and baseline homophily for CC, CD, and DD links are shown by dark blue, light blue, and magenta lines, respectively. $k = 1, 3, 5$ is shown by triangles, circles, and squares, respectively. Average degree by round for cliques (*C*) and random (*D*) initial conditions: $r = 1$ and $k = 0, 1, 3, 5$.

than cut ties to the defectors, resulting in a defection cascade (see *SI Appendix*, Fig. S12 for an illustrative example).

Although surprising in light of simulation and theoretical models that assume punitive partner deletion (10, 29–33), its relative rarity can be understood in terms of two related conditions of the payoff matrix. First, because the benefit to a cooperator forming an additional tie with another cooperator (four points) outweighed the loss of maintaining a tie with a defector (minus one point), when a cooperator was faced with a choice to sever a tie with a defector or add new tie to a cooperator, players rationally chose the latter. Second, because defectors gained a positive amount from all interactions (i.e., seven points when interacting with cooperators and one point when interacting with defectors), defecting players rationally preferred to maintain all their ties, even with other defecting players.

The upshot of these two conditions of the payoff matrix was that early on in the game, when most players were cooperating, all players wished to add links as fast as possible, even if that required tolerating occasional defectors. Defectors, therefore, were not punished for their actions, thereby persisting and encouraging cooperators to switch. Finally, as the end game neared, and the availability of other cooperators diminished, all players preferred to defect rather than cooperate and cut ties with defectors. One can conclude that dynamic partner updating combined with these payoffs resulted in cooperation being promoted, but not sustained.

To quantify the effect of these conditions on cooperation and assortativity, we conducted a new set of experiments with modified payoffs such that cooperators were penalized five points when interacting with defectors instead of just one point, and defectors lost one point each when interacting with each other, rather than gaining one point each. These payoffs still fulfilled the usual requirements of the PD, but had two additional properties: (*i*) maintaining a tie with a defector was more costly for a cooperator than creating a new tie with another cooperator; and (*ii*) defectors could no longer profit by interacting with other defectors. In addition, because the results from the $k = 3$ condition were not substantively different from $k = 1, 5$, we replaced this condition with $k = (N - 1) = 23$ to explore more of the parameter space (as before, a static condition was also included for comparison). By permitting every player to propose or delete an edge with every other player each round, the $k = 23$ condition effectively removed any budget constraint on deleting links, hence further increasing the likelihood of deleting links to defectors. Also, because it was previously established that initial conditions did not have a qualitative impact on cooperation levels, only the cliques initial condition was studied for the modified payoffs.

For these new payoffs, equilibrium analysis again predicts that all players will defect on all rounds; however, in contrast to the initial payoffs, the analysis predicts that players will delete as many links as possible, leading to an empty graph for the parameters we tested (see *SI Appendix* for details). Clearly, we did not expect all players to defect on all rounds; hence the equilibrium prediction regarding the network is again best viewed as a baseline for comparison rather than a hypothesis. We emphasize, however, that although the new payoffs were designed to make interacting with defectors less attractive, it is not obvious that they would lead either to higher levels of cooperation or to longer persistence. The reason is that the punishment for a cooperator interacting with a defector was also greater than in the original payoffs and could easily have increased the prevalence of strategic defections (i.e., in anticipation of others defecting or to exploit cooperators), especially for low rewiring rates.

It is therefore striking that, as Fig. 4*A* shows, the modified payoffs led to significantly higher rates of cooperation and longer persistence of cooperation for all rewiring rates. Cooperation levels approached and in many treatments reached 100% and were sustained at that level until close to the end of the game. In fact, in seven of the nine trials performed with the modified payoffs, cooperation levels were over 90% in 9 of the 12 rounds. Fig. 4*B*, *Upper* summarizes the effects, showing that a small amount of dynamic partner updating, $r = 1$, $k = 1$, resulted in

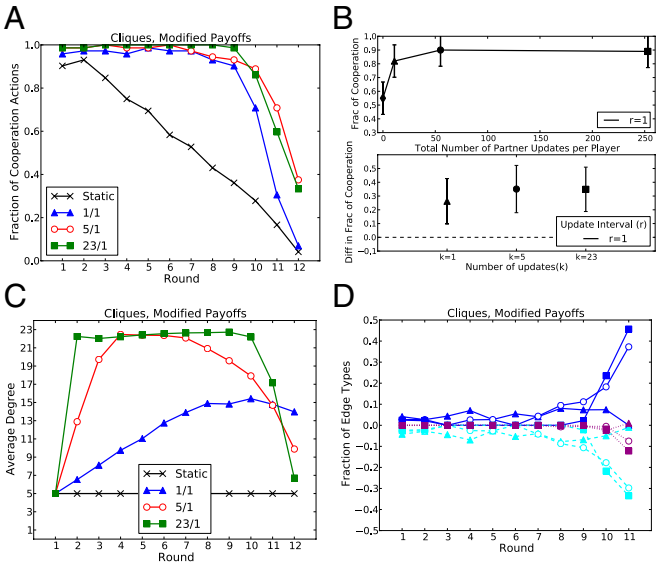

Fig. 4. Results for modified payoffs, $r = 1$ and $k = 0, 1, 5, 23$, cliques initial condition. (*A*) Fraction of cooperation by round. (*B*) Aggregate fraction of cooperation as a function of total number of partner updates (*Upper*) and estimated difference in fraction of cooperation from the static case as a function of k (*Lower*). (*C*) Average degree by round. (*D*) CC (dark blue), CD (light blue), and (DD) (magenta) assortativity by round, where $k = 1, 5, 23$ is shown by triangles, circles and squares, respectively. In *B* the error bars are 95% confidence intervals (*Materials and Methods*).

large increases in cooperation levels over the static case. As before, further increases in update rate resulted in smaller increases in cooperation levels, where, interestingly, the highest rate of rewiring, $k = 23$ did not lead to appreciably more cooperation or higher average payoffs than were obtained in the $k = 5$ condition, consistent with our hypothesis that the earlier payoff structure, not the constraints on partner updating imposed by the update rate, was responsible for the previous absence of punitive deletions. Fig. 4*B*, *Lower*, meanwhile, shows that the differences between the average cooperation levels of the $r = 1$, $k = 1, 5, 23$ and the static case were significant at the 5% level (see *Materials and Methods* for more details and *SI Appendix*, Fig. S13 for the corresponding analysis of player earnings).

Fig. 4 *C* and *D* also shows interesting similarities with and differences from the original payoffs: In early rounds, players added edges much as they did previously, leading in fact to even denser networks; however, unlike in the previous case, they deleted edges as rapidly as possible toward the end of the game as defection began to spread. In fact, in the modified payoffs 22% of the partner update actions were deletions, compared with only 10% in the original payoffs (see *SI Appendix*, Table S2 and Fig. S14 for more details). The increased amount of deletions suggests that the amount of exploitation of cooperators by defectors was mitigated as a result of these new payoffs. CC and CD assortativity, meanwhile, had the same signs as before, but much lower magnitudes in the early rounds for the simple reason that cooperation levels were close to 100%, and hence baseline CC assortativity was also close to 100%. As the end of the game approached, however, the assortativity measure increased dramatically as cooperators attempted to segregate themselves from defectors, and defectors also began cutting ties with other defectors—an effect that is illustrated graphically in *SI Appendix*, Fig. S15. Thus, in the modified payoffs cooperators and conditional cooperators were able to separate themselves from defectors, further helping cooperation to be promoted and sustained.

Discussion

At a high level, our results are consistent with prior work (8–10, 30–35), in that allowing players to form new ties and sever existing ones generates assortative mixing between cooperators along with increased cooperation. However, our results advance upon previous work in three key respects.

First, in spanning a wide range of update rates, Fig. 2 goes beyond the qualitative claim that updating should aid cooperation, revealing the functional form of the relationship. Interestingly, for both cooperation and payoffs the effect of updating is strongly concave, meaning that small increases in the update rate near zero correspond to much larger effects than subsequent increases (i.e., the marginal return to increasing update rate is strongly decreasing). This finding therefore helps clarify previous theoretical and experimental work, which has made conflicting claims regarding the importance of the update rate. In some cases (31, 32) cooperation has been claimed to increase smoothly with update rate, whereas in others (9, 10, 34, 35) partner updating has been claimed to impact cooperation only when the rate exceeds a critical threshold. We found that both cooperation and payoffs were sensitive to the update rate across the entire range and that the effects became very large and significant at much lower rates than had been found previously. Thus, high rates of updating are not required to realize measurable improvements.

Second, our results revealed that the effect of dynamic partner selection on cooperation levels depends on a novel condition of the PD payoffs. Unlike in static games, where only the relative payoffs matter, in games with partner updating the absolute payoff that a tie yields becomes highly relevant. The rationale behind this result is obvious—when faced with a choice between adding a profitable link and severing a costly one, players rationally chose the action that yielded the higher aggregate payoff. Cooperators did not sever ties to defectors when the result of cutting a tie to a defector yielded less than the gain of a link to another cooperator. But the resulting failure to punish defectors also had a less obvious consequence: Network density increased even as defectors proliferated, thereby further increasing the temptation to defect, leading ultimately to a sudden and irreversible collapse in cooperation. Because cooperators could not segregate themselves from defectors, cooperation was promoted early but not sustained. Only when the cost of interacting with defectors outweighed the benefits to cooperators of forming new ties with each other did we see punitive deletion of links with defectors. This increased willingness to sever links with defectors in turn isolated from cooperators, leading ultimately to higher levels of cooperation that were sustained until almost the end of the game.

Third, we attribute the much greater sensitivity and effectiveness of the update rate, relative to previous results (8, 9), to two features of our design that capture important elements of real-world social networks: first, that players could choose which others they wish to make or break ties with (as opposed to having those choices imposed exogenously); and, second, that new partnerships required the consent of both partners, whereas existing partnerships could be terminated unilaterally. Allowing participants to choose their partners allowed cooperators to separate themselves from defectors more effectively than with random partner choices. As a result, cooperation levels approached and in some cases were sustained at nearly 100%, a rate far higher than prior work which showed only a slight increase in cooperation over the baseline (9).

In closing, we note that our focus on dynamic partner updating complements previous experimental work that has explored related mechanisms for increasing cooperation, such as punishment (36), reward (6), assortative group formation (21), and ostracism (22, 37). Although clearly analogous in some respects, dynamic partner updating is distinct in others. First, in contrast to explicit punishment and reward mechanisms, fully endogenous partner updating of the kind we have studied effectively uses implicit punishment, by link deletion, and implicit reward, by proposing or maintaining links. Clearly it is not always feasible for individuals to choose with whom they interact, in which case explicit mechanisms may be required; however, our results suggest that when they are free to choose, other, more explicit, forms of punishment and reward may be unnecessary. Second, in contrast to assortative group formation and ostracism, both of which require coordination among a group of individuals, partner updating can be accomplished in an entirely distributed manner, via the natural process of individuals making and breaking ties with their choice of others. For both these reasons, along with the frequently large size of the effects we observe, dynamic partner updating deserves to be considered among the most promising levers for eliciting cooperation between humans, especially in informal settings. Nevertheless, the specific conditions under which different forms of feedback—punishment, reward, ostracism, or dynamic partner selection—are most realistic and/or effective in practice remain an important question for future work.

Materials and Methods

This research was reviewed and approved by Yahoo! Labs' Human Subjects Research process. Correspondingly, informed consent was obtained from all participants (see *SI Appendix* for informed consent statement). All experiments were conducted online using Amazon's Mechanical Turk, a crowd-sourcing platform that is increasingly used to conduct experimental behavioral research (9, 23, 38–41). Over the course of 4 wk, a total of 108 unique subjects participated in a total of 94 experiments (82 for the initial payoffs and 12 for the modified payoffs), where each experiment required 24 subjects to participate simultaneously (see *SI Appendix* text and *SI Appendix*, Figs. S1 and S2 for details of recruiting). One consequence of this recruiting strategy was that some individuals played many games, whereas others played only once; hence the possibility arises that overrepresented individuals will bias our results, either because they are systematically different from those who play rarely or because they learn to play differently with experience. In addition, it is well known that cooperation levels in iterated games of cooperation exhibit temporal dependencies, in the sense that random differences in initial cooperation levels persist over many rounds. To mitigate potential interactions between treatment and other (e.g., learning, time of day) effects, the order in which the various treatments were applied was randomized. In our analysis, moreover, we accounted for the various forms of nonindependence across observations (repeated observations of individual subjects, game effects, and learning effects), by fitting the data to the following nonnested, multilevel model (27):

$$y_i = \beta^{\text{treatment}}_{t[i]} + \beta^{\text{subject}}_{s[i]} + \beta^{\text{game}}_{g[i]} + \beta^{\text{history}}_{h[i]} + \varepsilon_i, \quad [1]$$

where y_i is the expected cooperation level for the ith observation (i = 1, … n_{obs}, and $t[i]$, $s[i]$, $g[i]$, and $h[i]$ are all index variables that map the ith observation to a particular treatment, subject, game, and experience level, respectively. Each observation refers to the average contribution of a particular player in a single game; hence for the initial payoffs, we have n_{obs} = 82 × 24 = 1,968, and for the modified payoffs, n_{obs} = 12 × 24 = 288). Moreover, $\beta^{\text{treatment}}_{t[i]}$ is a dummy variable for treatment, where $t[i]$ = 1, … , 20 for the initial payoffs (we have two initial conditions, a static case for each, and r = 1, 3, 6 and k = 1, 3, 5; hence 2 × 3 × 3 + 2 = 20 treatments), and for the modified payoffs $t[i]$ = 1, … , 4 (we have one initial condition, a single static case, and r = 1 and k = 1, 5, 23, and hence 4 treatments); $\beta^{\text{subject}}_{s[i]} \sim \mathcal{N}(0, \sigma^2_{\text{subject}})$ is a group-level predictor for subjects; $\beta^{\text{game}}_{g[i]} \sim \mathcal{N}(0, \sigma^2_{\text{game}})$ is a group-level predictor for games; and $\beta^{\text{history}}_{h[i]} \sim \mathcal{N}(0, \sigma^2_{\text{history}})$ is a group-level predictor for the number of games played by player $s[i]$ at the time of the $g[i]$th game. Note that unlike for subject, game, and history effects, we do not model the treatment effects, preferring the simpler and more conservative approach of using a dummy variable for each treatment and hence avoiding the need to worry about potentially erroneous distributional assumptions (27). To test for significance, we computed 95% confidence intervals for the difference between each treatment and its corresponding static case; hence if zero is not contained in the interval, then the null hypothesis that they are the same can be rejected at the 5% level. As described in the main text, for the initial payoffs the null hypothesis can be rejected for all treatments except for r = 6 and k = 1, 3, 5 for the cliques initial conditions and r = 6 and k = 1 for the random initial condition. For the modified payoffs, all treatments had a significant and positive effect.

ACKNOWLEDGMENTS. The authors thank Andrew Gelman for advice on statistical modeling and Winter Mason and Daniel Goldstein for helpful conversations.

1. Axelrod R, Hamilton WD (1981) The evolution of cooperation. *Science* 211:1390–1396.
2. Ledyard J (1995) Public goods: A survey of experimental research. *The Handbook of Experimental Economics*, eds Kagel JH, Roth AE (Princeton University Press, Princeton, NJ), pp 111–194.
3. Chaudhuri A (2011) Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature. *Exp Econ* 14:47–83.
4. Fischbacher U, Gächter S, Fehr E (2001) Are people conditionally cooperative? Evidence from a public goods experiment. *Econ Lett* 71:397–404.
5. Fehr E, Gächter S (2000) Cooperation and punishment in public goods experiments. *Am Econ Rev* 90:980–994.
6. Sefton M, Shupp R, Walker J (2007) The effect of rewards and sanctions in provision of public goods. *Econ Inq* 45:671–690.
7. Kossinets G, Watts DJ (2006) Empirical analysis of an evolving social network. *Science* 311:88–90.
8. 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.
9. Rand D, Arbesman S, Christakis N (2011) Dynamic social networks promote cooperation in experiments with humans. *Proc Natl Acad Sci USA* 108:19193–19198.
10. Pacheco JM, Traulsen A, Nowak MA (2006) Active linking in evolutionary games. *J Theor Biol* 243:437–443.
11. Kossinets G, Watts D (2009) Origins of homophily in an evolving social network1. *Am J Sociol* 115:405–450.
12. Newcomb T (1961) *The Acquaintance Process* (Holt, Rinehart & Winston, New York, NY).
13. Feld S (1981) The focused organization of social ties. *Am J Sociol* 86:1015–1035.
14. Granovetter M (1985) Economic action and social structure: The problem of embeddedness. *Am J Sociol* 91:481–510.
15. Felmlee D, Sprecher S, Bassin E (1990) The dissolution of intimate relationships: A hazard model. *Soc Psychol Q* 53:13–30.
16. Hauk E (2003) Multiple prisoner's dilemma games with (out) an outside option: An experimental study. *Theory Decis* 54:207–229.
17. Hauk E, Nagel R (2001) Choice of partners in multiple two-person prisoner's dilemma games. *J Conflict Resolut* 45:770–793.
18. Coricelli G, Fehr D, Fellner G (2004) Partner selection in public goods experiments. *J Conflict Resolut* 48:356–378.
19. List J (2006) Friend or foe? A natural experiment of the prisoner's dilemma. *Rev Econ Stat* 88:463–471.
20. Slonim R, Garbarino E (2008) Increases in trust and altruism from partner selection: Experimental evidence. *Exp Econ* 11(2):134–153.
21. Page T, Putterman L, Unel B (2005) Voluntary association in public goods experiments: Reciprocity, mimicry and efficiency*. *Econ J* 115:1032–1053.
22. Cinyabuguma M, Page T, Putterman L (2005) Cooperation under the threat of expulsion in a public goods experiment. *J Public Econ* 89:1421–1435.
23. Suri S, Watts DJ (2011) Cooperation and contagion in web-based, networked public goods experiments. *PLoS ONE* 6:e16836.
24. Traulsen A, Semmann D, Sommerfeld RD, Krambeck HJ, Milinski M (2010) Human strategy updating in evolutionary games. *Proc Natl Acad Sci USA* 107:2962–2966.
25. Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393:440–442.
26. Newman M (2003) The structure and function of complex networks. *SIAM Rev* 45: 167–256.
27. Gelman A, Hill J (2007) *Data Analysis Using Regression and Multilevel/Hierarchical Models* (Cambridge Univ Press, New York), Vol 1.
28. McPherson M, Smith-Lovin L, Cook J (2001) Birds of a feather: Homophily in social networks. *Annu Rev Sociol* 27:415–444.
29. Skyrms B, Pemantle R (2000) A dynamic model of social network formation. *Proc Natl Acad Sci USA* 97:9340–9346.
30. Immorlica N, Lucier B, Rogers B (2010) Cooperation in anonymous dynamic social networks. *Proceedings of the 11th ACM Conference on Electronic Commerce* (Association for Computing Machinery, New York), pp 241–242.
31. Santos FC, Pacheco JM, Lenaerts T (2006) Cooperation prevails when individuals adjust their social ties. *PLOS Comput Biol* 2:e140.
32. Hanaki N, Peterhansl A, Dodds P, Watts D (2007) Cooperation in evolving social networks. *Manage Sci* 53:1036–1050.
33. Bilancini E, Boncinelli L (2009) The co-evolution of cooperation and defection under local interaction and endogenous network formation. *J Econ Behav Organ* 70(1-2): 186–195.
34. Eshel I, Cavalli-Sforza LL (1982) Assortment of encounters and evolution of cooperativeness. *Proc Natl Acad Sci USA* 79:1331–1335.
35. Cavalli-Sforza LL, Feldman MW (1983) Paradox of the evolution of communication and of social interactivity. *Proc Natl Acad Sci USA* 80:2017–2021.
36. Ostrom E, Walker J, Gardner R (1992) Covenants with and without a sword: Self-governance is possible. *Am Polit Sci Rev* 86:404–417.
37. Maier-Rigaud FP, Martinsson P, Staffiero G (2010) Ostracism and the provision of a public good: Experimental evidence. *J Econ Behav Organ* 73:387–395.
38. Mason WA, Watts DJ (2009) Financial incentives and the performance of crowds. *Proceedings of the ACM SIGKDD Workshop on Human Computation* (Association for Computing Machinery, New York), pp 77–85.
39. Horton J, Rand D, Zeckhauser R (2011) The online laboratory: conducting experiments in a real labor market. *Exp Econ* 14:399–425.
40. Paolacci G, Chandler J, Ipeirotis PG (2010) Running experiments on Amazon Mechanical Turk. *Judgm Decis Mak* 5:411–419.
41. Mason W, Watts DJ (2012) Collaborative learning in networks. *Proc Natl Acad Sci USA* 109:764–769.

Cooperation and assortativity with dynamic partner updating
Supporting Information

Jing Wang[1], Siddharth Suri[2], Duncan J. Watts[2]

[1] IOMS Department, Leonard N. Stern School of Business
New York University, 44 West 4th Street, New York, NY 10012, USA
[2] Microsoft Research New York City
1290 Avenue of the Americas, 6th Floor, New York, NY 10104

To whom correspondence should be addressed:
E-mail: jwang5@stern.nyu.edu, suri@microsoft.com, duncan@microsoft.com

1 The Experiment

This section provides additional details on the Social Networking Game experiment, conducted on Amazons Mechanical Turk (AMT). All participants were recruited on AMT by posting a HIT for the experiment, entitled "The Social Networking Game (fun + bonus!)", a neutral title that was accurate without disclosing the purpose of the experiment. Before launching the experiment, we submitted to and complied with Yahoo!s internal human subjects review process (see Section 6 for terms of service). All data collected in the experiment could be associated only with participants Amazon Turker ID, not with any personally-identifiable information; thus all players remained anonymous.

1.1 Amazon's Mechanical Turk

Amazon's Mechanical Turk (`http://mturk.com`) is a web-based labor market with large volumes of small tasks (called human intelligence tasks, or HITs) offered for small reward. Typical tasks include image labeling, sentiment analysis, or classification of URLs, and wages are typically on the order of \$0.01–\$0.10 per HIT. Mechanical Turk is becoming increasingly popular with behavioral science researchers, in part because it allows experiments to be run faster and more cheaply, and in part because it provides access to a potentially much broader cross-section of the population than is typical of university-based lab experiments [1, 2, 3, 4, 5]. Accordingly, we posted each of our experimental sessions as a HIT and recruited workers to participate in the experiment. After seeing a screenshot of the experiment with explanatory text, workers could choose to accept the HIT, at which point the work was officially assigned to them and they could begin participating in the study. General overviews of AMT and its application to behavioral science experiments are available in [4, 6].

1.2 Recruiting Participants

Using the Mechanical Turk API, we restricted our participant pool to US workers. A total of 108 unique players participated in the 94 games reported in this study. Of these, 57 reported their gender as male and 51 reported as female. The median reported age was 30 years old, with quartiles of 26 and 37.5 years old. The modal response for annual household income was \$30k-\$40k, and the modal response for the highest level of education attained was high school graduate. The majority of participants played in 20 or fewer games, though there were a few that played in as many as 70 (see Figure 1). As noted in the main text, players received an average of \$1.63 per game; however, there was considerable variation around this figure, depending on the level of cooperation in the game and the specific actions of the player. Fig. S2 therefore shows the full distribution of compensation per game per player over all games.

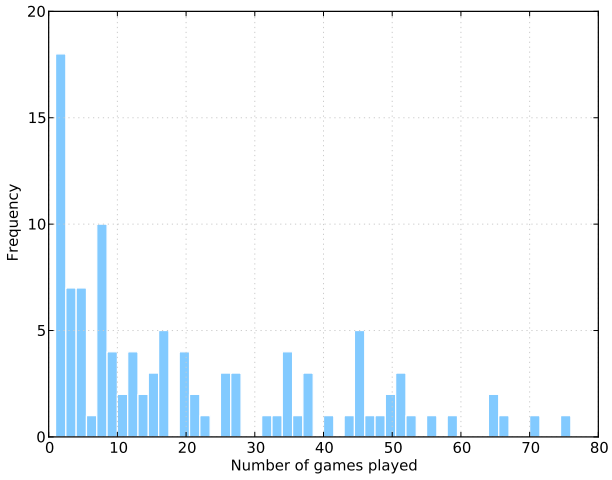

Figure 1: Histogram of number of games played by each panel member

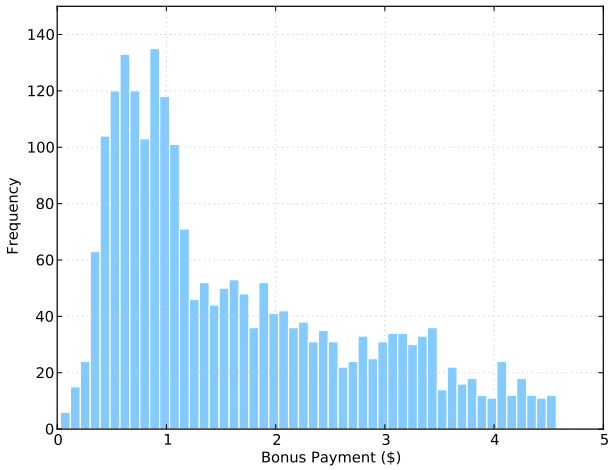

Figure 2: Histogram of compensation per player per game

Consistent with previous networked experiments run on AMT [7, 5] we solved the problem of recruiting many players simultaneously by creating a virtual "waiting room." Once arriving participants accepted the HIT, read the instructions, and passed the quiz (see below for instructions and player quiz) they were directed to a screen informing them that the experiment had not yet filled, along with how many remaining players were required. Once all positions had been filled, participants in the waiting room were informed that the game was about to commence. Also consistent with previous work [7, 5], posting the HIT on AMT would have been insuffi-

cient to fill networks of size $n = 24$ in a reasonable time, resulting in participants abandoning the waiting room and the HIT being terminated. To alleviate this problem, we ran a series of experiments with simple networks comprising $n = 4$ subjects. After participating in at least one experiment, the subjects were asked to report their demographics, and then given the opportunity to opt-in to a standing panel of experienced players who were familiar with the rules of the game, and were willing to be contacted for future games. The experiments reported in this manuscript were then conducted by recruiting from this standing panel.

1.3 Description of The Game

After accepting the HIT and agreeing to the terms of use, participants were provided with the following instructions given for the experiments where $T' = 7$, $R' = 4$, $P' = 1$, $S' = -1$. Instructions for the modified payoffs were suitably changed.

Welcome to the Social Networking Game!

Because the amount of money you can earn depends on your decisions in the game, it is important that you read these instructions with care. At the end of these instructions, you will be asked to take and pass a quiz to ensure that you understand the instructions. If you answer any questions incorrectly, you will get a second chance. If you answer a question incorrectly twice, you will not be allowed to play the game and will not receive payment for the HIT.

Overview

In the Social Networking Game you will be placed in a network with 23 other Turkers. However, you will only be linked to some of these players, which we call your "network neighbors". It is important to understand that your payoff from the game is determined only by how you interact with these network neighbors (i.e. Turkers directly linked to you in the network).

Throughout the game, we will show you who are your neighbors and who are not. The initial network will be specified by us; however, you might have opportunity to adjust your

links as the game proceeds.

This game consists of 12 “rounds”. During each round, you and other players will be able choose a strategy. As a result of these choices, you will receive a “payoff” which depends on both your chosen strategy and the strategies of your neighbors. Your total payoff for the game is the sum of your payoffs from each round.

Every so often, you might have the opportunity to “rewire” links, either severing a link to one of your existing neighbors, or proposing a new link to someone who is not currently your neighbor. In this way you can change with whom you play the game. How frequently you get to rewire links, and how many links you get to rewire at each opportunity, may change from game to game. Sometimes you will get no opportunity while at other times you will get many. Regardless, we will tell you the frequency of rewiring opportunities and the number of rewirings per opportunity at the start of each game.

During the game we, will not report your earnings in terms of dollars and cents but rather in terms of points. At the end of the game the total amount of points you have earned will be converted to dollars at the rate of 2 points = 1 cent. The amount you earn from the game will be the bonus for this HIT.

Regardless of how you perform in the game, you will earn the base rate of 25 cents for this HIT by correctly/successfully passing the quiz at the end of these instructions.

How the game works

Please read the following rules carefully. You will be tested on your understanding, and must answer all questions correctly, before you can play.

1. In each round, you will see a “Choose your Strategy” message. There are two strategies you can choose between, which we label “A” and “B”. After you choose your strategy, click the “Submit” button. You will have 45 seconds to make your decision. If you do not submit your decision before the end of a round, you will not get paid for that particular

round and the system will play the strategy that you chose on the previous round. If you do not submit a strategy choice in the first round, the system will make a choice for you.

2. There are two types of players from your view: those who are linked to you in the network (neighbors), and those who are not (non neighbors). Your payoff in each round will depend only on your choice of strategy and the choices of your network neighbors. In any given round, you will not be able to see the choices of your neighbors until you have submitted your choice and the round has completed. However, you will be able to see the past choices of all 23 other players (your neighbors and also your non-neighbors) for up to the 5 previous rounds with the strategies for more recent rounds on the right.

3. Your payoff for each round is determined by your strategy and the strategies of your neighbors. Your payoff from one interaction (with one of your neighbors) can be summarized by the payoff table below:

		Your Neighbor's Strategy	
		A	**B**
Your Strategy	**A**	4, 4	-1, 7
	B	7, -1	1, 1

In each cell, the first number represents the payoff to you, and the second number represents the payoff to your neighbor. Another way to understand this table is as follows:

(a) If both you and your neighbor choose A, you receive 4 points each **[top-left cell]**;

(b) If you choose A and your neighbor chooses B, you receive -1 point (i.e. you lose 1 point) and your neighbor receives 7 points **[top-right cell]**;

(c) If you choose B and your neighbor chooses A, you receive 7 points and your neighbor receives -1 point **[bottom-left cell]**;

(d) If both of you choose B, you receive 1 point each **[bottom-right cell]**.

4. Your payoff for each round is summed over your interactions with all of your neighbors. To illustrate, consider the following examples:

   Example 1: Suppose you choose A, and that you have three neighbors, X1, X2, and X3, where: X1 chooses A, X2 chooses B and X3 chooses A.

   Your total payoff for this round will be 4 - 1 + 4 = 7.

   Example 2: Now suppose that in Example 1 (i.e. with the same three neighbors making the same choices as above) you had chosen B instead of A: Your payoff would be 7 + 1 + 7 = 15.

   Example 3: Next, suppose that you have chosen A and your three neighbors have made the following choices: X1 chooses A, X2 chooses A and X3 chooses A.

   Your payoff would be 4 + 4 + 4 = 12.

   Example 4: Next, suppose that have chosen B and your three neighbors have made the following choices: X1 chooses B, X2 chooses B and X3 chooses B.

   Your payoff would be 1 + 1 + 1 = 3.

   Example 5: If you have no neighbors in a given round, your payoff in that round will be 0 regardless of what strategy you choose.

5. Your total payoff is summed over all 12 rounds. The payoff for each person in the network is calculated in the same way. Notice that you might get negative payoff in some round. If at the end, your total payoff is negative, we will still pay you the base rate but you will simply receive zero bonus. Given that you might be able to adjust your links (i.e. choose your neighbors) from time to time, this is unlikely to happen.

6. During the game, you might have the opportunity to "rewire" links. Each opportunity will be signaled by the title "Link Rewiring". You can then rewire your links by following the

two phases described below:

(a) **Phase 1**:

- You can "Delete" a link between you and an existing neighbor. Deletions are unilateral, meaning that you can delete an edge without the permission of your neighbor.
- You can "Propose" a new link with a player who is not currently your neighbor. In order to form a new link, a proposal must be agreed to by the other player, as explained in Phase 2, below.

You will have 45 seconds to make all your rewiring decisions. If you do not submit all your decisions before the phase ends, you will be automatically redirected to the next phase.

(b) **Phase 2**:

You can "Accept" as many link proposals as you want from other players. After that, you will be linked to them. There will be 30 seconds in this phase. If you fail to submit before the phase ends, you will not form new links with any player.

7. After all rounds have been completed your payoffs over all 12 rounds will be summed, and you will be paid a bonus corresponding to your total points earned.

8. Upon your completion of this game, you will be asked to complete a short survey.

1.4 Participant Quiz

Finally, participants were required to pass a quiz, thus demonstrating that they had understood the instructions.

Quiz

To make sure you have read and understood the instructions, you must answer the following questions correctly. If you answer any questions incorrectly, you will get a second chance. If you answer a question incorrectly twice, you will not be allowed to play the game and will not receive payment for the HIT. The answers to all of the questions below are in terms of points.

1. Suppose you have 3 neighbors, if you play B and everyone else plays A, what would your payoff be?

2. Suppose you have 3 neighbors, if everyone plays B what would your payoff be?

3. Suppose you have 3 neighbors, if everyone plays A what would your payoff be?

4. What would be your bonus if you have a negative total payoff in points in the end?

2 Game Specification

To specify the game, we note that a two player PD is defined in terms of the following payoff matrix

	C	D
C	R, R	S, T
D	T, S	P, P

where $T > R > P > S$, and for iterated games the inequality $2R > T + S$ must also be satisfied. For an n-player PD, we make the standard assumption [8, 9, 10, 11, 12, 13, 14] that each player plays the same strategy with respect to all n partners, from which it follows that player i's payoff π_i^u for playing strategy $u \in C, D$ is,

$$\begin{aligned} \pi_i^C &= n_c R + (n - n_c)S - n\gamma \\ \pi_i^D &= n_c T + (n - n_c)P - n\gamma \end{aligned}$$

where n_c is the number of cooperators among i's neighbors, and γ is the cost of maintaining a tie [12]. Rewriting these expressions to incorporate tie cost into the game payoffs, we have

$$\begin{aligned}\pi_i^C &= n_c(R-\gamma)+(n-n_c)(S-\gamma) = n_cR' + (n-n_c)S' \\ \pi_i^D &= n_c(T-\gamma)+(n-n_c)(P-\gamma) = n_cT' + (n-n_c)P'\end{aligned}$$

To ensure that cooperators had reason to sever ties with defectors, but that pairs of defectors had reason to retain ties, we set $S < \gamma < P$ (where we modify the second inequality in the second series of experiments), obtaining finally the modified payoffs $T' = 7$, $R' = 4$, $P' = 1$, $S' = -1$. In the second series of experiments the payoffs were $T' = 7$, $R' = 4$, $P' = -1$, $S' = -5$. Both sets of payoffs satisfy the standard iterated PD conditions. Players accumulated points over the course of the game, and were then compensated in money upon completion, where the exchange rate was \$0.005 per point. At this exchange rate, players earned an average of \$1.63 per game, and as much as \$4.57 per game (see S2 for full distribution).

For the initial set of payoffs, with the exception of one treatment, we conducted four to five realizations each for $r = 1, 3, 6$, $k = 1, 3, 5$, and the two choices of initial networks. For comparison, we also conducted four to five realizations of each corresponding static network (i.e. same initial condition, but with no rewiring allowed). Finally, we conducted an additional 12 experiments for the modified payoffs—four each for $r = 1$ and $k = 1, 5, 23$. Thus, we conducted a total 94 experiments with 24 players each, where the order of the experiments was randomized (within each set of payoffs) to negate any possible learning effects. Moreover, to avoid selection effects, players were only informed of the parameters after they had entered the experiment and passed the quiz, but before the game itself commenced.

3 Screenshots

The following screenshots illustrate a typical players view at three stages of the play:

1. During the strategy update part of a round (Fig. S3)

2. During the partner deletion and partner proposal part of a partner updating turn (Fig. S4)

3. During the proposal acceptance part of a partner updating turn (Fig. S5)

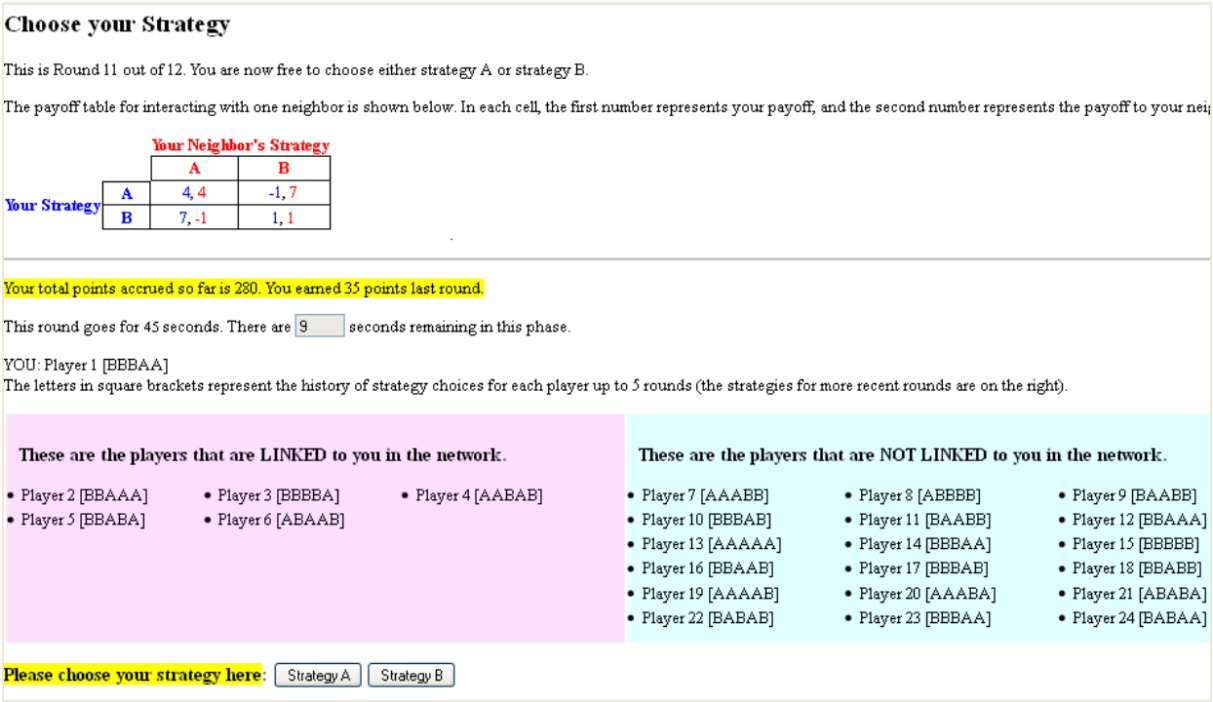

Figure 3: Screenshot during strategy update

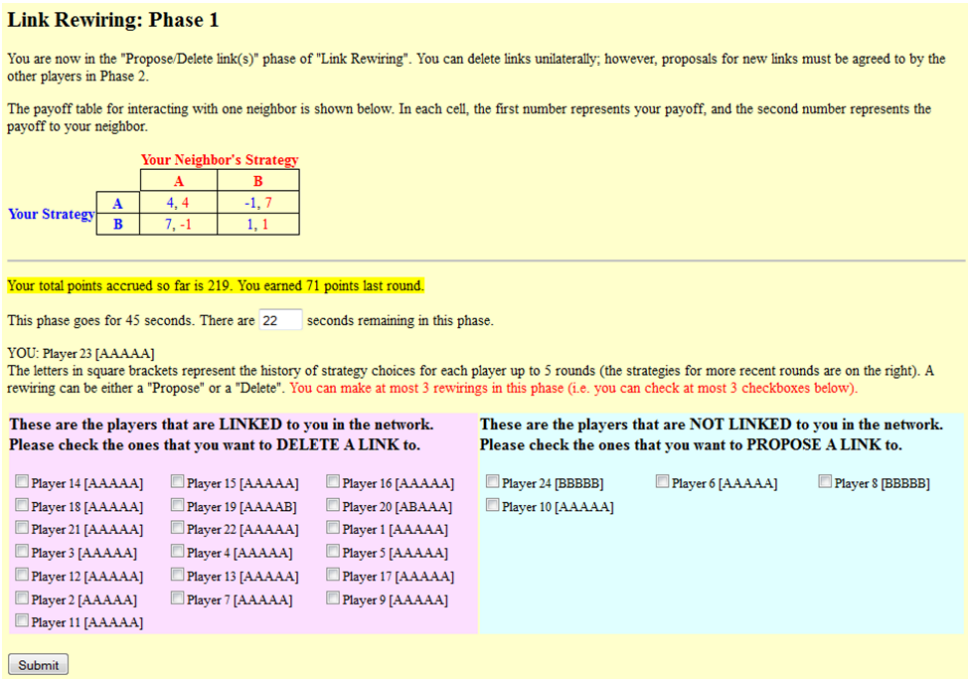

Link Rewiring: Phase 1

You are now in the "Propose/Delete link(s)" phase of "Link Rewiring". You can delete links unilaterally; however, proposals for new links must be agreed to by the other players in Phase 2.

The payoff table for interacting with one neighbor is shown below. In each cell, the first number represents your payoff, and the second number represents the payoff to your neighbor.

		Your Neighbor's Strategy	
		A	B
Your Strategy	A	4, 4	-1, 7
	B	7, -1	1, 1

Your total points accrued so far is 219. You earned 71 points last round.

This phase goes for 45 seconds. There are 22 seconds remaining in this phase.

YOU: Player 23 [AAAAA]
The letters in square brackets represent the history of strategy choices for each player up to 5 rounds (the strategies for more recent rounds are on the right). A rewiring can be either a "Propose" or a "Delete". You can make at most 3 rewirings in this phase (i.e. you can check at most 3 checkboxes below).

These are the players that are LINKED to you in the network. Please check the ones that you want to DELETE A LINK to.

Player 14 [AAAAA] Player 15 [AAAAA] Player 16 [AAAAA]
Player 18 [AAAAA] Player 19 [AAAAB] Player 20 [ABAAA]
Player 21 [AAAAA] Player 22 [AAAAA] Player 1 [AAAAA]
Player 3 [AAAAA] Player 4 [AAAAA] Player 5 [AAAAA]
Player 12 [AAAAA] Player 13 [AAAAA] Player 17 [AAAAA]
Player 2 [AAAAA] Player 7 [AAAAA] Player 9 [AAAAA]
Player 11 [AAAAA]

These are the players that are NOT LINKED to you in the network. Please check the ones that you want to PROPOSE A LINK to.

Player 24 [BBBBB] Player 6 [AAAAA] Player 8 [BBBBB]
Player 10 [AAAAA]

Submit

Figure 4: Screenshot during partner deletion and new partner proposal

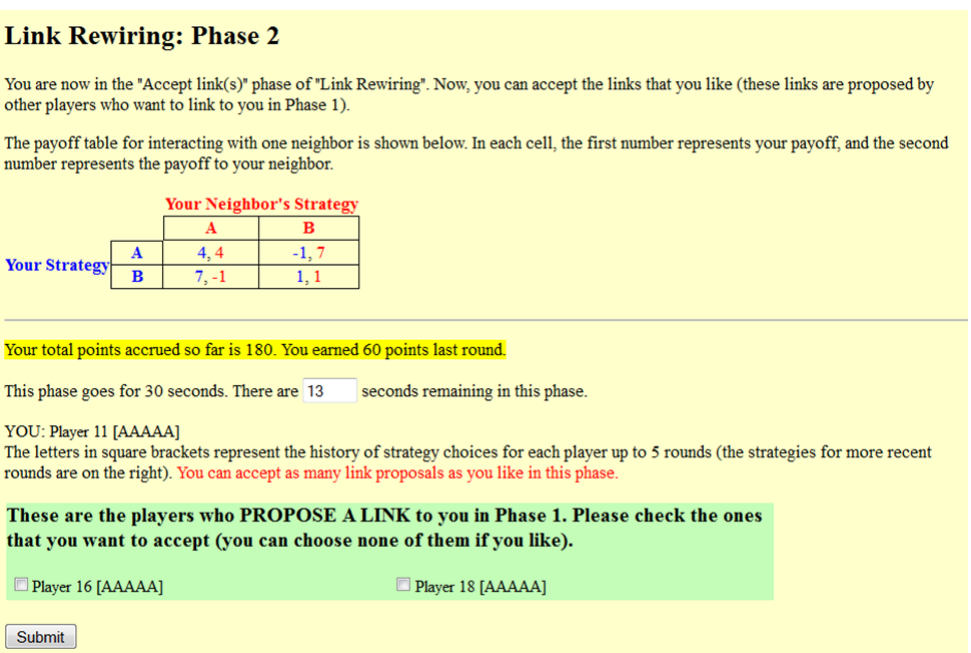

Link Rewiring: Phase 2

You are now in the "Accept link(s)" phase of "Link Rewiring". Now, you can accept the links that you like (these links are proposed by other players who want to link to you in Phase 1).

The payoff table for interacting with one neighbor is shown below. In each cell, the first number represents your payoff, and the second number represents the payoff to your neighbor.

		Your Neighbor's Strategy	
		A	B
Your Strategy	A	4, 4	-1, 7
	B	7, -1	1, 1

Your total points accrued so far is 180. You earned 60 points last round.

This phase goes for 30 seconds. There are 13 seconds remaining in this phase.

YOU: Player 11 [AAAAA]
The letters in square brackets represent the history of strategy choices for each player up to 5 rounds (the strategies for more recent rounds are on the right). You can accept as many link proposals as you like in this phase.

These are the players who PROPOSE A LINK to you in Phase 1. Please check the ones that you want to accept (you can choose none of them if you like).

Player 16 [AAAAA] Player 18 [AAAAA]

Submit

Figure 5: Screenshot during proposal acceptance/rejection

As described in the instructions given in Section 1.3, players were given 45 seconds to update their strategies, 45 seconds to sever and propose links, and 30 seconds to accept proposed links. If a player did not update their strategy in the time allotted, the strategy they used in the previous round was used by the system (a random choice was used if this happened in round 1). Only trials where *both* of the following two conditions, which were also used in prior work [7], were included in this analysis. First, over 90% of the strategy updates had to have been made by humans. Second, more than 50% of the strategy updates for each node were conducted by the human controlling that node.

4 Equilibrium Analysis

4.1 Original Payoffs

The following theorem says that, for the original payoffs, the Nash equilibrium is for all nodes to defect and add as many edges as possible. In general, it may not be possible for each node to add an edge to *every* node it is not already connected to. For example, nodes may only be allowed one link update for the whole experiment but have degree $d << n$ at the beginning. So the equilibrium network need not be unique, but the equilibrium actions of each player in term of whether to cooperate/defect or add/delete links is unique. Corollary 1 which follows the theorem below shows which of our treatments had a unique Nash equilibrium network.

Theorem 1. *In a finitely repeated PD with $P > 0$, η rounds, and partner updating turns every r rounds (where $r|\eta$), the Nash equilibrium is for all players to defect during each strategy update round, to accept all link proposals, and to propose as many links as possible.*

Proof. Let $s_1, \ldots, s_\eta$ denote the strategy update rounds. Assume that the nodes are allowed to update up to k links every r strategy update rounds. We will break up the strategy update rounds and partner updating turns into phases. Define phase p_0 to be the first r strategy update rounds.

Define phase p_i, $0 < i \leq \frac{\eta}{r} - 1$ to be partner update turn e_i and the r strategy update rounds that immediately follow, $s_{ir+1} \ldots s_{(i+1)r}$.

For the base case of our backwards induction, consider the last phase which is $p_{\frac{\eta}{r}-1}$. Since the players are playing a PD and D is the dominant strategy, the Nash equilibrium for $s_{\eta-r+1} \ldots s_\eta$ is for all players to defect for each of those rounds. Knowing this and since $P > 0$, during the second phase of partner update turn $e_{\frac{\eta}{r}-1}$ all players will accept all link proposals. Since all link proposals will be accepted, all nodes will defect, and $P > 1$, during the first phase of partner update turn $e_{\frac{\eta}{r}-1}$ each node will make as many link proposals as possible.

Next assume the theorem holds for phases $p_{\frac{\eta}{r}-1} \ldots p_{i+1}$ and consider phase p_i. By the induction hypothesis all players will defect for $s_{(i+2)r+1}, \ldots s_\eta$ and since the players are playing a PD with D as the dominant strategy they will also defect for rounds $s_{(i+1)r+1} \ldots s_{(i+2)r}$. Knowing this and since $P > 0$, during the second phase of e_i all players will accept all link proposals. Since all proposals will be accepted, all nodes will defect, and $P > 1$, during the first phase of partner update turn e_i each node will make as many link proposals as possible. Finally, since all players defect in rounds $s_{r+1}, \ldots, s_\eta$ and since this is a PD with a dominant strategy of D, all players will defect for p_0. □

The following corollary, also for the original payoffs, says that for the $3/1$ and $5/1$ treatments the unique Nash equilibrium network is the complete network.

Corollary 1. *If the initial network has minimum degree d, and $(\frac{\eta}{r} - 1)k \geq n - 1 - d$, the unique Nash equilibrium network is the complete network.*

Proof. Theorem 1 says that players will add as many edges as possible every partner updating turn. Each player will experience $(\frac{\eta}{r} - 1)k$ edge actions over the course of a game. So a node of minimum initial degree d would need $(\frac{\eta}{r} - 1)k \geq n - 1 - d$ edge actions to add edges to all nodes it was not initially connected to. □

4.2 Modified Payoffs

The following theorem says that, for the modified payoffs, the Nash equilibrium is for all nodes to defect and delete as many edges as possible. In general, it may not be possible for each node to delete every edge incident on it. For example, nodes may only be allowed one link update for the whole experiment but have degree $d > 1$ at the beginning. So the equilibrium network need not be unique, but the equilibrium actions of each player in terms of whether to cooperate/defect or add/delete links is unique. Corollary 2 which follows the theorem below shows the unique Nash equilibrium network for the specific parameters our subjects experienced.

Theorem 2. *In a finitely repeated PD with $P < 0$, η rounds, and partner updating turns every r rounds (where $r|\eta$), the Nash equilibrium is for all players to defect during each strategy update round and to delete as many edges as possible during each partner updating turn.*

The proof works by dividing the strategy update rounds into phases. Define a phase to be a partner updating turn and the ensuing r strategy update rounds. Then one does backward induction on the phases showing that nodes only delete links and defect during each phase.

Proof. Let $s_1, \ldots, s_\eta$ denote the strategy update rounds. Assume that the nodes are allowed to update up to k links every r strategy update rounds. We will break up the strategy update rounds and partner updating turns into phases. Define phase p_0 to be the first r strategy update rounds. Define phase p_i, $0 < i \leq \frac{\eta}{r} - 1$ to be partner update turn e_i and the r strategy update rounds that immediately follow, $s_{ir+1} \ldots s_{(i+1)r}$.

For the base case of our backwards induction, consider the last phase which is $p_{\frac{\eta}{r}-1}$. Since the players are playing a PD and D is the dominant strategy, the Nash equilibrium for $s_{\eta-r+1} \ldots s_\eta$ is for all players to defect for each of those rounds. Knowing this and since $P < 0$, players will delete as many edges as possible during edge update turn $e_{\frac{\eta}{r}-1}$.

Next assume the theorem holds for phases $p_{\frac{\eta}{r}-1} \dots p_{i+1}$ and consider phase p_i. By the induction hypothesis all players will defect for $s_{(i+2)r+1}, \dots s_\eta$ and since the players are playing a PD with D as the dominant strategy they will also defect for rounds $s_{(i+1)r+1} \dots s_{(i+2)r}$. Knowing this and since $P < 0$ all players will delete as many edges as possible during e_i which concludes the inductive step. Finally, since all players defect in $s_{r+1}, \dots, s_\eta$ and since this is a PD with a dominant strategy of D, all players will defect for p_0. □

The following corollary, also for the modified payoffs, says that for the $1/1$, $5/1$ and $23/1$ treatments the unique Nash equilibrium network is the empty network.

Corollary 2. *If the initial network has maximum degree d, and $(\frac{\eta}{r} - 1)k \geq d$, the unique Nash equilibrium network is the empty network.*

Proof. Theorem 2 says that players will delete as many edges as possible every partner updating turn. Each player will experience $(\frac{\eta}{r} - 1)k$ edge actions over the course of a game. So a node of maximum initial degree d would need $(\frac{\eta}{r} - 1)k \geq d$ edge actions to delete all incident edges. □

5 Supplementary Results

5.1 Player Earnings, Original Payoffs

Fig. S6A shows average earnings per player as a function of partner update rate for cliques (dashed lines) and random (solid lines) initial conditions, for $r = 1, 3, 6$, $k = 0, 1, 3, 5$ (symbols indicate different different values of k ($k = 0$ diamonds, $k = 1$ triangles, $k = 3$ circles, $k = 5$ squares; error bars are ± 2 standard errors). Fig. S6B shows the differences between the treatment effects and the corresponding control conditions. Error bars are ± 2 standard errors, corresponding to 95% CI. Hence when zero is not included in the CI, the null hypothesis that the treatment cannot be distinguished from zero can be rejected at the 5% level. The error bars and statistical tests were computing using the non-nested, multi-level model described in the Materials and Methods section of the main text.

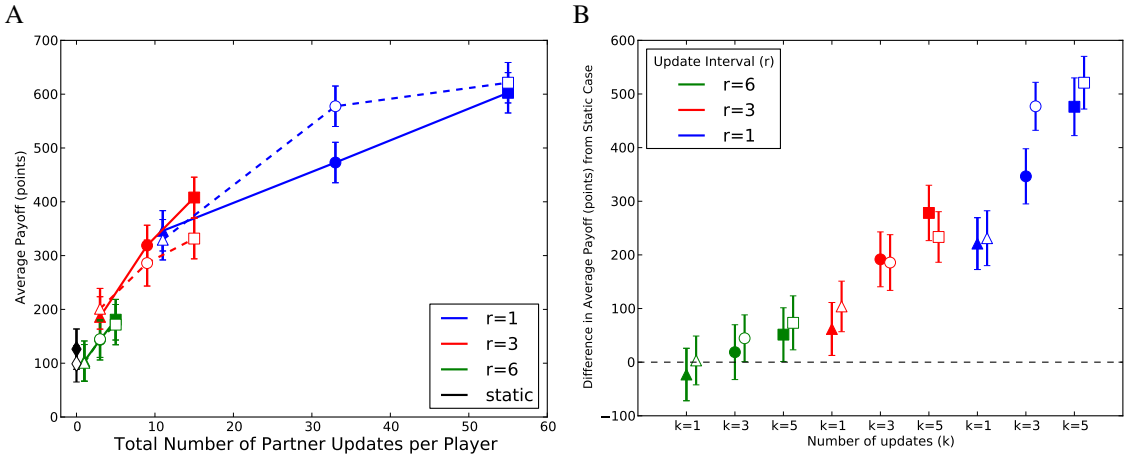

Figure 6: Average Earnings, Original Payoffs

5.2 Results for $r = 3, 6$, Original Payoffs

Figs. S7 and S8 correspond to Fig. 3 in the main text for $r = 3, 6$.

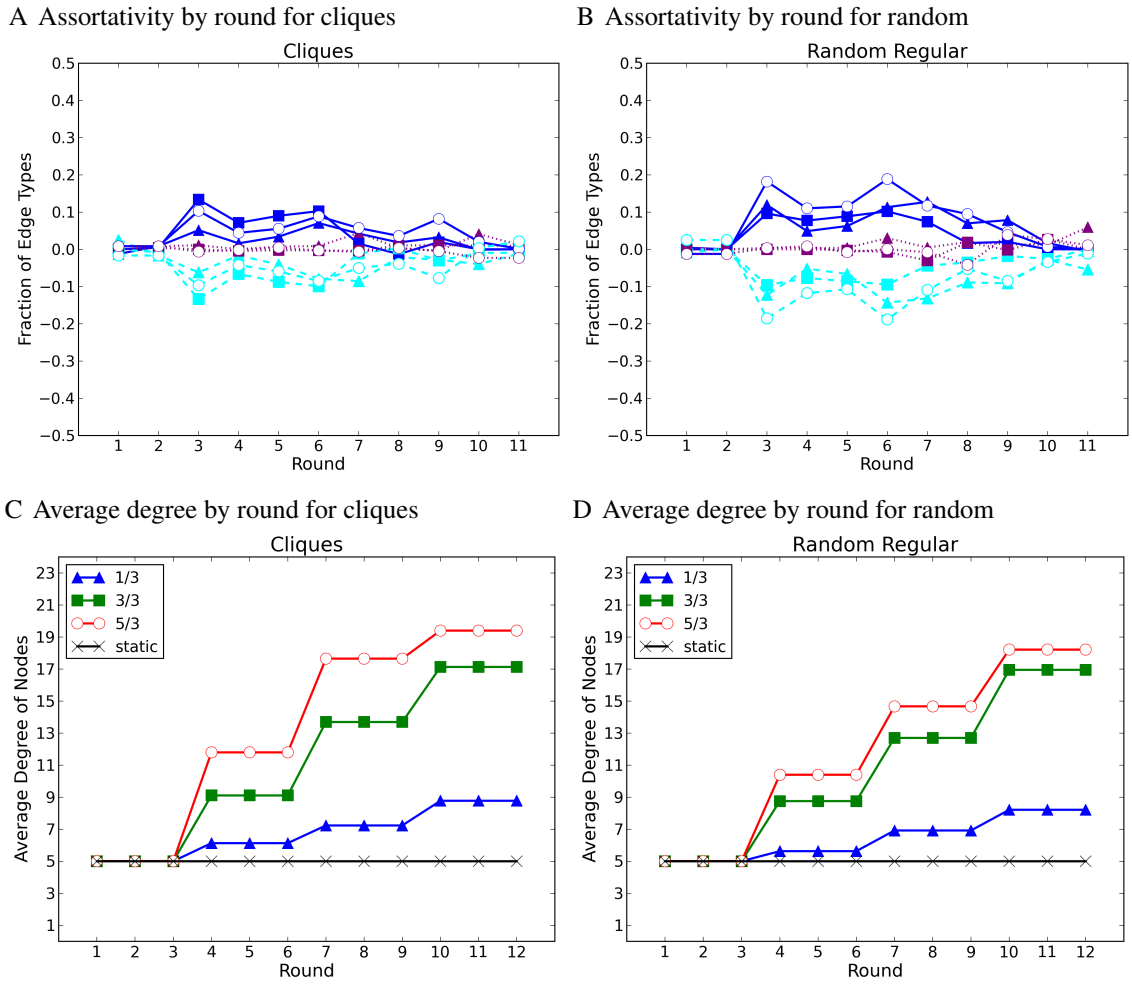

Figure 7: Average assortativity and degree for cliques and random initial conditions, $r = 3$.

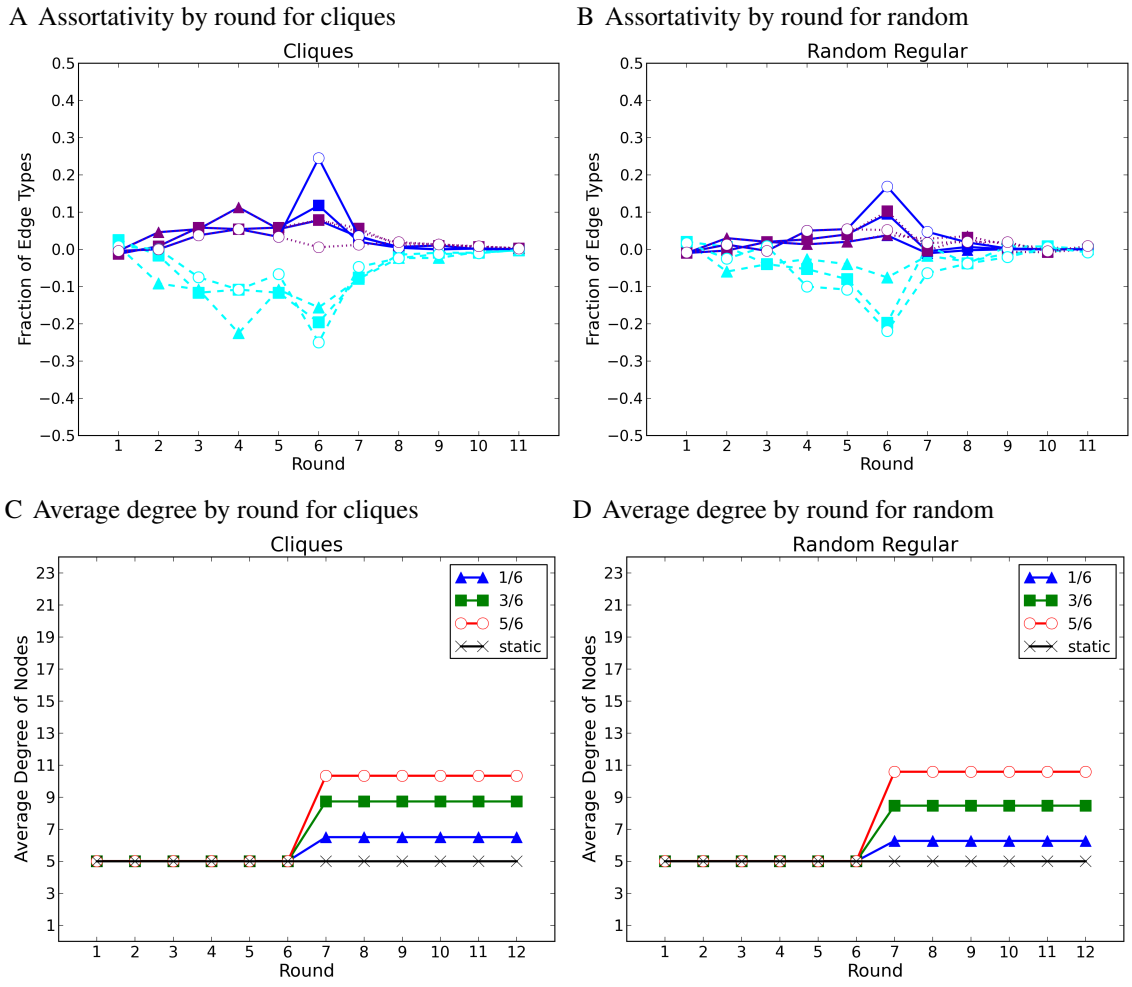

Figure 8: Average assortativity and degree for cliques and random initial conditions, $r = 6$.

5.3 Original Payoffs: Tie Deletions, Proposals and Acceptances

As stated in the main text, for the original payoffs $T' = 7$, $R' = 4$, $P' = 1$, $S' = -1$, we observed relatively little tie deletion. Figs. S9–S11 confirm this, showing average number of ties proposed, accepted, and deleted. Almost all ties that were proposed were accepted, suggesting that defectors did not try to propose ties to cooperators, and very few ties were ever deleted. Table 1 corroborates this finding by showing that most proposals were from cooperators to other cooperators and from defectors to other defectors and these were both accepted over

Table 1: Tie Proposals, Deletions, and Acceptances for Original Payoffs for $r = 1, k = 1$. To clarify notation, C $\rightarrow$ C proposals corresponds to cooperators proposing ties to other cooperators with whom they are not currently partnered, whereas C $\rightarrow$ C deletions corresponds to cooperators deleting links with existing partners who are also cooperators. Fraction proposals and fraction deletions denote the fraction of the total edge actions of each type, where the sum of these two columns is bounded above by 1. % acceptances shows, of the proposals of a given type, the percentage that were accepted.

tie type	fraction deletions	fraction proposals	% acceptances
C $\rightarrow$ C	0	0.52	97.8%
C $\rightarrow$ D	0.10	0.02	93.8%
D $\rightarrow$ C	0	0.10	40.4%
D $\rightarrow$ D	0	0.24	90.5%

90% of the time. It also shows that deletions only made up 10% of the partner update actions.

A Cliques, $r = 1$, $k = 1$

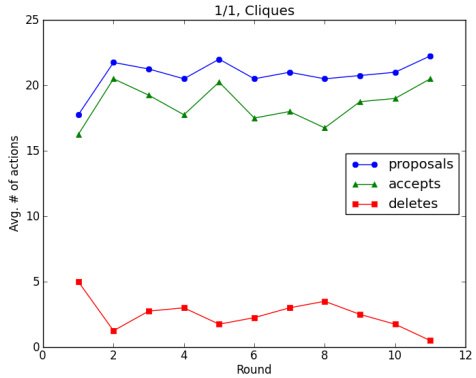

B Random Reg., $r = 1$, $k = 1$

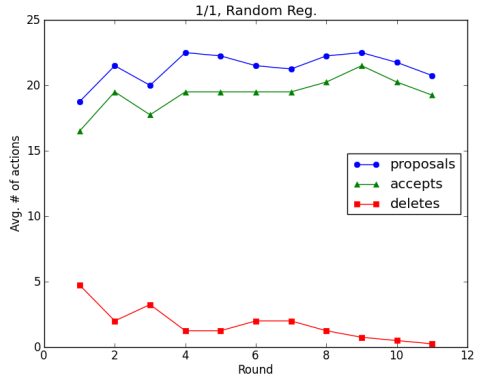

C Cliques, $r = 1$, $k = 3$

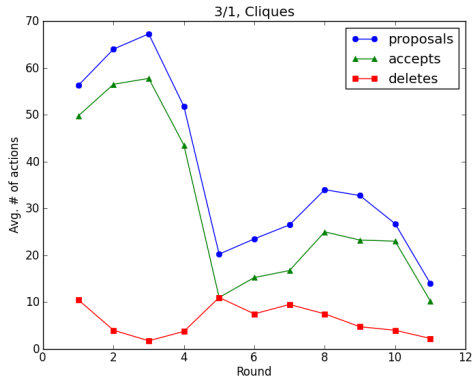

D Random Reg., $r = 1$, $k = 3$

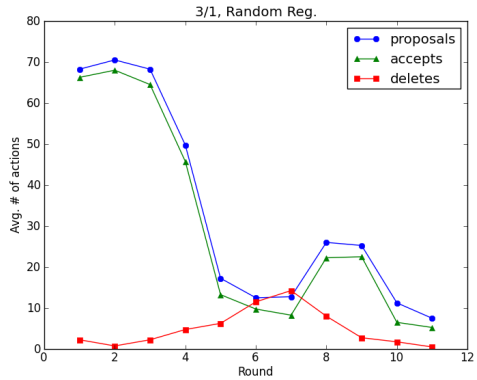

E Cliques, $r = 1$, $k = 5$

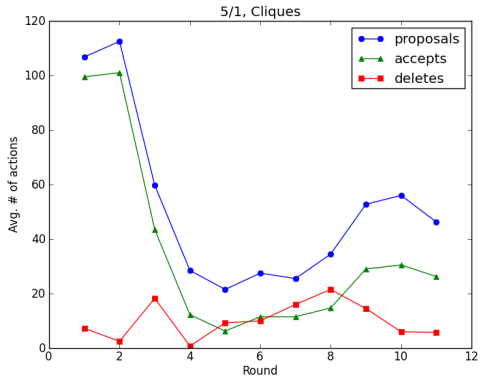

F Random Reg., $r = 1$, $k = 5$

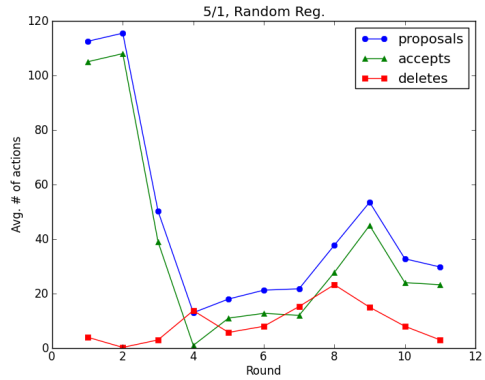

Figure 9: Number of proposed, accepted, and deleted edges by round for $r = 1$, $k = 1, 3, 5$, original payoffs.

A Cliques, $r = 1$, $k = 1$

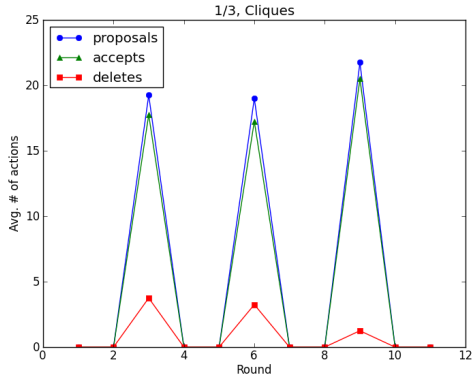

B Random Reg., $r = 1$, $k = 1$

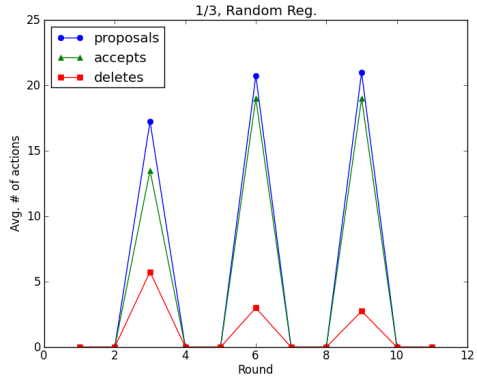

C Cliques, $r = 1$, $k = 3$

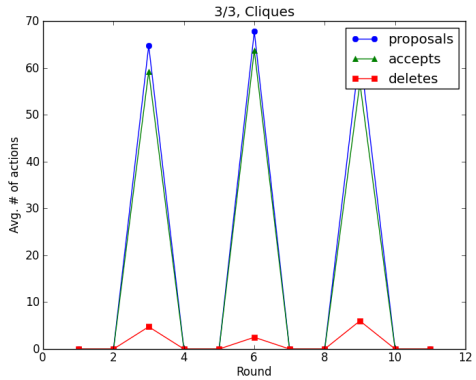

D Random Reg., $r = 1$, $k = 3$

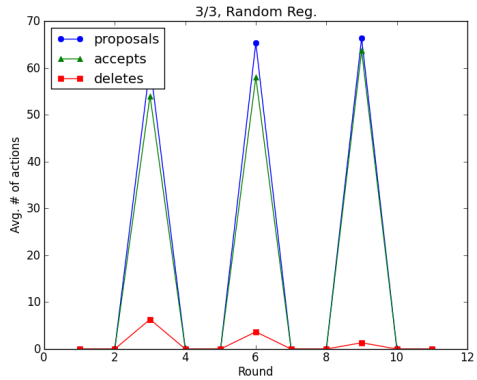

E Cliques, $r = 1$, $k = 5$

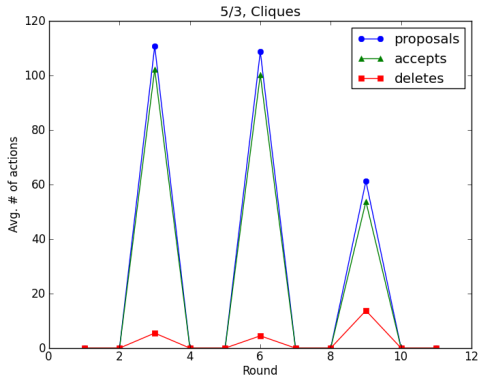

F Random Reg., $r = 1$, $k = 5$

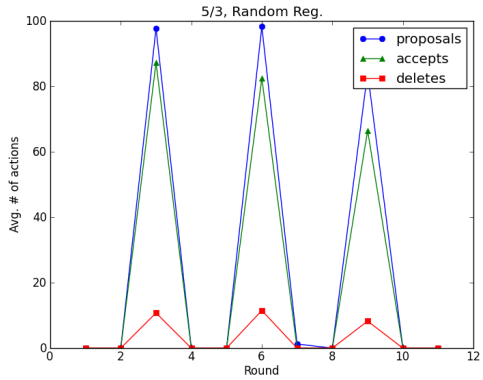

Figure 10: Number of proposed, accepted, and deleted edges by round for $r = 3$, $k = 1, 3, 5$, original payoffs

A Cliques, $r = 1$, $k = 1$

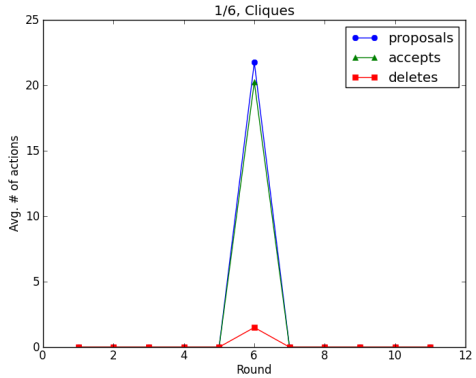

B Random Reg., $r = 1$, $k = 1$

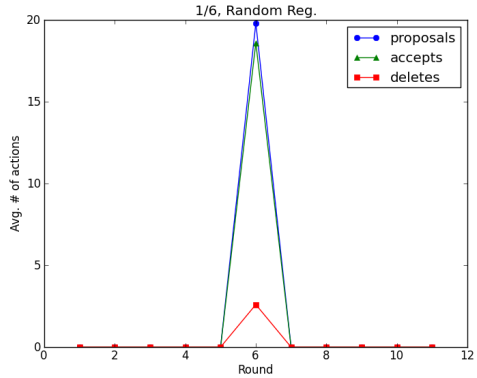

C Cliques, $r = 1$, $k = 3$

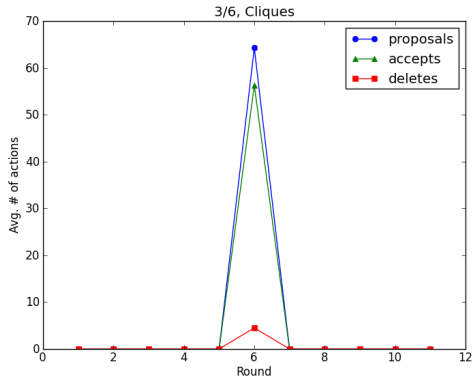

D Random Reg., $r = 1$, $k = 3$

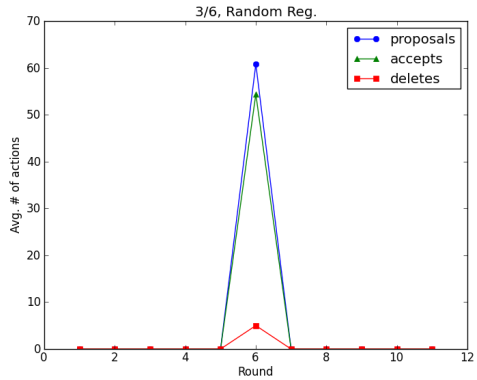

E Cliques, $r = 1$, $k = 5$

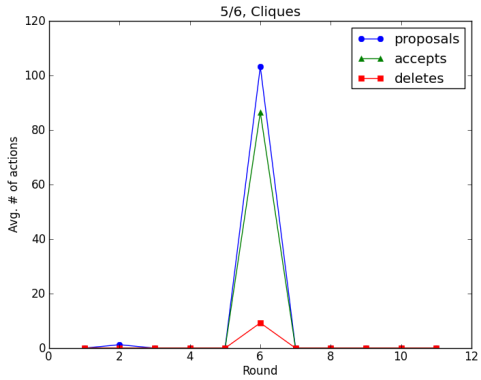

F Random Reg., $r = 1$, $k = 5$

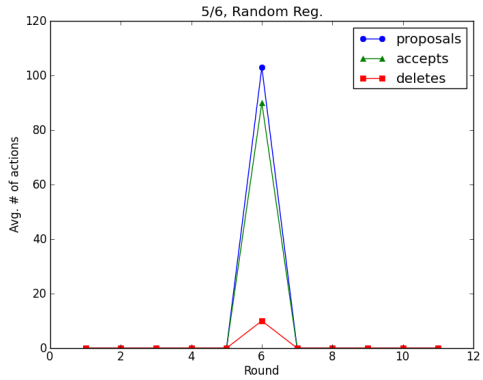

Figure 11: Number of proposed, accepted, and deleted edges by round for $r = 6$, $k = 1, 3, 5$, original payoffs

5.4 Defection Cascade Example, Original Payoffs

Figure S12 is an illustrative example of the dynamics under the original payoffs. The initial topology was 4 disjoint cliques each consisting of 6 nodes and $r = 1, k = 5$. Early on nodes were largely cooperating and adding edges as quickly as possible. During round 7 there was a group of 4 defectors. Instead of the cooperators severing their ties to these 4 defectors, the cooperators started to defect. As one can see by looking at the later rounds, this resulted in a defection cascade.

5.5 Player Earnings, Modified Payoffs

Fig. S13 (top) shows average earnings per player as a function of partner update rate for $r = 1$, $k = 0, 1, 5, 23$ (error bars are ± 2 standard errors). Fig. S13 (bottom) shows the differences between the treatment effects and the corresponding control condition, symbols indicate different different values of k ($k = 1$ triangles, $k = 5$ circles, $k = 23$ squares). Error bars are ± 2 standard errors, corresponding to 95% CI. Hence when zero is not included in the CI, the null hypothesis that the treatment cannot be distinguished from zero can be rejected at the 5% level. The error bars and statistical tests were computing using the non-nested, multi-level model described in the Materials and Methods section of the main text.

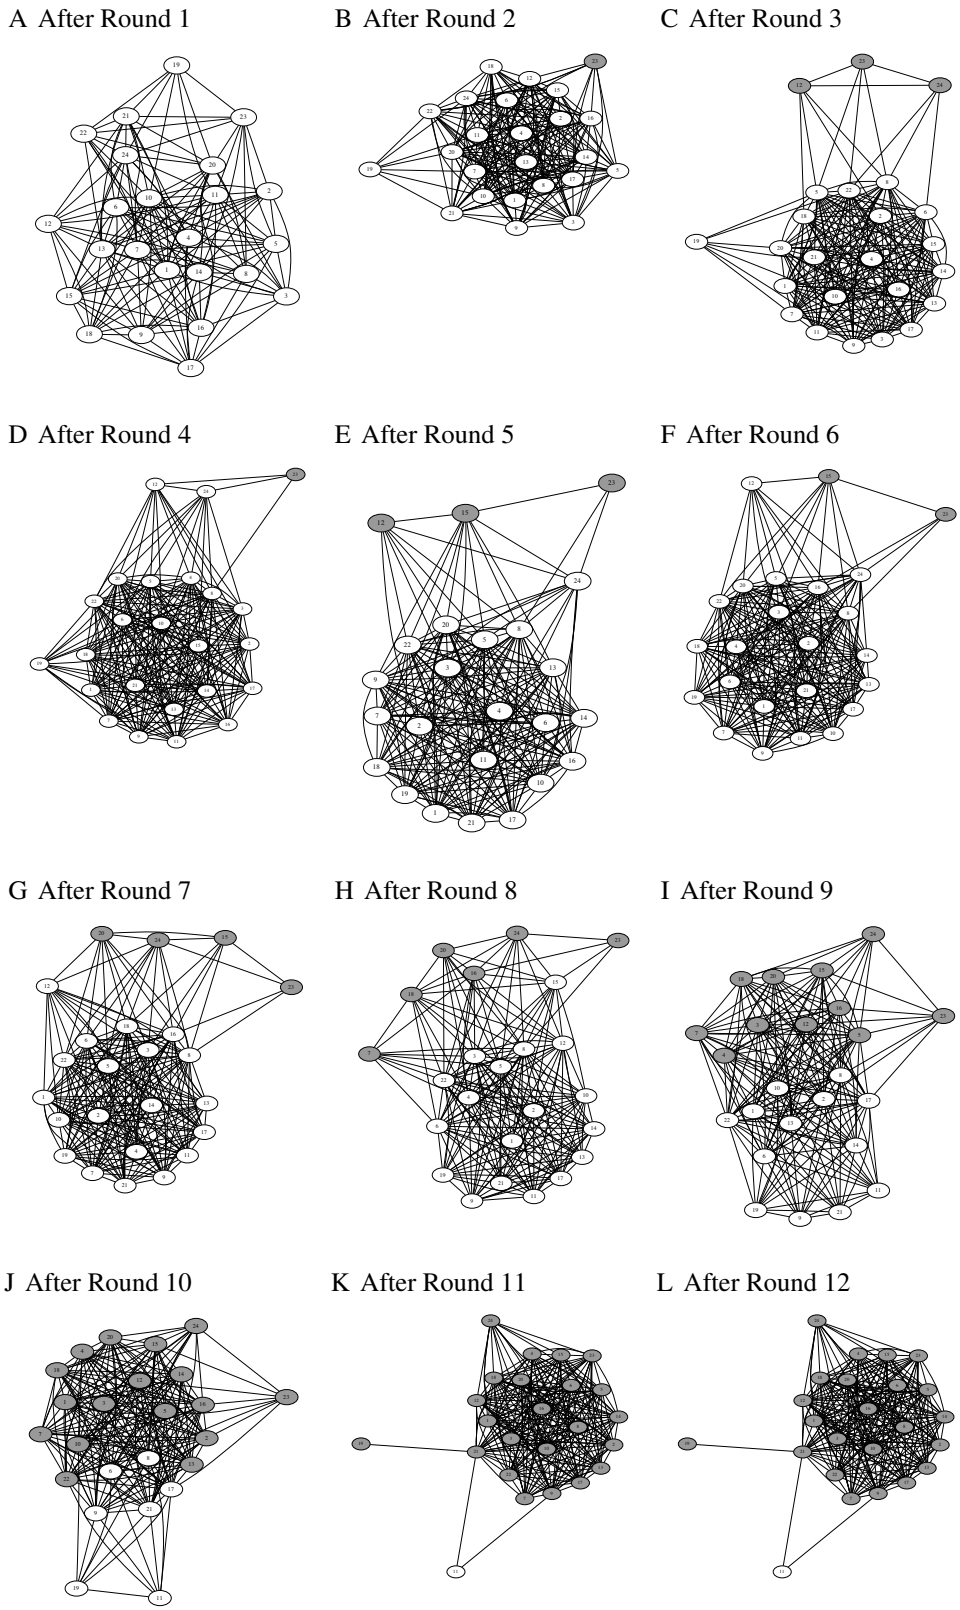

Figure 12: Gray nodes indicate defectors. White nodes indicate cooperators.

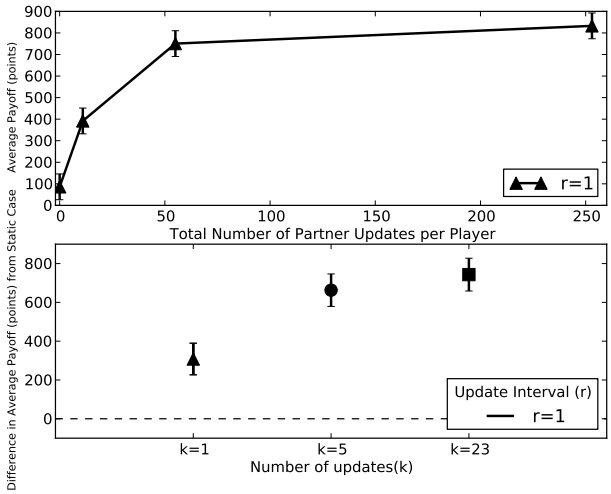

Figure 13: Average Earnings, Modified Payoffs

5.6 Modified Payoffs: Tie Deletions, Proposals and Acceptances

For the modified payoffs, $T' = 7$, $R' = 4$, $P' = -1$, $S' = -5$, the situation is quite different. Figure 14 shows that in the modified payoffs a substantial number of edge deletions occurred. These results confirm our conclusion in the main text that for original payoffs assortativity arises mostly out of cooperators avoiding defectors, whereas for the modified payoffs tie deletion plays an important role. Table 2 shows that most proposals were from cooperators to other cooperators. It also shows that there were roughly twice as many deletions as in the original payoffs (22% vs. 10%).

Table 2: Tie Proposals, Deletions, and Acceptances for Modified Payoffs for $r = 1, k = 1$. Fraction proposals and fraction deletions denote the fraction of the total edge actions of each type. The sum of these two columns is bounded above by 1. % acceptances shows, of the proposals of a given type, the percentage that were accepted.

tie type	fraction deletions	fraction proposals	% acceptances
C $\rightarrow$ C	0	0.70	94.73%
C $\rightarrow$ D	0.16	0.003	100%
D $\rightarrow$ C	0	0.04	0%
D $\rightarrow$ D	0.06	0.01	25%

A Cliques, $r = 1$, $k = 1$

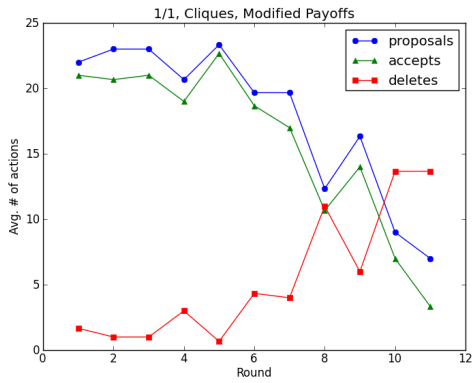

B Cliques, $r = 1$, $k = 3$

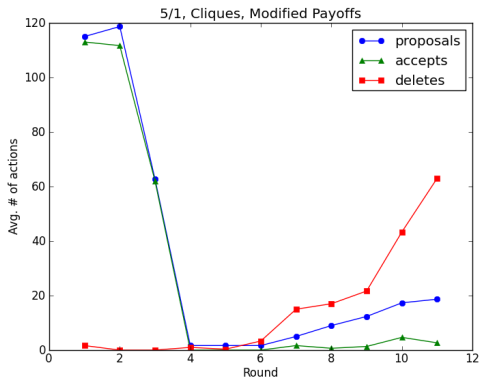

C Cliques, $r = 1$, $k = 4$

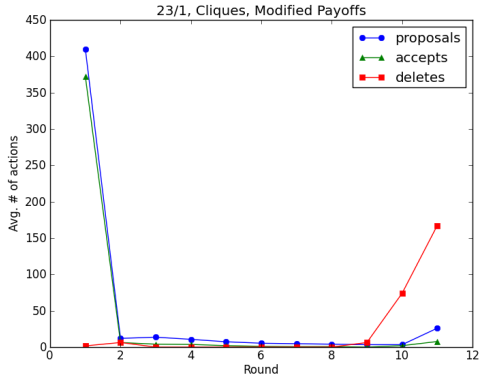

Figure 14: Number of proposed, accepted, and deleted edges by round for $r = 1$, $k = 1, 5, 23$, modified payoffs

5.7 Modified Payoffs: Sustained Cooperation via Implicit Punishment

Figure S15 is an illustrative example of the dynamics under the modfied payoffs. The initial topology was 4 disjoint cliques each consisting of 6 nodes and $r = 1, k = 5$. In the modified payoffs players added edges as quickly as possible until the graph essentially became a clique. Cooperation levels stayed at 100% until round 9 when one node defected. All but one of the links to that node was cut during the partner updating turn just after round 9. There were a few more defectors during round 10 who got isolated. These are examples of the cooperators implicitly punishing the defectors by deleting links. Finally in round 12, the end game effect kicked in and nodes defected.

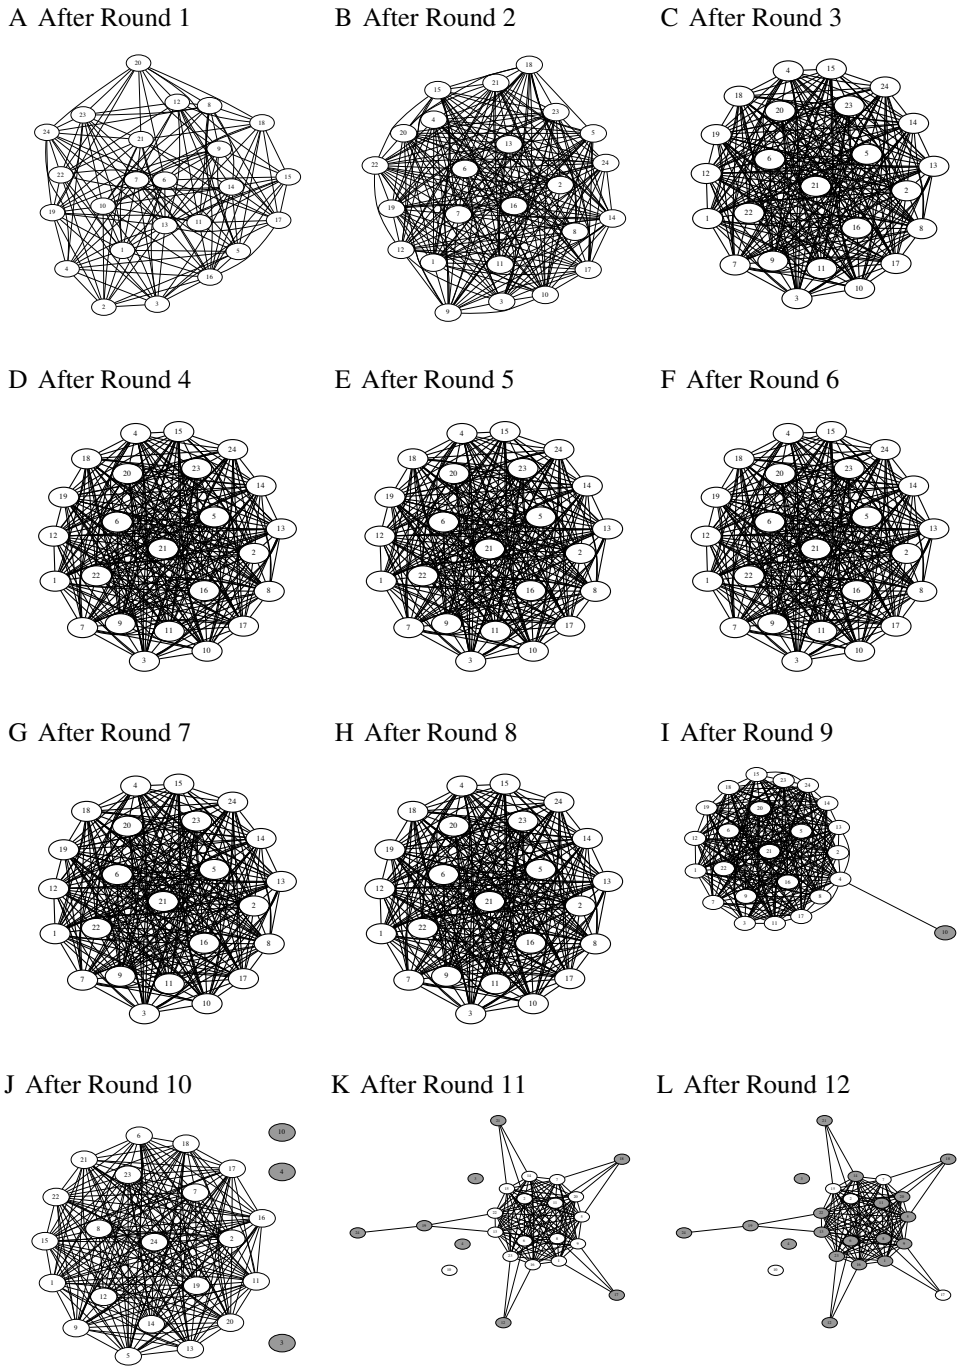

Figure 15: Gray nodes indicate defectors. White nodes indicate cooperators.

6 Terms of Use Agreement

As stated above, our experiment was reviewed and approved by Yahoo! Labs' Human Subjects Review (HSR) process. Correspondingly, before participating all subjects were required to read and agree to the following terms of use agreement, equivalent to an Informed Consent Statement.

The Social Networking Game: Terms of Use

You will be paid $0.25 for completing the HIT plus a bonus depending on your performance in the game. By clicking the "I Agree" button below you affirm that: 1) you have read, understood, and agree to the Yahoo! Research Social Networking Game Terms of Use and Social Networking Game description and agree to comply with and be bound by its terms, and 2) you are 18 years of age or older. If you do not agree to these Terms of Use, click the "I Don't Agree" button below to complete the HIT and get paid $0.25 for completing the quiz. Doing so will result in skippping the game. If you have any questions at any time, please contact: Social-Networking-Game team at Yahoo! Research, 111 W. 40th St., New York, NY, or by email at socialnetwork-game@yahoo-inc.com.

YAHOO! RESEARCH SOCIAL NETWORKING GAME TERMS OF USE

THE FOLLOWING GAME IS INTENDED FOR PLAY IN THE FIFTY UNITED STATES AND THE DISTRICT OF COLUMBIA ONLY (COLLECTIVELY, THE "ELIGIBILITY AREA") AND SHALL ONLY BE CONSTRUED AND EVALUATED ACCORDING TO UNITED STATES LAW. DO NOT PARTICIPATE IN THIS GAME IF YOU ARE NOT LOCATED IN, AND A LEGAL RESIDENT OF, THE ELIGIBILITY AREA. PLEASE DO NOT PROCEED ANY FURTHER IF YOU DO NOT AGREE TO THESE TERMS OF USE OR IF PARTICIPATION IN THE GAME IS CONTRARY TO THE LAWS OF YOUR COUNTRY. THIS GAME IS VOID WHERE PROHIBITED OR OTHERWISE RESTRICTED IN A MANNER

NOT SATISFIED BY THE PUBLICATION OF THESE OFFICIAL RULES. SEE BELOW FOR ADDITIONAL ELIGIBILITY RESTRICTIONS. NO PURCHASE REQUIRED. A PURCHASE DOES NOT IMPROVE YOUR CHANCE OF WINNING.

1. Welcome to the Yahoo! Research Social Networking Game ("Project"). This Project is a game of skill, not a game of chance. By participating in the Project you are entering into a legally binding agreement with Yahoo!, Inc., ("Yahoo!," "we," "our," and "us"). This agreement is comprised solely of these Terms of Use ("Agreement" or "Terms"), including anything explicitly incorporated by reference. If you do not agree to these Terms, please do not participate.

2. The Project is offered to individuals registered as "workers" with Amazon, Inc.'s "Mechanical Turk" service (`http://www.mturk.com/mturk/welcome`).

3. Your participation in the Project as a worker is governed by Amazon, Inc.'s Mechanical Turk's conditions of use (`http://www.mturk.com/mturk/conditionsofuse`) in addition to the following Yahoo! terms:

   (a) Description of Project. The Project is intended to collect data on how well people play this game.

   (b) Eligibility. You must be a legal resident of the Eligibility Area and at least 18 years or older. Void in overseas US territories, possessions, commonwealths and military installations, and where prohibited by law.

   (c) How to Enter. The Project begins on Aug 9, 2011 and ends on September 9, 2011. To participate in the Project, you must review the instructions contained in the sections entitled Overview and How the Game Works ("Description") and successfully pass a short quiz demonstrating your understanding of these instructions. If you

successfully pass the quiz as described in the Description, you will be allowed to enter the experiment.

(d) How to play. In addition to these Terms, additional rules regarding game play and scoring (which are fully incorporated into these Terms by reference) can be found in the Description. After you complete your participation in the Project, you will be asked to complete a short survey.

(e) Payment. You will be paid $0.25 plus a bonus (as described in the Description) depending on your skill level in the Project. All payments will be made to You through the Mechanical Turk service as detailed in the Mechanical Turk conditions of use.

(f) Work Product/Ownership. You agree to perform the tasks provided in the Project and to be compensated for the completion of each task as set forth in E above. You also agree that Yahoo!, and not You, shall own all work product from your participation in the Project.

(g) Relationship of the Parties. The Parties are independent contractors. Nothing in these Terms shall be construed as creating any agency, partnership, or other form of joint enterprise between the Parties and neither Party may create any obligations or responsibilities on behalf of the other Party.

(h) Termination.

   i. By You. You may terminate Your participation in the Project by clicking the Return HIT button at any time.

   ii. By Yahoo!. We may suspend or terminate the Project at any time, with or without notice, for any reason or no reason. In the event of such termination, Yahoo! will pay You for all tasks fully completed by You in accordance with

these Terms prior to termination.

(i) Contact. If you have any questions at any time, please contact Social-Networking-Game team at Yahoo! Research, 111 W. 40th St., New York, NY, or by email at socialnetwork-game@yahoo-inc.com

(j) Confidentiality. You will not disclose or use Yahoo!'s Confidential Information. "Confidential Information" means any information disclosed or made available to You by Yahoo!, directly or indirectly, whether in writing, orally or visually, other than information that: (a) is or becomes publicly known and generally available other than through Your action or inaction or (b) was already in Your possession (as documented by written records) without confidentiality restrictions before you received it from Yahoo!. Confidential Information includes, but is not limited to, all information contained within the Project, these Terms, the Policies, and any other technical or programming information Yahoo! discloses or makes available to you.

(k) Indemnity. You will defend, indemnify and hold harmless Yahoo! Inc., and its affiliated companies, ("Indemnified Parties") from and against any and all claims, liabilities, losses, costs, and expenses, including reasonable attorneys' fees, which the Indemnified Parties suffer as a result of claims that arise from or relate to your activities under or in connection with this Agreement, including but not limited to claims that allege or arise from: (i) a violation a third party's right of privacy, or infringement of a third party's copyright, patent, trade secret, trademark, or other intellectual property rights, (ii) any breach of your obligations, covenants, warranties or representations as set forth in this Agreement, including any breach of any applicable policies, (iii) any violation of applicable laws, rules, and regulations by you, including, without limitation, privacy laws, and (iv) any breach of this Agreement. You shall not enter into any settlement that affects any Indemnified Party's rights or

interest, admit to any fault or liability on behalf of any Indemnified Party, or incur any financial obligation on behalf of any Indemnified Party without that Indemnified Party's prior written approval.

(l) No Warranty. YOU EXPRESSLY AGREE TO THE FOLLOWING WARRANTY DISCLAIMER. YOU ARE PARTICIPATING IN THE PROJECT AT YOUR OWN RISK. YOU REPRESENT AND WARRANT THAT BY PARTICIPATING IN THIS PROJECT THAT YOU WILL COMPLY WITH ALL APPLICABLE LAWS. THE PROJECT AND EVERYTHING PROVIDED UNDER THIS AGREEMENT IS PROVIDED "AS IS." YAHOO! DOES NOT WARRANT THAT THE PROJECT WILL OPERATE UNINTERRUPTED OR ERROR-FREE. YAHOO AND ITS LICENSORS ARE NOT RESPONSIBLE FOR ANY CONTENT PROVIDED HEREUNDER. TO THE EXTENT ALLOWED BY LAW, YAHOO! AND ITS LICENSORS MAKE NO WARRANTY OF ANY KIND, WHETHER EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING WITHOUT LIMITATION WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT. YAHOO! MAKES NO WARRANTY AND NO REPRESENTATION ABOUT THE AMOUNT OF MONEY YOU WILL EARN THROUGH THE PROGRAM. THIS WARRANTY DISCLAIMER SHALL APPLY TO THE MAXIMUM EXTENT PERMITTED BY LAW.

(m) Limitation of Liability. YOU EXPRESSLY AGREE TO THE FOLLOWING LIMITATION OF LIABILITY. YAHOO! WILL NOT BE LIABLE FOR ANY LOST PROFITS, COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY OTHER INDIRECT, SPECIAL, INCIDENTAL, EXEMPLARY, PUNITIVE OR CONSEQUENTIAL DAMAGES ARISING OUT OF OR IN CONNECTION WITH THIS AGREEMENT, HOWEVER CAUSED, AND UN-

DER WHATEVER CAUSE OF ACTION OR THEORY OF LIABILITY BROUGHT, EVEN IF YAHOO! HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. YAHOO! WILL NOT BE LIABLE FOR DIRECT DAMAGES IN EXCESS OF ANY AMOUNT THAT YAHOO! HAS ALREADY PAID TO YOU FOR YOUR PARTICIPATION IN THE PROJECT. IF YOU ARE DISSATISFIED WITH ANY ASPECT OF THE PROJECT, OR WITH ANY OF THESE TERMS OF USE, YOUR SOLE AND EXCLUSIVE REMEDY IS TO DISCONTINUE YOUR PARTICIPATION IN THE PROJECT. This limitation of liability shall apply to the maximum extent permitted by law.

(n) Privacy. The information that you provide based on your participation in this Project is subject to Yahoo!s privacy policy, which can be found at `http://info.yahoo.com/privacy/us/yahoo/details.html`.

(o) No Public Statements. You may not issue any press release or other public statement regarding the Agreement, Yahoo!, and/or Yahoo! Inc., its affiliates, or partners or advertisers without the prior written consent of an authorized person at Yahoo!.

References and Notes

[1] Mason WA, Watts DJ (2009) *Financial incentives and the performance of crowds* pp 77–85.

[2] Ipeirotis PG, Provost F, Wang J (2010) *Quality management on Amazon Mechanical Turk* pp 64–67.

[3] Paolacci G, Chandler J, Ipeirotis PG (2010) Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making* 5:411–419.

[4] Horton J, Rand D, Zeckhauser R (2010) The online laboratory: conducting experiments in a real labor market. *Experimental Economics* pp 1–27.

[5] Mason W, Watts DJ (2012) Collaborative learning in networks. *Proceedings of the National Academy of Sciences* 109:764–769.

[6] Mason W, Suri S (2011) Conducting behavioral research on amazons mechanical turk. *Behavior Research Methods* pp 1–23.

[7] Suri S, Watts D (2011) Cooperation and contagion in web-based, networked public goods experiments. *PloS One* 6:e16836.

[8] Eshel I, Cavalli-Sforza L (1982) Assortment of encounters and evolution of cooperativeness. *Proceedings of the National Academy of Sciences of the United States of America* 79:1331.

[9] Cavalli-Sforza L, Feldman M (1983) Paradox of the evolution of communication and of social interactivity. *Proceedings of the National Academy of Sciences* 80:2017.

[10] Santos F, Pacheco J, Lenaerts T (2006) Cooperation prevails when individuals adjust their social ties. *PLoS Computational Biology* 2:e140.

[11] Pacheco J, Traulsen A, Nowak M (2006) Active linking in evolutionary games. *Journal of theoretical biology* 243:437–443.

[12] Hanaki N, Peterhansl A, Dodds P, Watts D (2007) Cooperation in evolving social networks. *Management Science* 53:1036–1050.

[13] Bilancini E, Boncinelli L (2009) The co-evolution of cooperation and defection under local interaction and endogenous network formation. *Journal of Economic Behavior & Organization* 70:186–195.

[14] Immorlica N, Lucier B, Rogers B (2010) *Cooperation in anonymous dynamic social networks* (ACM), pp 241–242.