

Supporting Information

for

Dynamic Social Networks Promote Cooperation in Experiments with Humans

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

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

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

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

1. Details of the experimental setup.....	2
2. Additional analysis of reciprocity via cooperative action.....	15
3. Additional analysis of reciprocity via network rewiring	16
4. Additional analysis of network properties and cooperation	17
5. Supporting References	21

1. Details of the experimental setup

A total of 785 subjects participated in our incentivized economic game experiments. Subjects were recruited using the online labor market Amazon Mechanical Turk (AMT) (1-3). AMT is an online labor market in which employers contract with workers to complete short tasks for relatively small amounts of money. Workers receive a baseline payment and can be paid an additional bonus depending on their performance. Thus, incentivized experiments are easy to conduct using AMT: the baseline payment corresponds to the traditional ‘show-up fee,’ and the bonus payment is determined by the number of points earned during the experimental session.

Issues exist when running experiments online which do not exist in the traditional laboratory. Running experiments online naturally implies some loss of control, since the workers cannot be directly monitored as in the traditional lab; hence, experimenters cannot be certain that each observation is the result of a single person (as opposed to multiple people making joint decisions at the same computer), or that one person does not participate multiple times (although AMT goes to great lengths to try to prevent this, and, based on IP address monitoring, it seems to happen very infrequently). Moreover, the sample of subjects in AMT experiments is restricted to people that participate in online labor markets (although most physical lab studies are restricted to college undergraduates, who are also far from representative).

To address these potential concerns, recent studies have explored the validity of data gathered using AMT (for an overview, see (2)). Most pertinent to our study are two direct replications using economic games. The first shows quantitative agreement in contribution behavior in a repeated public goods game between experiments conducted in the physical lab and those conducted using AMT with approximately 10-fold lower stakes (3). The second replication again found quantitative agreement between the lab and AMT with 10-fold lower stakes, this time in cooperation in a one-shot Prisoner’s Dilemma (1).

Our experiments add another set of replications to the growing literature on AMT. In our random condition, we see the same behavior so often observed among a wide range of subject pools in multi-player cooperation games, namely, initial high levels of contribution which quickly decay over time. And the comparison between our random and fixed conditions is again consistent with behavior in the physical laboratory (4, 5), where fixed interaction structure does little to prevent the breakdown of cooperation.

Our participants interacted anonymously over the internet using custom software playable in a browser window. The initial environment consisted of a countdown timer of 15 minutes, at which time a “Go” button became visible and the participants clicked this to participate. Upon clicking, subjects were taken to a website external to AMT designed to implement our experiments. For each experiment, each subject was asked to perform a tutorial, after which the actual game would begin.

If a subject did not click “Go” and enter our custom website within 100 seconds, they were dropped from the game. If they did not complete the tutorial within 600 seconds, they were dropped. After 600 seconds from the beginning of the tutorial, all participants (who completed the tutorial) began to play. At any point during the game, if a subject was inactive for 180 seconds, they were warned about being dropped. If they still remained inactive after 360 seconds, they were dropped.

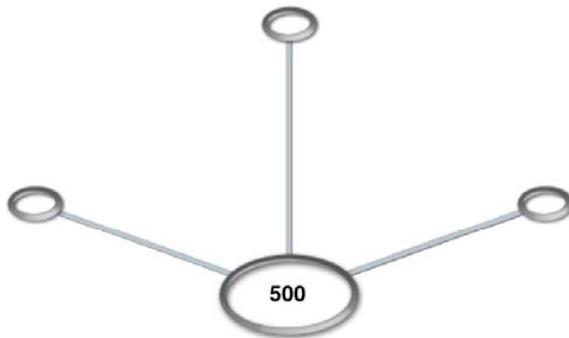
A total of 38 subjects were dropped at some point after the first round of game play, and this dropout rate did not vary significantly between the fluid condition and the other conditions (logistic regression clustered on session, all pairwise comparisons $p > 0.10$). An additional 39

subjects were dropped in the very first round of play. First-round dropout rates were somewhat higher in the strategic updating conditions compared to the random and fixed conditions (logistic regression clustered on session, $p=0.014$), likely because of the increased wait-time caused by the rewiring round. Importantly, however, first round dropout rates did not differ significantly between the viscous and fluid conditions (logistic regression clustered on session, $p=0.194$). Thus differential dropout rates are unlikely to explain the differences in behavior we observe in our experiment.

Once beginning the game, future interactions occurred with probability 0.8. To control for variation in game lengths across conditions, we pre-generated a set of 10 game lengths from a geometric distribution with success probability 0.8, and used the same set of 10 game lengths for the 10 games in each condition (as in (6, 7)). Thus, the differences across conditions we observe in our experiment cannot be explained by certain conditions having games that lasted for more or less time.

Below are screenshots from the initial description of the tutorial where rewiring is allowed (random and no rewiring simply eliminate the rewiring round, and if relevant, note that rewiring is randomized after each round). We also show the first of three practice rounds.

Round: Tutorial



How to play:

You will be playing this game with other Mechanical Turk workers.

At any particular point in the game you will be connected to some of the other players who are also playing.

The image to the left shows the players you are connected with.



Each of these players is represented by a small ring.



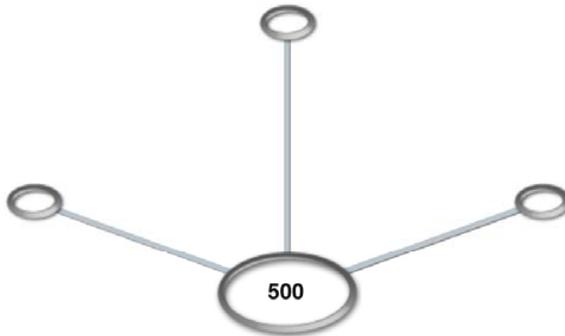
You are represented by the large ring.

In the game, you and the other players will make a number of decisions. These decisions will cause you to gain or lose points. You start with 500 points.

At the end of the game you will be paid a bonus of 1 cent for every 10 points in your account.

Next

Round: Tutorial



How to play:

The game will be played over a series of rounds.

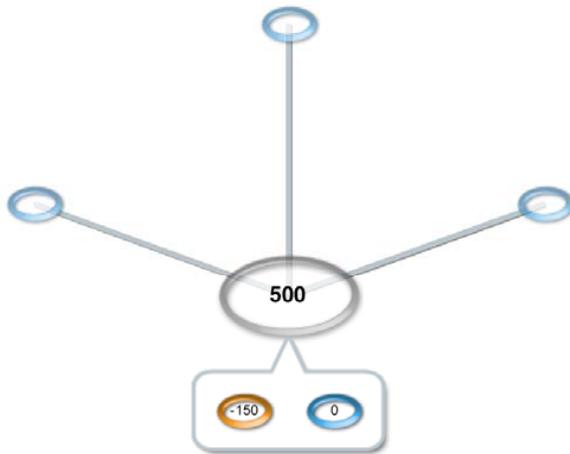
In every round you make a choice about whether to pay to give points to the other players you are connected to.

In some rounds you also make a choice about whether or not to be connected to a particular player.

After every round there is a 80%% chance that the game continues on to another round. There is a 20% chance that the game ends and you receive your payment.

We will now describe the game in more detail.

Next



How to play:

In every round you choose whether to pay to give points to the people you are connected to.

 If you click the **orange ring**, you pay 50 points for each player you are connected to and each of them gains 100 points.

 If you click the **blue ring**, you do not pay any points and do not change the points of the players you are connected to.

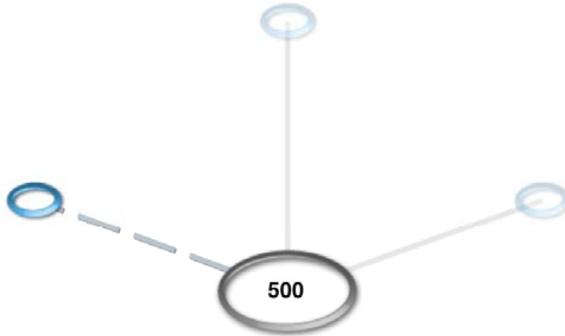
Each player you are connected to has the same choice. For each of them that chooses the **orange ring**, you gain 100 points.

Once everyone makes a decision the results are displayed. You will be shown the choices of each player you are connected to and how many points in total you gained or lost.

The order in which your neighbors are displayed may change at any time: you will not be able to keep track of your neighbors' actions beyond the most recent turn.

Remember, for every 10 points you have at the end of the game, we will add 1 cent to your bonus.

Next



How to play:

In some rounds, you also have the chance to change who you are connected to.

When this happens, you may be shown one randomly selected player in the game. We will show you:

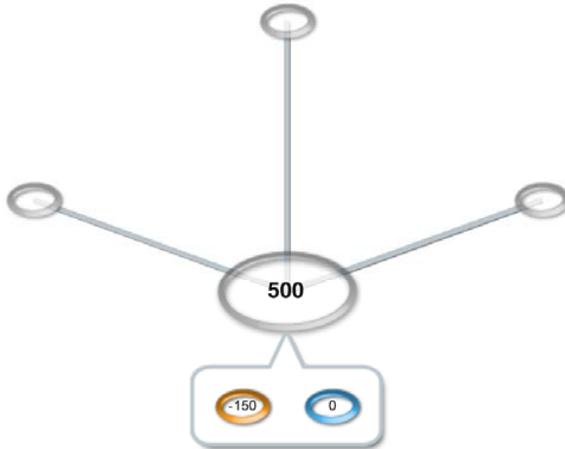
- Whether or not you are currently connected to this player
- Which choice they made in the previous round: **orange** (pay to give points to others) or **blue** (keep points for self).

If you are currently connected to this player you can choose to cut the connection. If you cut the connection with this player you will *no longer* pay 50 points for that player to gain 100 points when you choose **orange** and they will *no longer* pay 50 points for you to gain 100 points when they choose **orange**.

If you are not currently connected to this player you can choose to make a connection. If you make a connection with this player you will pay 50 points to give this player 100 points when you choose **orange** and they will pay 50 points to give you 100 points when they choose **orange**.

Next

Round: Tutorial



Practice round 1/3:

These rounds will not change your score. Your score will be reset before the game starts.

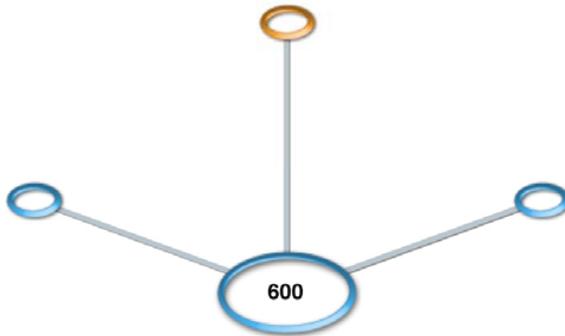
 If you click the **orange ring**, you pay 50 points for each player you are connected to and each of them gains 100 points.

 If you click the **blue ring**, you do not pay any points and do not change the points of the players you are connected to.

Each player you are connected to has the same choice. For each of them that chooses the **orange ring**, you gain 100 points.

Click a ring to continue.

Round: Tutorial



Practice round 1/3:

These rounds will not change your score. Your score will be reset before the game starts.

Last round 1 player(s) you are connected to paid 50 each to contribute a total of 100 points to you and everyone else they are connected to.

Last round you paid 0 to contribute 0 points to each player you are connected to.

Next

Round: Tutorial



Practice round 1/3:

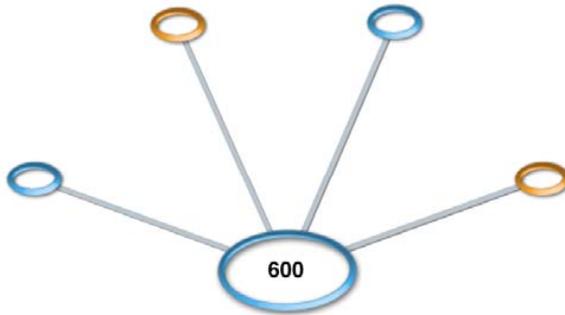
These rounds will not change your score. Your score will be reset before the game starts.

You are not currently connected to this player; you can choose to make a connection. If you make a connection with this player you will pay 50 points to give this player 100 points when you choose **orange** and they will pay 50 points to give you 100 points when they choose **orange**

Do you want to make a connection with this player?

Yes No

Round: Tutorial



Practice round 1/3:

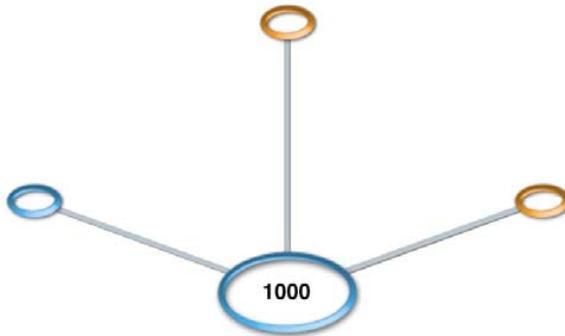
These rounds will not change your score. Your score will be reset before the game starts.

This round:

- you made 1 **new connections** with players.
- you broke 0 connections with players.
- 0 players made **new connections** with you.
- 0 players broke their connection with you.

Next

Round: Tutorial



You have now completed the tutorial.

We expect the task will take about half an hour to complete and we need your active participation until it is completed.

If you do not have about half an hour to allocate to this task, please close your browser window now.

Once the task has begun you will be dropped if you are idle for longer than 30 seconds. Your HIT will be rejected if you are dropped, and you will not be paid.

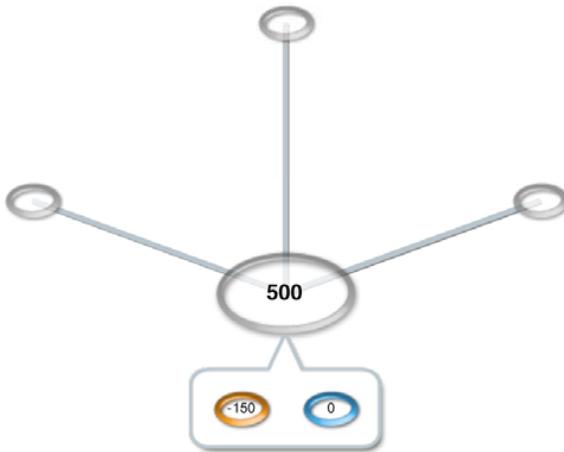
If you have understood the material in the tutorial, click 'Begin'.

Begin

Please note that even though we say that you will be dropped after 30 seconds, we are much more lenient (waiting 360 seconds), due to server and client-side delays.

Below are screenshots from the first two rounds of a sample game, where rewiring is allowed. For the randomized ties, no rewiring round is shown, and the players are told that “After every round, the connections between players are randomly shuffled.” For no rewiring, the rewiring rounds are skipped in the game, as well as the experiment.

Round: 1



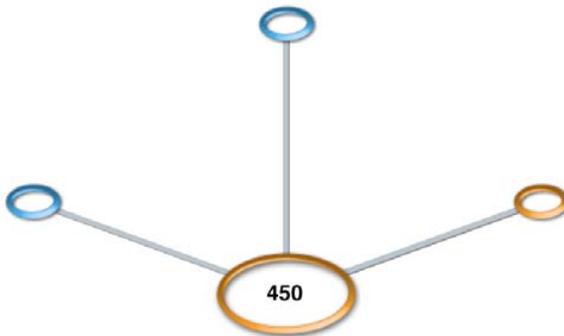
If you click the **orange ring**, you pay 50 points for each player you are connected to and each of them gains 100 points.



If you click the **blue ring**, you do not pay any points and do not change the points of the players you are connected to.

Each player you are connected to has the same choice. For each of them that chooses the **orange ring**, you gain 100 points.

Round: 1

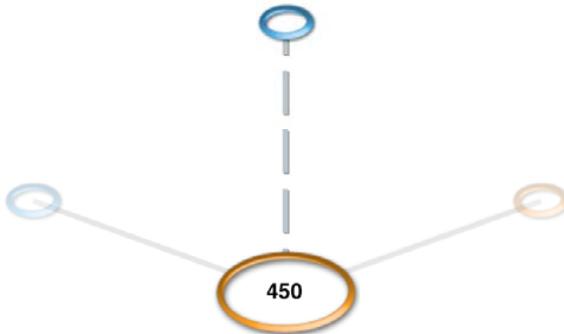


Last round 1 player(s) you are connected to paid 50 each to contribute a total of 100 points to you and everyone else they are connected to.

Last round you paid 150 to contribute 100 points to each player you are connected to.

Next

Round: 2

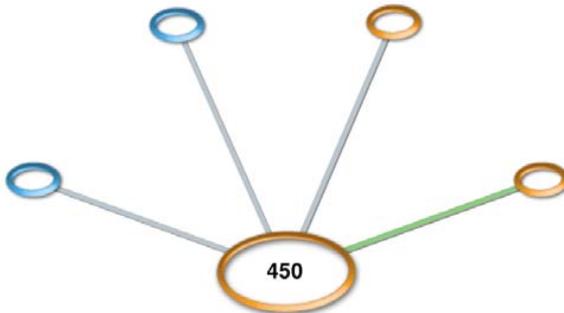


You are currently connected to this player; you can choose to cut the connection. If you cut the connection with this player you will *no longer* pay 50 points for that player to gain 100 points when you choose **orange** and they will *no longer* pay 50 points for you to gain 100 points when they choose **orange**

Do you want to break your connection with this player?

Yes No

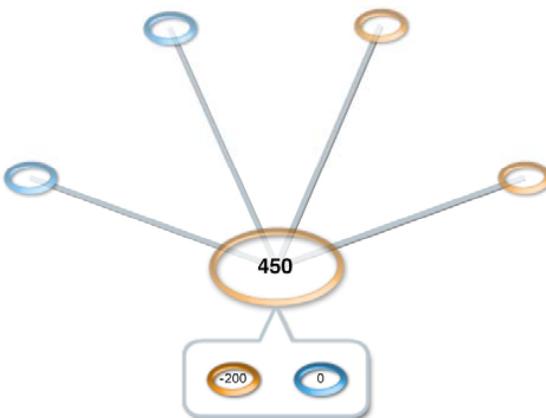
Round: 2



This round:

- you made 0 **new connections** with players.
- you broke 0 connections with players.
- 1 player made **new connections** with you.
- 0 players broke their connection with you.

Round: 2



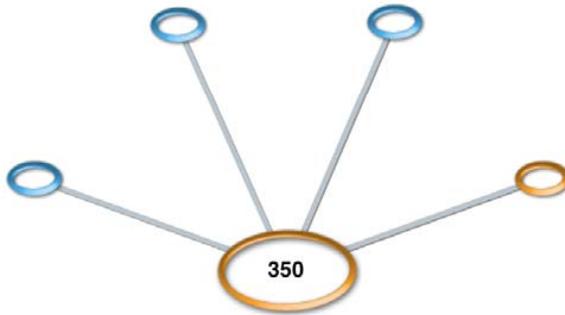
If you click the **orange ring**, you pay 50 points for each player you are connected to and each of them gains 100 points.



If you click the **blue ring**, you do not pay any points and do not change the points of the players you are connected to.

Each player you are connected to has the same choice. For each of them that chooses the **orange ring**, you gain 100 points.

Round: 2



Last round 1 player(s) you are connected to paid 50 each to contribute a total of 100 points to you and everyone else they are connected to.

Last round you paid 200 to contribute 100 points to each player you are connected to.

Next

2. Additional analysis of reciprocity via cooperative action

As reported in the main text, reciprocity via change in cooperation action occurs in all conditions, despite the varied success of cooperation across conditions. Thus, this type of reciprocity alone cannot explain the cooperation we observed in our fluid treatment. The relationship between cooperation in the current round and the behavior of one's interaction partners in the previous round is visualized in Figure S1.

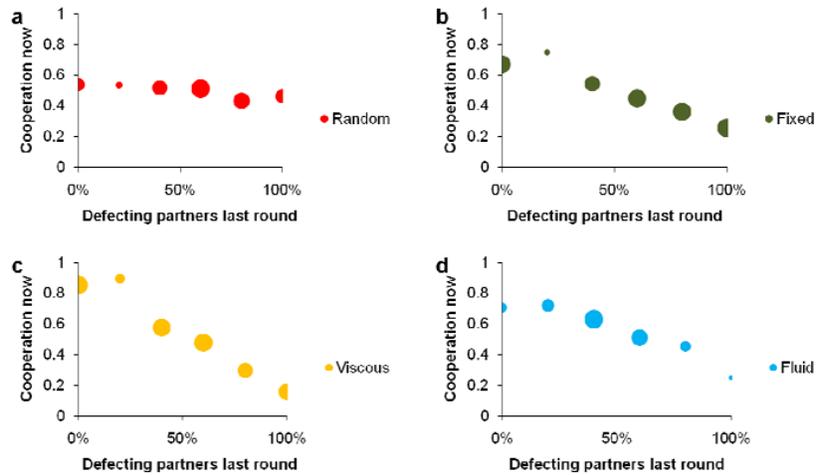


Figure S1. Relationship between a person's current behavior and the behavior of the person's previous interaction partners. Dot size reflects number of observations. A significant negative correlation exists in all treatments, including the random condition, although the slope of the correlation is significantly smaller in the random condition.

We note that in the random condition, unlike the other conditions, direct reciprocation is limited as partners are reshuffled every round, and therefore the slope of the correlation is much shallower (although still negative and statistically significant). What we observe in the random condition is instead a form of generalized (or upstream) reciprocity, in which subjects 'pay it forward.' As a result, the relationship between previous partners' defection and a person's cooperation is significantly weaker in the random condition ($p < 0.001$ for all 3 [condition] X [% defecting partners] interaction terms), though still negative.

We also note that the decrease in cooperation over time in the random, fixed, and viscous treatments, and the stability of cooperation in the fluid treatment shown in main text Figure 1a, are robust to the exclusion of subjects with no connections. As the initial network is determined randomly, there is some chance that a given player may have no neighbors; and in the dynamic network conditions, all of a given player's neighbors might break their connections with her. In such cases, the isolated player is still asked each round whether she would like to cooperate or defect, but since she has no neighbors, her response has no payoff consequences (and therefore is essentially meaningless). Excluding cooperation decisions made by subjects with 0 connections has virtually no impact on the results reported in Figure 1a: cooperation still decreases over time in the random (coeff=-0.11, $p < 0.001$), fixed (coeff=-0.19, $p < 0.001$) and viscous conditions (coeff=-0.22, $p = 0.011$), and is stable in the fluid condition (coeff=-0.03, $p = 0.444$).

3. Additional analysis of reciprocity via network rewiring

In the main text, we show that subjects in the fluid condition preferentially break ties with defectors and make new ties with cooperators. In Figure S2, we examine how the frequency of breaking and making links might change over time. Specifically, we ask what fraction of network updates result in the formation of a new link, the breaking of an existing link, or no change in the network. There is a significant increase over time in the fraction of network updates in which no action occurs (coeff=0.146, $p<0.001$). We see that this is driven by a dramatic decline in the fraction of updates which establish new links (coeff=-0.122, $p<0.001$); while conversely there is no significant change over time in breaking of existing links (coeff=-0.018, $p=0.433$). Both of these relationships are robust to also controlling for the player's current number of existing links (probability of making new links: coeff=-0.108, $p=0.021$; probability of breaking existing links, coeff=0.010, $p=0.713$). The network is randomly initialized, so, at first, many new ties are formed as pairs of cooperative players find and connect to each other. Over time, however, the network approaches a dynamic equilibrium, with the number of new ties created each round roughly equaling the number of existing ties broken.

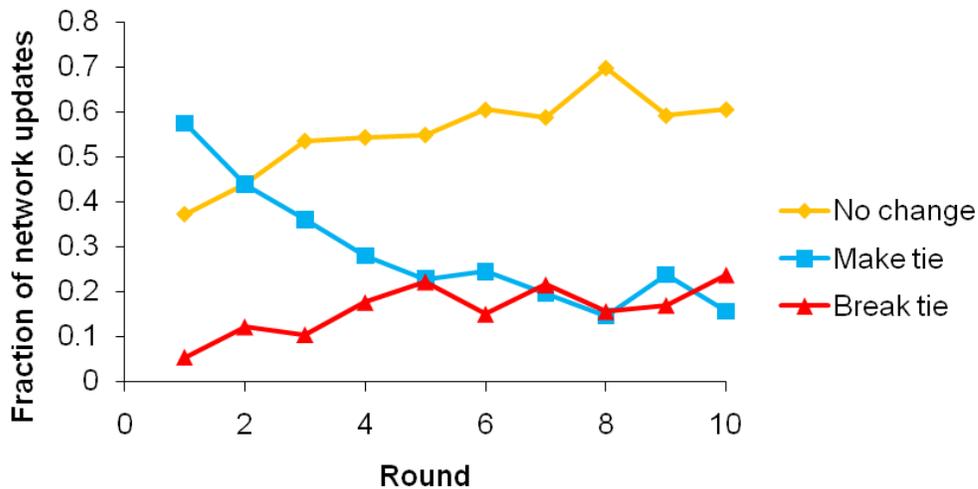


Figure S2. Fraction of network update events resulting in the formation of new ties, the breaking of old ties, or no change to the network.

In the main text, we also report that defectors are encouraged to switch to cooperation when others break links with them, but are unaffected by the formation of new links; and that cooperators are unaffected by either the making or breaking of links. Here, we report that this remains true when considering the fraction of possible links broken and formed rather than the absolute number, both for defectors (multivariate logistic regression clustered on subject and session, taking cooperation as the independent variable and including both fraction of possible links broken, coeff=2.23, $p=0.001$, and formed, coeff=-1.30, $p=0.55$) and for cooperators (multivariate logistic regression clustered on subject and session, taking cooperation as the independent variable and including both fraction of possible links broken, coeff=-3.08, $p=0.14$, and formed, coeff=-1.41, $p=0.17$).

4. Additional analysis of network properties and cooperation

Here we extend the analysis of the properties of the networks in our different conditions shown in Figure 1b. To capture differences which might emerge over time, we now compare the state of the networks in each condition in round 7 of play (excluding sessions which did not reach round 7). In Figure S3, we show the degree distribution for the 4 conditions, in a somewhat different presentational format compared to main text Figure 1. We see that not only is there greater degree heterogeneity in the fluid condition, but also that the average degree is higher in the fluid condition. To demonstrate the latter observation formally, we conduct a linear regression (clustered on subject and session) over all rounds, taking number of connections as the dependent variable, using the fluid condition as the baseline, and taking binary variables (i.e., ‘dummies’) for the other 3 conditions as independent variables. All 3 dummies are highly significant ($p < 0.001$ for all), indicating that subjects in the fluid condition have significantly higher degree on average than subjects in all 3 other conditions.

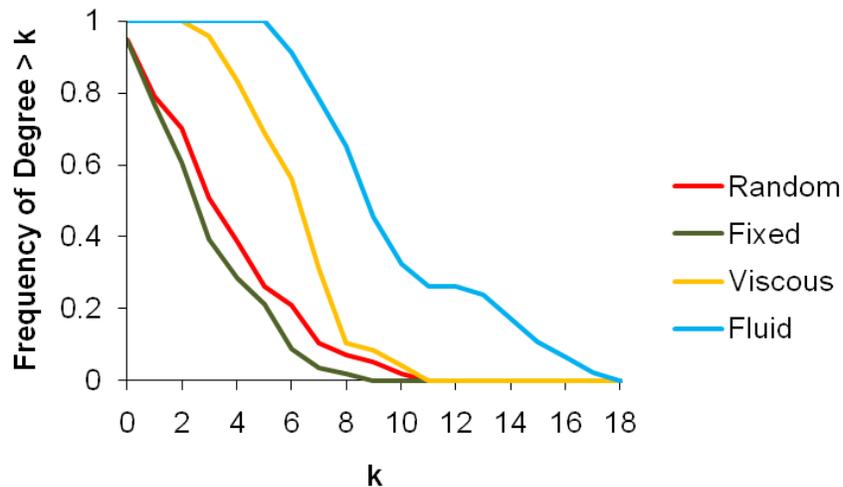


Figure S3. Degree distribution for each condition at round 7.

This raises a possible alternative explanation for the success of cooperation in the fluid condition in addition to the incentives created by the making and breaking of links. Namely, we see both higher average degree (as explicitly allowed by the experiment) and higher cooperation in the fluid condition. Perhaps, therefore, high degree subjects might somehow behave more cooperatively, simply by virtue of having higher degree. Since dynamic networks allow subjects to reach higher degree, might the potentially indiscriminate formation of new links somehow itself be a source of cooperation?

To provide evidence that this is not that case, we show that a positive association between cooperative behavior and degree emerges over time in the fluid condition, but not in the other conditions when regressing cooperation against *fraction* of possible connections (rather than total number of connections), which removes noise introduced by variation across sessions in maximum number of possible connections (fraction of possible connections X round interaction: random, coeff=-0.001, $p=0.99$; fixed, coeff=0.05, $p=0.83$; viscous, coeff=-0.28, $p=0.30$; fluid,

coeff=0.62, $p < 0.001$). Furthermore, we perform the same analysis, but we restrict analysis to subjects with 12 or fewer ties (which results in the omission of only 16% of the observations), because no subjects in the non-fluid conditions ever had more than 12 connections. In this subset, it still remains the case that a significant relationship between cooperation and fraction of possible connections emerges over time in the fluid condition (coeff=0.546, $p = 0.034$), with a coefficient even larger than when analyzing all subjects. This suggests that the effect in the fluid condition is not preferentially driven by subjects outside of the degree range observed in the other conditions.

Figure S4 helps visualize the fact that an association between cooperation and degree emerges only in the fluid condition, showing the degree distributions in round 7 for cooperators and defectors in each condition.

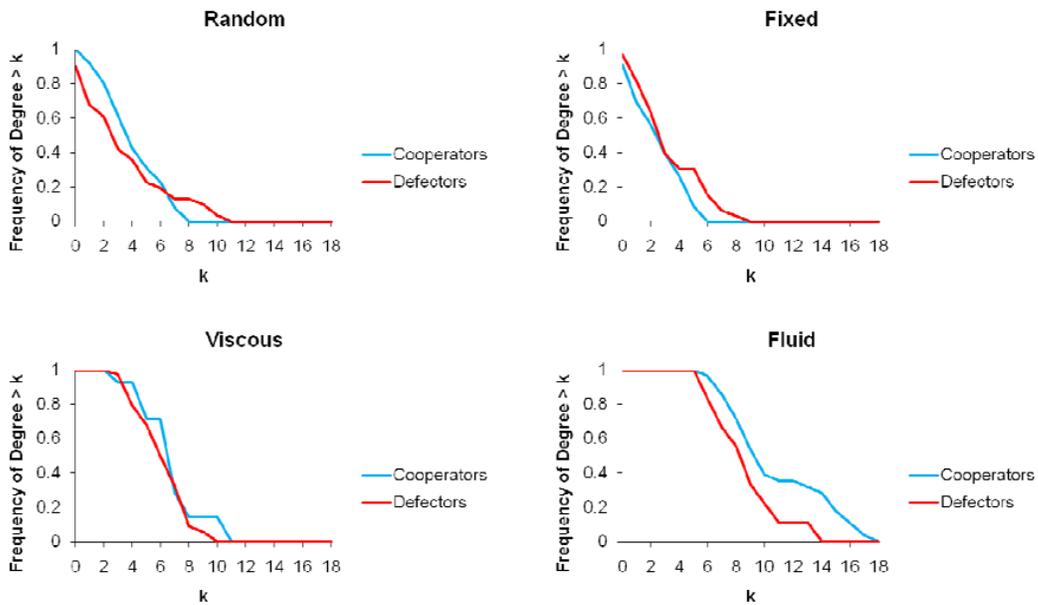


Figure S4. Degree distribution for cooperators and defectors in each condition at round 7.

We also observe more clustering in the fluid condition than the other conditions, among both cooperators and defectors, as shown in Figure S5. To explore this formally, we conduct a linear regression (clustered on subject and session) over all rounds, taking clustering coefficient as the dependent variable, using the fluid condition as the baseline, and including dummies for the other 3 conditions as independent variables. All 3 dummies are highly significant ($p \leq 0.001$ for all).

We define the clustering coefficient as the fraction of possible triangles that exist around a given node, where $T(v)$ is the number of triangles that go through node v :

$$c_v = \frac{2T(v)}{\deg(v)(\deg(v) - 1)}$$

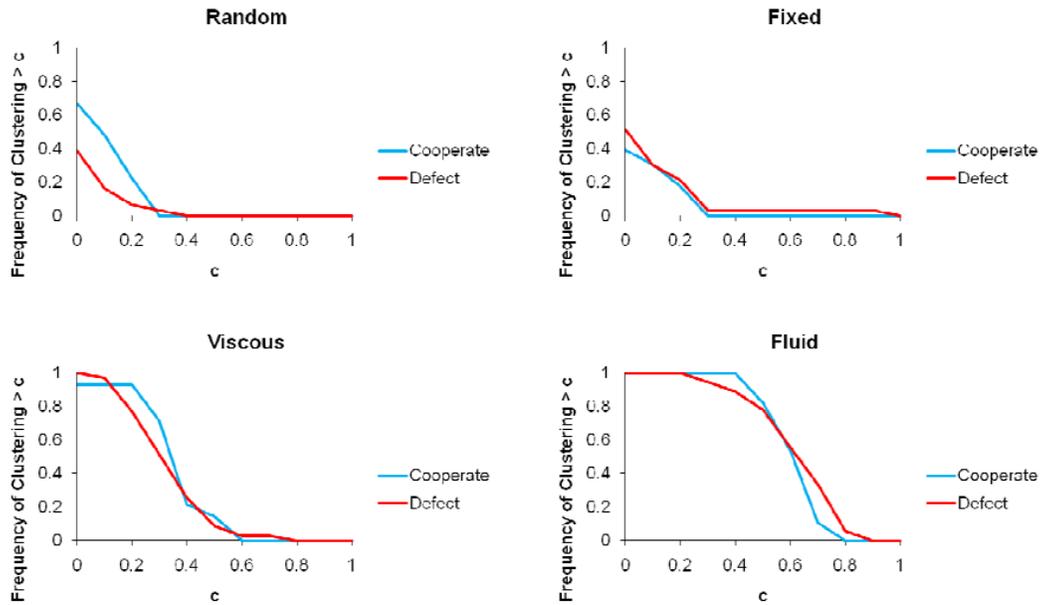


Figure S5. Clustering in round 7 of each condition, among cooperators and defectors.

We now turn to considering the consistency of cooperation decision across conditions. To do so, we first consider round-to-round consistency in each round of play. That is, in each round, we ask what fraction of subjects (i) consistently played C in both the previous and current round, (ii) consistently played D in both the previous and current round, or (iii) switched action between the previous and current round. The results are shown in Figure S6. We see a significant decrease over round in the number of consistent cooperators in the random (coeff=-0.122, $p=0.034$), static (coeff=-0.188, $p<0.001$), and viscous (coeff=-0.245, $p=0.026$) conditions, whereas there is no change over time in consistent cooperators in the fluid condition (coeff=-0.046, $p=0.272$). Conversely, there is a significant increase over round in the number of consistent defectors in the random (coeff=0.145, $p<0.001$), static (coeff=0.231, $p<0.001$), and viscous (coeff=0.297, $p=0.002$) conditions, whereas once again there is no change over time in consistent defectors in the fluid condition (coeff=0.017, $p=0.701$). Interestingly, the number of players changing their action does not change significantly over time in the random condition (coeff=-0.035, $p=0.42$), decreases over round in the static (coeff=-0.089, $p=0.025$) and viscous (coeff=-0.082, $p=0.003$) conditions, and increases over round in the fluid condition (coeff=0.043, $p=0.047$). Further exploration of these differences in switching of action is an interesting direction for future research.

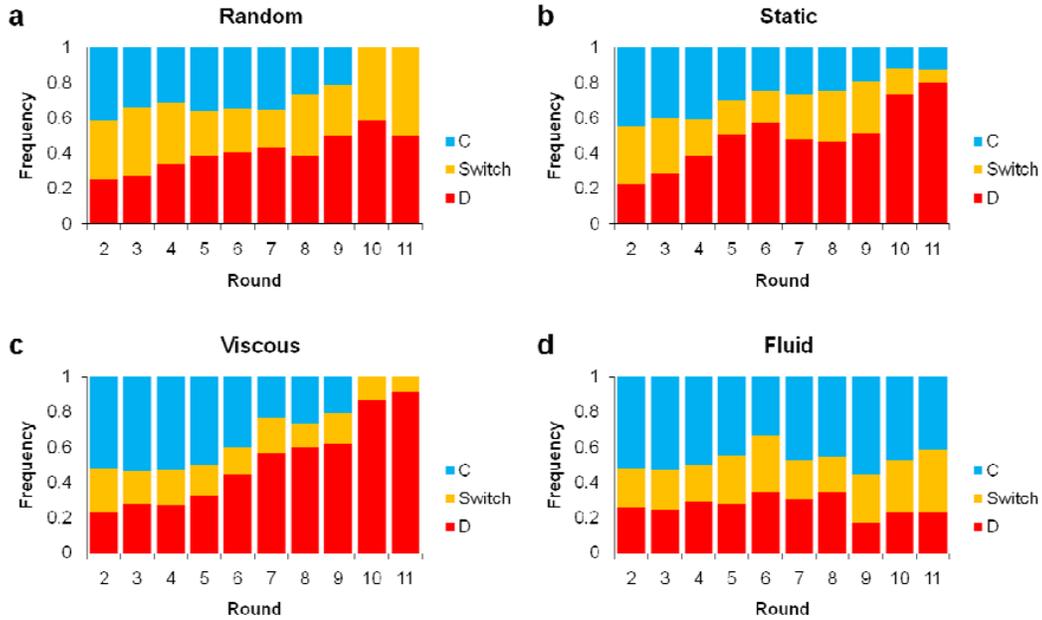


Figure S6. Consistency of play across rounds. Shown are the fraction of players consistently choosing C in both the previous and current round (blue), switching actions between the previous and current round (yellow) or consistently choosing D in the previous and current round (red).

Finally, we examine consistency over a longer window. To allow for some learning, we ask what fraction of subjects chose the same action in the latter half of the rounds (rounds 7-11), or the subset of those rounds for which their particular session lasted (recall that because of the stochastic end game rule, game length varied across session). We find that the fluid dynamic condition leads to significantly more consistency compared to the random and static conditions, but not more consistency than the viscous condition (R vs S, $p=0.821$; R vs V, $p=0.098$; R vs F, $p=0.004$; S vs V, $p=0.142$; S vs F, $p=0.004$; S vs F, $p=0.006$; V vs F, $p=0.279$).

5. Supporting References

1. Horton JJ, Rand DG, & Zeckhauser RJ (2011) The Online Laboratory: Conducting Experiments in a Real Labor Market. *Experimental Economics*:doi:10.1007/s10683-10011-19273-10689.
2. Rand DG (2011) The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of theoretical biology*:doi: 10.1016/j.jtbi.2011.1003.1004.
3. Suri S & Watts DJ (2011) Cooperation and Contagion in Web-Based, Networked Public Goods Experiments. *PLoS ONE* 6(3):e16836.
4. Grujić J, Fosco C, Araujo L, Cuesta JA, & Sánchez A (2010) Social Experiments in the Mesoscale: Humans Playing a Spatial Prisoner's Dilemma. *PLoS ONE* 5(11):e13749.
5. Traulsen A, Semmann D, Sommerfeld RD, Krambeck H-J, & Milinski M (2010) Human strategy updating in evolutionary games. *Proceedings of the National Academy of Sciences* 107(7):2962-2966.
6. Dreber A, Rand DG, Fudenberg D, & Nowak MA (2008) Winners don't punish. *Nature* 452(7185):348-351.
7. Fudenberg D, Rand DG, & Dreber A (In press) Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World. *American Economic Review*.