Cooperation and Contagion in Web-Based, Networked Public Goods Experiments

SIDDHARTH SURI and DUNCAN J. WATTS
Yahoo! Research

A longstanding idea in the literature on human cooperation is that cooperation should be reinforced when conditional cooperators are more likely to interact. In the context of social networks, this idea implies that cooperation should fare better in highly clustered networks such as cliques than in networks with low clustering such as random networks. To test this hypothesis, we conducted a series of experiments on Amazon Mechanical Turk, in which 24 individuals played a local public goods game arranged on one of five network topologies that varied between disconnected cliques and a random regular graph. In contrast with previous work, we found that network topology had no significant effect on average contributions. This result implies either that individuals are not conditional cooperators, or else that cooperation does not benefit from positive reinforcement between connected neighbors. We then tested both of these possibilities in two subsequent series of experiments in which artificial "seed" players were introduced, making either full or zero contributions. First, we found that although players did generally behave like conditional cooperators, they were as likely to decrease their contributions in response to low contributing neighbors as they were to increase their contributions in response to high contributing neighbors. Second, we found that positive effects of cooperation did not spread beyond direct neighbors in the network. In total we report on 113 human subjects experiments, highlighting the speed, flexibility, and cost-effectiveness of web-based experiments over those conducted in physical labs.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Economics, Behavioral Economics, Social Networks

Additional Key Words and Phrases: Cooperation, Public Goods, Social Networks, Contagion

1. INTRODUCTION

The idea that cooperative behavior might arise as a consequence of the population structure was initially proposed in the context of evolutionary biology [Hamilton 1964]. This notion has particular relevance for social dilemmas among human actors, where the total population is large, but the effects of any one individual's actions fall disproportionately on a relatively small set of neighbors determined either by spatial or social proximity. For example, smog or acid rain causing pollutants disproportionately impact geographically proximate populations; thus one can think of the game as playing out on some approximation of a spatial lattice. Correspondingly, the benefit derived from social networking sites (e.g. Facebook) is highly dependent on the activities and contributions of one's immediate social acquaintances, whose identities in turn depend some complicated mixture of social and spatial distance [Watts et al. 2002]. Because in either case an individual's neighbors are themselves connected to others, who are in turn connected to others still, and so on, the dynamics of social dilemmas can be thought of as taking place on extended networks [Newman 2003; Strogatz 2001]. In these settings, out-

Authors' addresses: `suri@yahoo-inc.com`, `djw@yahoo-inc.com`

comes of interest, such as aggregate levels of cooperation, plausibly depend on the structure of the network as well as on the strategies of the individuals in the population [Nowak et al. 2010].

There are two main reasons to suspect that cooperation should depend on network structure. The first reason is that many theoretical models of social dilemmas assume that cooperation is conditional, in the sense that an individual will only cooperate on the condition that its partners are also cooperating. Arguably the clearest example of the principle of conditional cooperation is the celebrated Tit-For-Tat strategy, which has consistently been shown to outperform more exploitative strategies in a range of simulation studies, in large part because it performs well when interacting with other cooperative strategies [Axelrod 1984]. In addition, related strategies have also been proposed that generalize the idea of conditional cooperation to multi-player settings [Watts 1999; Glance and Huberman 1993], usually by specifying some form of threshold requirement—i.e. "I will cooperate if at least X of my neighbors cooperated last round, else I will defect." Regardless of the specifics of the rule, the implication of these results for networks is that networks characterized by high levels of local clustering [Watts and Strogatz 1998], meaning that an individual's neighbors are also likely to be neighbors of each other, ought to sustain higher aggregate levels of cooperation than populations in which individuals are randomly mixed [Axelrod and Hamilton 1981]. Put another way, local reinforcement would imply that when an individual's neighbors also interact with each other, they are in a better position to reinforce one another's pro-social behavior, and so may be expected to resist "invasion" by defecting strategies better than when each neighbor interacts with a different set of others.

The second reason to suspect that network structure should impact cooperation is that cooperation in networks might be "contagious". Specifically, if A is a conditional cooperator surrounded mostly by cooperating neighbors, A will cooperate more; but then A's increased cooperation may cause its remaining neighbors to cooperate more as well. These neighbors may in turn cause their neighbors to cooperate more as well, and so on, leading to a cascade of cooperation that sustains itself over multiple steps. In fact, recently it has been claimed that cooperation is characterized by a "three degrees of influence" rule [Fowler and Christakis 2010], meaning that an individual who increases his or her level of cooperation can positively impact the contribution of an individual who is three steps removed from them in the network. Because the number of individuals who can be reached within three degrees of a cooperating individual will in general depend on the non-local structure of the network [Watts and Strogatz 1998], the presence of social contagion would imply that network features other than local clustering should also impact aggregate cooperation levels.

2. EXPERIMENTAL DESIGN

In contrast with standard public goods games, in which participants' contributions are shared among members of the same group, here participants are arranged in a network and their payoff is only affected by the actions of their neighbors. To reflect this change, players' payoffs are defined by the function $\pi_i = e_i - c_i + \frac{a}{k+1}\sum_{j \in \Gamma(i)} c_j$, where the last term is summed over $\Gamma(i)$, the network neighborhood of i (which we define to include i itself), and k is the vertex degree (all nodes in all networks have the same degree). Therefore, i's contributions are, in effect, divided equally among the edges of the graph that are incident on i, where payoffs are correspondingly summed over i's edges. Critically, when $1 < a < k+1$ meaning that the marginal per capita return $M = \frac{a}{k+1}$ lies in the range $0 < M < 1$ players face a social dilemma in the sense that social welfare is maximized when all individuals contribute the maximum amount, but players have a selfish incentive to free ride on the contributions of others. Aside from adapting the payoff function to the network setting, our experimental design was kept as similar as possible to previous work in order to make comparisons possible [Fehr and Gachter 2000]. For example, we ran each experiment for 10 rounds and set $M = 0.4$.

We chose networks that spanned a wide range of possible structures between a collection of four disconnected cliques at one extreme, and a random regular graph at the other. All networks comprised $n = 24$ players, each with constant vertex degree $k = 5$; however, they varied with respect to three frequently studied structural parameters, (a) the clustering coefficient $C = \frac{2}{n}\sum_{i=i}^{n} \frac{K_i}{k(k-1)}$ where K_i is the number of completed triangles in node i's neighborhood; (b) the average path length $L = \langle d_{ij} \rangle$ where the average distance between all pairs of nodes is taken over each connected component; and (c) the diameter $D = \max d_{ij}$, which is the distance between the farthest two nodes. The clustering coefficient of node i is computed by dividing the number of triangles incident on i by the number of triangles possible given i's degree. The clustering coefficient of a network, which is the average clustering coefficient over all nodes, is therefore a local measure of structure that captures the extent to which the neighbors of i are also neighbors of each other. The average path length and diameter, by contrast, are global network measures that quantify the extent to which effects can propagate along chains of network ties.

If the "reinforcement" hypothesis, outlined above, is correct the actions of an individual's neighbors ought to be dependent on the actions of their neighbors. Hence the experience of the focal individual will depend on the density of interaction between his or her immediate neighbors, which is measured by the clustering coefficient. Correspondingly, if the "contagion" hypothesis is correct, the focal individual's experience will depend in addition on the actions of individuals by two or more steps away and hence on the distribution of path lengths. Thus our choice of topologies was specifically designed to highlight the importance both of local reinforcement and contagion.

3. RESULTS

In the first set of network experiments all positions in the network were filled by human players recruited from Amazon Mechanical Turk. Because individual contributions tended to vary considerably from one experiment to the next, and different players were likely to play at different times of day, we conducted multiple realizations of the experiment for each topology. In total, we conducted 23 experiments in random order. Surprisingly, we did not find significant differences in the average contributions in each round for the various topologies. Moreover, we did not find large differences in the distribution of individual contributions over each round, or in the average contributions of groups defined by the network topologies over each round. Thus we conclude that topology does not exert a noticeable impact on contributions at any level: individual, group, or aggregate.

The absence of topological dependency of contributions suggests that one or both of the hypotheses outlined above (reinforcement and contagion) must be wrong. We therefore conducted two further series of experiments designed to test the reinforcement and contagion hypotheses respectively. In the first series of 30 experiments we followed the same design as above, but with the key difference that in each experiment four nodes were selected and their contributions were all artificially fixed either at 10 (the "cooperative" condition) or 0 (the "defection" condition) for all rounds. Thus, we were able test the reinforcement hypothesis by directly measuring the positive/negative influence of unconditional cooperators/defectors on their immediate neighbors. The seed players were arranged in order to cover the network, meaning that each human player was adjacent to precisely one seed player (in the random regular case, a perfect cover arrangement did not exist for the selected network; thus a close approximation was used instead). An advantage of this arrangement, is that all human players were subjected to the same experimentally manipulated influence.

The presence of cooperating seeds stimulated consistently higher aggregate contributions from the remaining 20 players, while the presence of defecting seeds had the opposite effect. Possessing a high (or low) contributing neighbor therefore did increase (or decrease) the average contribution levels; thus our subjects were indeed behaving as conditional cooperators. Nevertheless the effect of the seed players was not consistently bigger in the graphs with the highest clustering. For example the effect of the seed nodes in a network of disjoint cliques, which had the maximum number of triangles incident on each node, was very similar to the effect of the seeds nodes in the random regular network, which had fewer than 1/10th as many triangles. This result implies that two nodes that form a triangle with a cooperating (or defecting) seed do not have an appreciably larger (or smaller) average contribution level then two disconnected nodes with a cooperating (or defecting) seed neighbor in common. Mutual reinforcement of the contributions among the neighbors of a seed node is largely absent, whether or not there is an edge between the neighbors.

Next, in a series of 20 experiments over 2 weeks, we tested the contagion hypothesis by keeping the number of unconditionally cooperating seeds constant at four per network (we did not introduce unconditional defectors in these experiments), but concentrating them together into two adjacent pairs. This arrangement of seeds exposed some human players to two unconditional cooperators as immediate neighbors, while others were not exposed to any seeds directly, but were connected indirectly to the seeds via a human intermediary. If positive contagion were present in the network, we would expect to see nodes at distance two from the seeds increase their contributions relative to the all-human (i.e. no seeds) condition. Quite to the contrary, in fact, the two-step neighbors of the cooperating seeds contributed slightly less than the nodes in the corresponding network positions contributed in the all-human experiments.

4. DISCUSSION

Returning to our original motivation, theoretical arguments in favor of an association between network structure and cooperation invoke two related ideas: first, that individuals are conditional cooperators, increasing their contributions in response to the increased contributions of their neighbors; and second, that positive effects of conditional cooperation should propagate through the network via a process of contagion. In this paper, we have tested the effects of network topology on contribution levels in a standard public goods game, finding no significant effects. In addition, we conducted two separate rounds of experiments—one to test for the presence of conditional cooperation, and the other to test for the possibility of positive contagion. Although we do find strong evidence of conditional cooperation, we do not find evidence of positive contagion in the standard sense of multi-step propagation along a sequence of ties in a static network.

Our explanation for these results is that the theoretical arguments cited above emphasize the positive aspect of conditional cooperation, yet conditional cooperation implies not only that players increase their contributions in response to cooperative neighbors, but also that they decrease their contributions in response to defecting neighbors. Although it is the case that highly clustered networks offer more opportunities for positive effects to reinforce each other than random networks, they also offer more opportunities for negative effects to reinforce each other as well. By contrast, in random graphs where there is very little clustering, neither cooperation nor defection get reinforced and seeds act as influence blockers preventing either positive or negative influence from propagating among neighbors.

Finally, we close by noting that in contrast with our results here, network structure has been found to have very striking effects in coordination games [Kearns et al. 2006; Kearns et al. 2009]. What could account for the for the difference? In the case of coordination games, if node A chooses an action that results in a lack of coordination with neighbor B, then B has a clear incentive to change its action. In turn, if this results in a lack of coordination with C which is a neighbor of B and not A, this can result in contagion. In a cooperation setting, by contrast, B need not change its action in response to A because the incentives do not enforce such a tight coupling of neighbors actions. Although subtle, we suspect that this difference between coordination and cooperation games is responsible for the respective presence and absence of network effects. This observation in turn leads to an interesting open question—under what theoretical conditions should one expect to see contagion over networks with fixed neighbors? In demonstrating that not all dynamic games on networks exhibit contagion we hope that our findings will provoke new theoretical hypotheses along these lines, as well as new experiments to test them.

REFERENCES

Axelrod, R. and Hamilton, W. D. 1981. The evolution of cooperation. *Science 211,* 4489, 1390–1396.

Axelrod, R. M. 1984. *The evolution of cooperation.* Basic Books, New York.

Fehr, E. and Gachter, S. 2000. Cooperation and punishment in public goods experiments. *The American Economic Review 90,* 4, 980–994.

Fowler, J. H. and Christakis, N. A. 2010. Cooperative behavior cascades in human social networks. *Proceedings of the National Academy of Sciences 107,* 9, 5334–5338.

Glance, N. S. and Huberman, B. A. 1993. The outbreak of cooperation. *Journal of Mathematical Sociology 17,* 4, 281–302.

Hamilton, W. D. 1964. The genetical evolution of social behaviour. *Journal of Theoretical Biology 7,* 1, 17–52.

Kearns, M., Judd, S., Tan, J., and Wortman, J. 2009. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences 106,* 5, 1347.

Kearns, M., Suri, S., and Montfort, N. 2006. An experimental study of the coloring problem on human subject networks. *Science 313,* 5788, 824–827.

Newman, M. E. J. 2003. The structure and function of complex networks. *Siam Review 45,* 2, 167–256.

Nowak, M. A., Tarnita, C. E., and Antal, T. 2010. Evolutionary dynamics in structured populations. *Philosophical Transactions of the Royal Society B: Biological Sciences 365,* 1537, 19.

Strogatz, S. H. 2001. Exploring complex networks. *Nature 410,* 6825, 268–276.

Watts, D. J. 1999. *Small Worlds : The Dynamics of Networks Between Order and Randomness.* Princeton University Press, Princeton.

Watts, D. J., Dodds, P. S., and Newman, M. E. J. 2002. Identity and search in social networks. *Science 296*, 1302–1305.

Watts, D. J. and Strogatz, S. H. 1998. Collective dynamics of 'small-world' networks. *Nature 393,* 6684, 440–442.