Cooperation, clustering, and assortative mixing in dynamic networks

David Melamed[a,1], Ashley Harrell[b], and Brent Simpson[c,1]

[a]Department of Sociology, The Ohio State University, Columbus, OH 43210; [b]Department of Organizational Studies, University of Michigan, Ann Arbor, MI 48109; and [c]Department of Sociology, University of South Carolina, Columbia, SC 29208

Edited by Alan Hastings, University of California, Davis, CA, and approved December 19, 2017 (received for review August 30, 2017)

Humans' propensity to cooperate is driven by our embeddedness in social networks. A key mechanism through which networks promote cooperation is clustering. Within clusters, conditional cooperators are insulated from exploitation by noncooperators, allowing them to reap the benefits of cooperation. Dynamic networks, where ties can be shed and new ties formed, allow for the endogenous emergence of clusters of cooperators. Although past work suggests that either reputation processes or network dynamics can increase clustering and cooperation, existing work on network dynamics conflates reputations and dynamics. Here we report results from a large-scale experiment (total $n = 2,675$) that embedded participants in clustered or random networks that were static or dynamic, with varying levels of reputational information. Results show that initial network clustering predicts cooperation in static networks, but not in dynamic ones. Further, our experiment shows that while reputations are important for partner choice, cooperation levels are driven purely by dynamics. Supplemental conditions confirmed this lack of a reputation effect. Importantly, we find that when participants make individual choices to cooperate or defect with each partner, as opposed to a single decision that applies to all partners (as is standard in the literature on cooperation in networks), cooperation rates in static networks are as high as cooperation rates in dynamic networks. This finding highlights the importance of structured relations for sustained cooperation, and shows how giving experimental participants more realistic choices has important consequences for whether dynamic networks promote higher levels of cooperation than static networks.

altruism | cooperation | dynamic networks | network science | reputation

Few issues have puzzled both social and biological scientists as much as the high level of cooperation observed among humans. Given that cooperation entails paying a cost for another to receive a benefit, such behavior poses an evolutionary paradox. Consistent with this logic, empirical research finds low levels of cooperation in "one-shot" interactions (1). But instead of one-shot interactions with strangers, much of human social life is embedded in networks (2). Research shows that repeated interactions in networks are characterized by much higher levels of cooperation than one-shot interactions (3–15). Further, clustered networks, characterized by dense regions, promote higher rates of cooperation than random networks (16), and dynamic networks, where relations can be dropped or new ties formed, promote even higher levels of cooperation (17–20). Clustering allows cooperation to thrive by insulating cooperators from exploitation by noncooperators. By extension, network dynamics allow conditional cooperators to shed their ties to defectors and seek out other cooperators, thus endogenously producing highly cooperative clusters (17).

While several studies have shown the powerful effects of dynamic networks on cooperation, it is unclear how much, if any, of this effect can be attributed to dynamics, per se. Another mechanism that promotes cooperation is reputations (11, 21). Several key studies showing that dynamic networks promote cooperation conflate dynamics with reputation effects (18, 19, 22). That is, when participants in these studies select new ties, they do so based on information about the past behaviors of potential partners. Participants use available reputational information to select new ties (21, 23–25), and are even willing to incur costs to do so (26, 27). As such, observed rates of cooperation may be driven by network dynamics, reputational processes, or both. Notably, dynamic networks promote cooperation even when new ties form at random (17, 28). What we cannot know from prior work is the benefit beyond dynamic network processes of knowing the reputations of others. While Gallo and Yan (21) showed that knowing the reputations of all others in the network promoted cooperation over not knowing, the networks in their study were created anew in each round with participants forming costless ties to as many others as they wanted. Despite this, other work has shown that reputations have no effect on cooperation (29).

Our study extends our understanding of dynamic networks and cooperation in three unique ways. First, we implement four distinct types of networks. As described more fully below, we study static networks where ties are unalterable, dynamic networks void of reputational information, dynamic networks where participants know the reputations of all others, and a more realistic dynamic network where participants only know the reputations of their neighbors' neighbors, i.e., those two links removed in the network. Unlike the full information condition, this last condition does not assume omniscience, but instead mimics referral processes, where we learn about trustworthy others through our existing ties. This enables us to distinguish reputational effects from network dynamics and to further discern whether and how these dynamics vary with "global" versus "local" reputational information. The global versus local reputational information conditions also enable us to examine the endogenous emergence of clustering under more realistic conditions. Second, we vary

Significance

Understanding the patterns and processes of human cooperation is of central scientific importance. Networks can promote cooperation when their existing or emergent topology allows conditional cooperators in the network to isolate themselves from exploitation by noncooperators. We do not know from prior work whether the emergent structures that promote cooperation are driven by reputation or can emerge purely via dynamics, i.e., the severing of ties to noncooperators and the formation of new ties irrespective of reputational information. Here we demonstrate, experimentally, that dynamic networks yield very high rates of cooperation even without reputational knowledge. Further, we identify realistic conditions under which static networks (where ties cannot be altered) yield cooperation rates as high as those in dynamic networks.

Author contributions: D.M. and B.S. designed research; A.H. performed research; D.M. analyzed data; and D.M., A.H., and B.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: melamed.9@osu.edu or bts@sc.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1715357115/-/DCSupplemental.

SOCIAL SCIENCES

EVOLUTION

whether initial networks are random or clustered. This complements our network manipulation by allowing us to assess whether initially clustered networks lead to more endogenous clustering in dynamic networks, and further promote cooperation.

Third, nearly all studies on the evolution of cooperation in networks force participants to simultaneously cooperate or defect with all alters, or participants to whom they are connected (18, 19, 22). But in real-world interactions, we are generally able to cooperate in one relationship while acting more selfishly in another. Similarly, we are typically able to reciprocate defection in one of our relations without simultaneously defecting in all our other relations. The standard design in experimental studies of networks and cooperation does not allow this. Importantly, we argue that whether we study cooperation in networks via the standard design (where a participant's decision to cooperate or defect necessarily applies to all alters) versus more realistic "targeted" choices (where a participant makes a decision about whether to cooperate or defect in each of his/her relationships) has important consequences for our understanding of when and why networks promote cooperation. The standard design will artificially inflate cooperation in dynamic networks if participants who are tempted to defect in some relations instead cooperate to reduce the risk of losing their ties. Similarly, the standard design will artificially suppress cooperation in static networks if participants have no other recourse to punish alters that have previously defected. In other words, the greater cooperation rates in dynamic versus static networks documented in prior work may be partly artifactual. Here we aim to decouple pure structural effects from how participants make their decisions about whether to cooperate.

To disentangle the effects of clustering, dynamics, and reputational processes, we conducted a large-scale behavioral experiment using workers from Amazon's Mechanical Turk (*Materials and Methods*). Workers were embedded within a network and played an iterated prisoner's dilemma (PD) game with each of those to whom they were connected. Participants were able to take part in multiple conditions, but could not participate in the same condition multiple times. Across all conditions (described below) there were 1,979 participants, 810 of whom were unique. Participants were randomly assigned to conditions and then to positions within networks (average network size = 24.7, SD = 2.7). Each network position was identified in interactions via a unique, randomly generated letter. In all conditions, participants began with an endowment of 1,000 monetary units (MUs). Following related work (18), cooperation entailed paying 50 MUs, which resulted in the alter gaining 100 MUs, while defection entailed paying nothing and generating no benefits. MUs were converted into dollars at the end of the study. Participants completed 16 rounds, but (to avoid end-game effects) were not told how many rounds to expect.

The networks were either static or dynamic. In dynamic networks, participants could drop one alter and initiate a tie with a new alter in each round; alters could then approve or decline new tie requests. Ties were not replaced for participants who were dropped (18, 22, 28, 30). The identifying labels of all prospective alters were displayed when participants had the opportunity to add a new tie (see *SI Appendix* for screenshots). Prospective new alters were those not tied to the participant, including those the participant had dropped in prior rounds. As long as a participant maintained at least one tie, he/she was eligible to be selected by others for a new tie. Any participant who lost all of his/her ties became a permanent isolate, and was excluded from both the network and the tie selection process for the remainder of the study.

We operationalized network dynamics in three ways. In the "no-reputation" condition, participants were given no information about potential alters' reputations when adding new ties, which is akin to replacement at random (17, 28). This condition allows us to assess the effects of pure network dynamics. In the "global-reputation" condition, participants knew the past cooperative behaviors of all alters, represented as a percentage of times they cooperated previously; this type of omniscience is standard in studies of dynamic networks (18, 19, 22). Finally, we introduced a novel "local-reputation" condition where participants only knew the past cooperative behaviors of alters that were two steps removed from them in the network, i.e., those who were connected to one or more of their own alters. As in the global-reputation condition, participants in the local-reputation condition were free to choose any alter when initiating a new tie, but only knew the reputations of their alters' alters. This local-reputation condition mimics the real-world process of referrals where we learn about trustworthy others through our existing ties and is consistent with how triadic closure is theorized to operate in social networks (31).

Initial networks were either random (Erdös–Rényi model) or clustered graphs. In both cases, the initial network density was 0.167, corresponding to an average of about four ties per node in round 1. The clustering coefficient (coeff) (32) for the random networks was the same as the density. The clustering coefficient for the clustered networks was 0.42 at round 1. Crossing random or clustered with static networks or one of the three types of dynamic networks yields eight network conditions. We ran 10 networks per condition, for a total of 80 networks.

We also counterbalanced on whether participants' decisions to cooperate or defect applied to all alters for a given round (18, 19, 21, 22) or were targeted, i.e., specific to each alter (17, 28). As noted above, we investigate whether this key methodological difference affects rates of cooperation and other outcome measures. As such, in each of our 8 main conditions, we studied five networks with "diffuse" choices where within a given round, participants made a single decision to cooperate or defect that applied to all alters (the standard procedure in the literature) and five networks with targeted choices for each alter, yielding 16 different conditions.

We expected clustering to matter more for cooperation in static—versus dynamic—networks, since clustering emerges endogenously in dynamic networks regardless of the initial structure (17, 18). Within the dynamic networks, we expected higher rates of cooperation in both reputation conditions, compared with networks with no reputational information (21). Even the local-reputation condition, where participants know the reputations of their alters' alters, has sufficient information for participants to select more cooperative partners. For the local-reputation condition, we expected that participants would be more likely to select new ties from the set of alters for whom they knew their reputations. We also expected clustering to endogenously increase in the local-reputation condition. Participants use reputations to select partners (21, 24). And only having referrals from alters, as often happens in real-world networks, increases the likelihood of triadic closure (31). Similarly, we anticipated a higher hazard of becoming isolated from the network for defectors in the reputation conditions. This is because defectors will not only be dropped by their interaction partners (19), they will also be avoided by prospective new partners, given their bad reputations.

Finally, we expected differences in cooperation in dynamic versus static networks to increase under the standard experimental design where participants make a single decision about whether to cooperate or defect that applies to all alters. As noted earlier, this could be because diffuse decisions increase cooperation in dynamic networks, decrease cooperation in static networks, or both.

Results

Fig. 1 shows the proportion of cooperation by round. What is immediately apparent is the very high level of cooperation in dynamic—versus static—networks. In the dynamic conditions, we find cooperation levels similar to those reported in related work on dynamic networks (19). Across the dynamic network conditions, round 1 cooperation rates were 86%. By round 3 there was 99% cooperation, and cooperation remained high for the remainder of the study. A similar rate of cooperation is observed early in the static networks, but it is not sustained. In round 1, there was 85% cooperation, but by round 16 it was down to 72%. Moreover, in the dynamic networks, it appears that reputations do not matter for cooperation: cooperation rates in the no-reputations conditions were similar to cooperation rates in those conditions

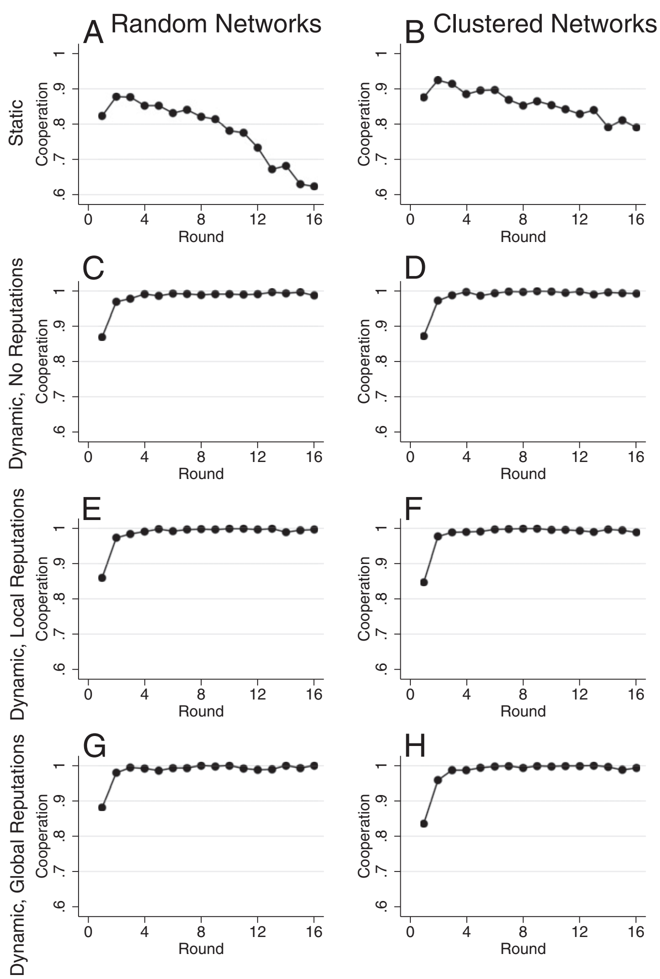

Fig. 1. Proportion of cooperation by round for (*A*) random static networks, (*B*) clustered static networks, (*C*) random dynamic networks with no reputations, (*D*) clustered dynamic networks with no reputations, (*E*) random dynamic networks with local reputations, (*F*) clustered dynamic networks with local reputations, (*G*) random dynamic networks with global reputations, and (*H*) clustered dynamic networks with global reputations.

where reputational information was available. Rather, what matters is that ties to defectors can be shed.

As expected, Fig. 1 also shows higher cooperation in the clustered static networks than in the random static networks. Clustering structurally insulates conditional cooperators, enabling them to maintain relatively higher rates of cooperation. While we expected clustering to have a stronger effect on static networks than dynamic networks, clustering appears to have no effect on the dynamic networks. Cooperation is high regardless of the underlying topology, so long as that topology is malleable. In essence, dynamic networks prompted noncooperators to become cooperative, or else they became isolated from the network, as elaborated in the hazard model below.

Fig. 2 shows the proportion of cooperation over time by whether the network was static or dynamic (aggregating over reputation information) and whether decisions were specific to each alter (targeted) or applied to all alters simultaneously (diffuse). Targeted or diffuse choices do not affect cooperation in the dynamic networks (column 2). Perhaps this is unsurprising as participants in dynamic networks are able to sever ties as a response to a partner's defection. But, as expected, whether choices are targeted or diffuse does matter for cooperation in static networks: cooperation remains high in networks with targeted decisions, but declines with time in networks with diffuse decisions.

To assess the significance of our manipulated factors on cooperation, we estimated generalized linear mixed models, with decisions nested in rounds, rounds nested in participants, and participants nested in networks (when decisions were diffuse, there was only one decision nested in each round; see *SI Appendix* for subsequent model specifications). We find that participants were more likely to cooperate in dynamic networks than static networks (coeff = 5.47, $P < 0.001$; *SI Appendix*, Table S1, model 1). This replicates several past findings (17–19). We also find that participants were more likely to cooperate in clustered networks (main effect = 1.28, $P = 0.031$), but that this effect is diminished in dynamic networks (interaction effect = −1.62, $P = 0.024$; *SI Appendix*, Table S1, model 1). The *SI Appendix* reports supplemental models of cooperation that include a variety of indicators of participant types, as well as how those types modify the effects of the manipulated factors.

Cooperation is significantly higher in static networks with targeted choices than in static networks with diffuse choices, while the type of choice does not matter for dynamic networks (coeff for interaction between targeted and dynamic = −3.36, $P < 0.001$; *SI Appendix*, Table S1, model 1). Fig. 3 shows that the probability of cooperating in static networks with targeted choices is as high as the probability of cooperating in dynamic networks with either targeted or diffuse choices (within the margin of error). The seemingly innocuous methodological detail of whether participants make targeted versus diffuse choices thus results in a substantial difference in cooperation. As a consequence, prior work may have overestimated the relative strength of dynamic versus static networks in promoting cooperation, due to how it operationalizes cooperation. In short, while our results highlight the importance of network dynamics for the evolution of cooperation that have been documented previously (17–19), they also highlight the importance of network topology in static networks and show that static networks can achieve similarly high rates of cooperation as dynamic networks when participants can make conditional, or targeted, choices.

Defection drove tie deletion in the dynamic networks. Participants opted to drop an alter only 6.3% of the time (*SI Appendix*, Fig. S3), but when they did so and at least one alter defected on the previous round, there was a 99% chance that participants dropped a tie to a defecting alter ($P < 0.001$; conditional logistic regression with SEs adjusted for networks). Once participants decided which alter to cut, they selected a new one. As noted earlier (Fig. 1), dynamics are more important than reputations to cooperation in our setting. However, when they

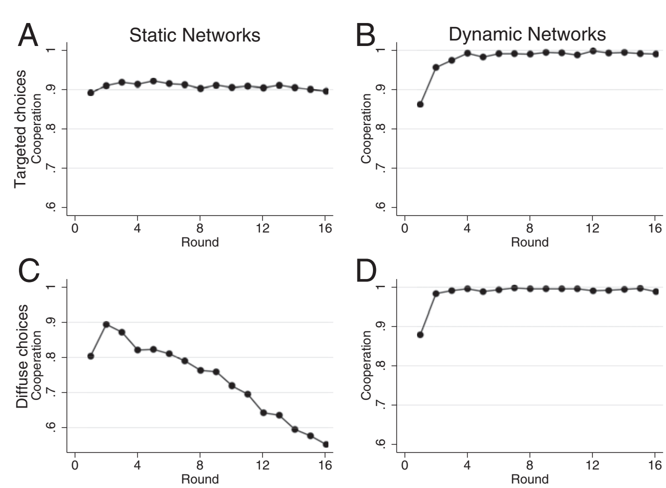

Fig. 2. Proportion of cooperation by round for (*A*) static networks with targeted choices, (*B*) dynamic networks with targeted choices, (*C*) static networks with diffuse choices, and (*D*) dynamic networks with diffuse choices.

SOCIAL SCIENCES

EVOLUTION

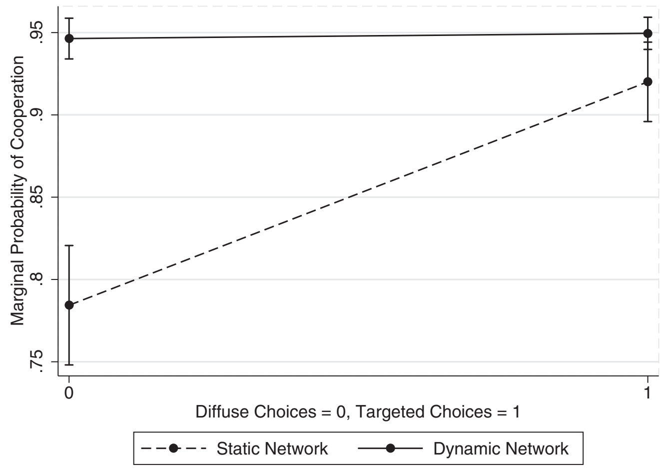

Fig. 3. Marginal probabilities of cooperation illustrating the interaction between static/dynamic networks and diffuse/targeted choices.

were available, reputations still played a role in partner selection. In the global-reputation condition, participants selected new alters with positive reputations (coeff = 3.84, $P < 0.001$; nested conditional logistic regression). In the local-reputation condition, participants were more likely to select an alter when they knew their reputation (coeff = 2.13, $P < 0.001$), and among those selected, participants chose alters with positive reputations (coeff = 5.79, $P = 0.003$).

The forgoing results show that reputations matter for partner selection but not for cooperation rates. To understand this otherwise paradoxical set of findings, note that reputations reduce the risk of participants being exploited by a newly acquired partner, relative to dynamic networks in which reputations of prospective partners are unknown. At the same time, participants in all dynamic network conditions quickly severed ties to defectors. Participants' use of this "out for tat" strategy (1, 33) thus reduced any effects that reputation-based partner selection might have otherwise had on cooperation rates. As a result, network dynamics alone were sufficient to maintain high levels of cooperation.

One alternative explanation for these paradoxical findings is that reputations may promote cooperation, but we are unable to observe the effect because the high rates of cooperation in dynamic networks created a ceiling effect. To rule out this alternative explanation, we ran three additional dynamic network conditions. All of these networks were initially random. Instead of updating ties after every round, participants were told that they would be able to update their ties "periodically (that is, in some rounds)" (19). In all conditions, participants were able to update partners every seven rounds. When they updated partners, they were given either (*i*) no-reputation information, (*ii*) local- reputation information, or (*iii*) global-reputation information, depending on the condition. Cooperation occurred less frequently in these networks than in the dynamic networks in the main experiment (coeff = 2.72, $P < 0.001$; *SI Appendix*, Fig. S8 and Table S7. This allows us to observe reputation effects if they occur over and above the significant effects of dynamic networks. But crucially, there were no differences in cooperation between the three new conditions, indicating that reputations do not promote cooperation beyond the effect of network dynamics, even when there is room for more cooperation. Thus, these results support our conclusions from the main experiment.

Turning back to our primary experiment, the cooperation and reputation results imply that through avoidance, defectors were more likely to become isolated (i.e., have no ties) from the network in the reputation conditions. Fig. 4 shows the number of participants that became isolates through time in each of the dynamic network conditions. More participants became isolated from networks in the local-reputation condition than in the condition with no-reputation information, and even more participants became isolates in the global-reputation condition. We modeled becoming an isolate with a Cox proportional hazard model with clustered SEs within networks. After controlling for participant defection as a time-varying covariate, which itself predicts becoming an isolate (coeff = 2.33, $P < 0.001$), results show that participants in the local-reputation condition are more likely to become isolated than participants in the no-reputation condition (coeff = 0.35, $P = 0.12$), and that participants in the global-reputation condition are significantly more likely to become isolated than participants in the no-reputation condition (coeff = 0.81, $P = 0.004$; *SI Appendix*, Table S3, model 1 and *SI Appendix*, Fig. S5). In sum, participants shed ties to defectors and then use reputations to select new alters. The result of this process is that those participants who defect are more likely to become isolated from the network as their reputation precedes them.

As noted above, participants in dynamic networks opted to replace relatively few of their ties. On average, they replaced a tie 6.3% of the time. Most of this change occurred within the first few rounds (*SI Appendix*, Fig. S4). After getting feedback about others' cooperation in round 1, more than 1 in 3 participants replaced a tie, but by round 4 only 1 in 20 did so. Thus, the dynamic networks quickly converged to a stable state of high cooperation (Fig. 1). This stabilization of the networks is apparent in the amount of clustering observed in the dynamic networks. As illustrated in Fig. 5, we find that the overall clustering levels off quickly. There is a sizable increase in clustering in the local-reputation condition (coeff. = 0.118, $P < 0.001$; Fig. 5*B* and *SI Appendix*, Table S5). As participants use reputations to form new ties, and they only know the reputations of their alters' alters, clustering increases. Thus, these results show strong

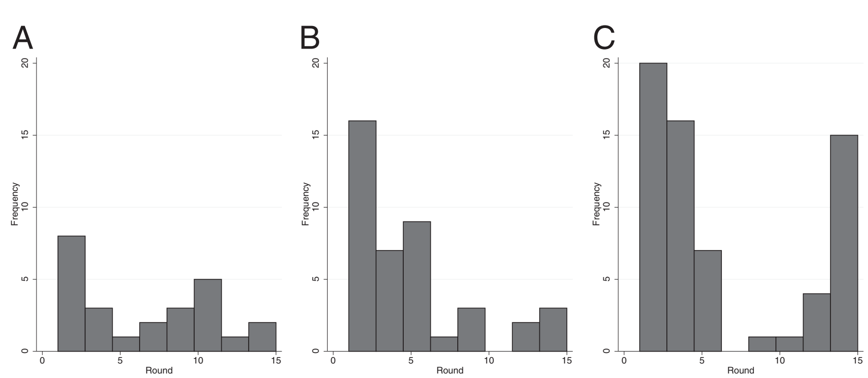

Fig. 4. Histograms of the number of participants who became isolates by round in dynamic networks with (*A*) no reputations, (*B*) local reputations, and (*C*) global reputations. There were 25 isolates in *A*, 41 in *B*, and 64 in *C*.

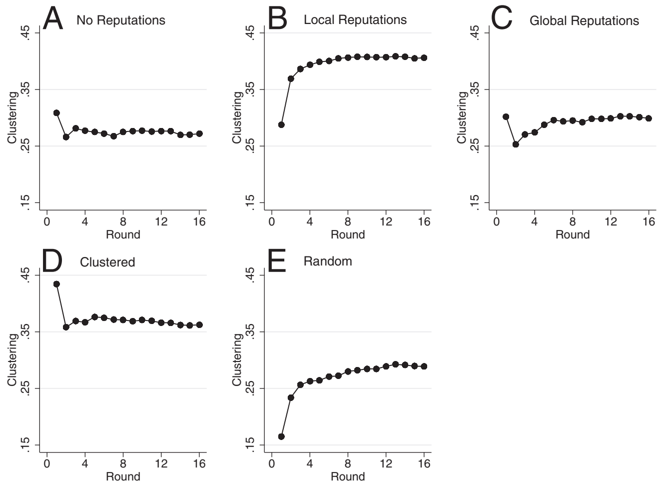

Fig. 5. Clustering through time in dynamic networks (*A*) with no reputations, (*B*) with local reputations, (*C*) with global reputations, (*D*) that were initially clustered, and (*E*) that were initially random.

evidence of triadic closure, a key micromechanism guiding the emergence of stable clusters (31). There is also a sizable difference in clustering between those networks that were initially clustered (Fig. 5*D*) and those that were initially random (Fig. 5*E*). However, the initially clustered networks (Fig. 5*D*) decrease in clustering early on, as defectors move to the periphery of the network, and the initially random networks increase in clustering (Fig. 5*E*), as cooperators seek out one another. This finding underscores the interplay between clustering and network dynamics. In particular, even in the initially random networks with no reputation information, the amount of clustering doubles during the experiment, but levels off as participants either become cooperative or get excluded from the network. As this occurs, the structure stabilizes, precluding further clustering.

The pattern of earnings by conditions mirrors the cooperation rates shown in Fig. 1 (*SI Appendix*, Fig. S7). Participants in dynamic networks earn substantially more than those in static networks (coeff = 14.86, $P < 0.001$; *SI Appendix*, Table S6, model 1), and in the dynamic networks, reputations do not impact earnings. Consistent with the patterns for cooperation in Fig. 1, whether the networks were initially random or clustered does not affect the earnings of participants in dynamic networks, but participants in clustered static networks earned more than those in random static networks (coeff = 3.63, $P = 0.035$). Also consistent with patterns of cooperation in Fig. 2, whether decisions were diffuse or targeted does not affect earnings in dynamic networks, but participants in static networks with targeted choices earned more than participants in static networks with diffuse choices (coeff = 8.36, $P < 0.001$).

To summarize, as in related work (19), we find very high levels of cooperation in our dynamic networks. Previous results from dynamic networks left open the possibility that high rates of cooperation were driven by either dynamics or reputation processes. But we find that only network dynamics, not reputation information, increases cooperation. When they are available, participants use reputations when seeking out new ties, but we found no evidence that the availability of reputational information increases cooperation or earnings above the effect of dynamics (29). Reputations, however, result in defectors quickly moving to the periphery of the network and to an increased hazard of being isolated from the network altogether. Earlier, we noted that reputations can impact partner selection and network dynamics without affecting cooperation rates because participants in dynamic networks employ an out-for-tat strategy (1, 33) or "tie reciprocity" (18, 22, 34), whereby they quickly sever ties to uncooperative alters. This results in high levels of cooperation, without the need for reputational information. We also find high rates of cooperation in clustered static networks, and in static networks where participants make targeted choices. These latter results highlight the importance of network topology and suggest that seemingly simple design decisions can have important consequences for our understanding of the extent to which dynamic networks promote cooperation.

Discussion

This research makes several key contributions to our understanding of the power of social networks to shape cooperation in human populations. First, several prior studies have demonstrated that dynamic networks promote cooperation, but we cannot know whether these effects are due to network dynamics, reputation processes, or both. Humans are more apt to cooperate when interacting with someone with a prosocial reputation (24, 35). Thus, when new relations form in the presence of reputational information, cooperation may be driven by either structural change or reputation-based expectations of cooperation. Past work has shown that networks with reputational information increase cooperation, but that study had a very high level of fluidity, where all possible ties could be formed or broken from round to round (21). That is, the networks were created anew each round and thus did not constitute a stable structuring of relationships. In contrast, we find that reputations do not affect decisions to cooperate (29). All our dynamic networks—regardless of reputational information—yielded higher levels of cooperation than our static networks with diffuse choices. The high level of cooperation we found in dynamic networks is consistent with other studies of PD games (36). While some studies of dynamic networks (18, 22) show somewhat lower levels of cooperation than we observe in the main experiment, those studies allowed far less frequent tie updates. Studies that allow participants to alter ties at more realistic rates reveal near total cooperation. These results are in line with ours, but we clarify when and how reputation processes combine with dynamics to promote cooperation.

Second, while we did not find evidence, either in the main experiment or in the supplemental conditions, that reputations matter for decisions to cooperate or defect, they were important for partner selection (24). In our global-reputation condition, we found that participants selected prospective alters as their reputations increased. Similarly, in the novel local-reputation condition we found that participants were more likely to select an alter about whom they had reputational information, and among them, they were more likely to select alters with prosocial reputations. The clusters that emerged were generated by triadic closure, bringing our understanding of dynamic cooperation networks in line with the clustering that often characterizes real-world social networks (31).

Third, our work underscores the overall importance of network topology and its interaction with network dynamics for the evolution of cooperation. In static networks, network topology has a large effect on cooperation. Clustered static networks showed substantially higher levels of cooperation than random static networks (18). In dynamic networks, the initial structure of relations is less important, but the fact that the topology is malleable has a very strong effect on cooperation. That is, in both static and dynamic networks, assortment seems to promote cooperation. Exogenous clustering in static networks, and endogenous clustering in dynamics ones, results in increased cooperation.

Finally, we show that the rules governing participants' decisions to cooperate or defect has important implications for cooperation. When our participants made a single decision that applied to all their partners, as is typical in the existing literature (18, 19, 22), we observed significant differences in cooperation between static and dynamic networks. However, when our participants made targeted choices for each of their partners, cooperation rates in static networks were as high as cooperation rates in our dynamic networks. As we argued above and our results showed, making a single decision to either cooperate or defect with all partners decreases cooperation in static networks, because participants have no other recourse to punish alters that

SOCIAL SCIENCES

EVOLUTION

have previously defected. Given that prior work on networks and cooperation generally requires that participants make a diffuse decision that applies to all actors, we argue that existing findings overstate the degree to which network dynamics promote cooperation, relative to static networks. In particular, we find that static networks, where people can make targeted decisions about whether to cooperate or defect in each of their relations yield cooperation levels that rival those typically observed in dynamic networks. Future research on networks and cooperation should mind this important design decision.

In closing, we provide evidence for the importance of network topology and decision processes that shape cooperation in situations characterized by conflicts between individual and collective interests. Our work builds on and extends related work on network dynamics, but highlights the importance of dynamics relative to reputational processes. In static networks, the ability to decide whether to cooperate or defect with each partner results in sustained cooperation, at rates similar to those we observed in dynamic networks. In dynamic networks, we find strong structural effects and very high levels of cooperation throughout our study, with rather negligible effects of reputation. That is, we show that networks—both static and dynamic—can yield sustained cooperation among humans.

Materials and Methods

The Institutional Review Board at the University of South Carolina reviewed and approved this research. There was no deception. The experiment was conducted using Amazon Mechanical Turk, an online crowd-sourcing platform that is frequently used for behavioral experiments (37–39). Workers from Amazon Mechanical Turk read an online informed consent form detailing the study procedures, approximate length of the study, and their expected payment for participating. Those who agreed to participate followed a link to a custom Web app that went over the instructions, embedded them within networks, and recorded their behaviors as they interacted with other workers. Several comprehension check items, with feedback, were included in the instructions (*SI Appendix*). Data were collected in the spring and summer of 2017. Data for the follow-up conditions were collected in the autumn of 2017. Each session lasted ∼30 min. Participants were paid a $2.50 show-up fee and were awarded a bonus based on how many MUs they acquired throughout the study (1,000 MUs = $1.00). The data were modeled using random intercept multilevel or mixed effects models. Cooperation was fit using the logistic link with alters nested in rounds, rounds nested in participants, and participants nested in networks. Clustering and earnings were fit using an identity link. The network-level clustering coefficient was measured at each round, and earnings for each round were nested in participants, who were nested in networks. Finally, participant choices of which alter to drop and which new alter to add were modeled using conditional logistic regression with SEs adjusted for multiple decisions within networks.

ACKNOWLEDGMENTS. We thank Mark McKnight for programming assistance and granting us access to Breadboard. The data reported in this paper are archived with Harvard Dataverse. The research reported here was funded in whole under Award W911-NF-15-1-0131 from the US Army Research Office/Army Research Laboratory. The views expressed are those of the authors and should not be attributed to the Army Research Office/Army Research Laboratory.

1. Kollock P (1998) Social dilemmas: The anatomy of cooperation. *Annu Rev Sociol* 24:183–214.
2. Granovetter M (1985) Economic action and social structure: The problem of embeddedness. *Am J Sociol* 91:481–510.
3. Rand DG, Nowak MA, Fowler JH, Christakis NA (2014) Static network structure can stabilize human cooperation. *Proc Natl Acad Sci USA* 111:17093–17098.
4. Nakamaru M, Matsuda H, Iwasa Y (1997) The evolution of cooperation in a lattice-structured population. *J Theor Biol* 184:65–81.
5. Axelrod R (1984) *The Evolution of Cooperation* (Basic Books, New York).
6. Hamilton WD (1964) The genetical evolution of social behaviour. II. *J Theor Biol* 7:17–52.
7. Nowak MA (2006) Five rules for the evolution of cooperation. *Science* 314:1560–1563.
8. Nowak MA, May RM (1992) Evolutionary games and spatial chaos. *Nature* 359:826–829.
9. Nowak MA, Sigmund K (1998) Evolution of indirect reciprocity by image scoring. *Nature* 393:573–577.
10. Ohtsuki H, Iwasa Y (2006) The leading eight: Social norms that can maintain cooperation by indirect reciprocity. *J Theor Biol* 239:435–444.
11. Rand DG, Nowak MA (2013) Human cooperation. *Trends Cogn Sci* 17:413–425.
12. Traulsen A, Nowak MA (2006) Evolution of cooperation by multilevel selection. *Proc Natl Acad Sci USA* 103:10952–10955.
13. Trivers RL (1971) The evolution of reciprocal altruism. *Q Rev Biol* 46:35–57.
14. Wilson DS (1975) A theory of group selection. *Proc Natl Acad Sci USA* 72:143–146.
15. Yamagishi T, et al. (2012) Rejection of unfair offers in the ultimatum game is no evidence of strong reciprocity. *Proc Natl Acad Sci USA* 109:20364–20368.
16. Assenza S, Gómez-Gardeñes J, Latora V (2008) Enhancement of cooperation in highly clustered scale-free networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 78:017101.
17. 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.
18. Rand DG, Arbesman S, Christakis NA (2011) Dynamic social networks promote cooperation in experiments with humans. *Proc Natl Acad Sci USA* 108:19193–19198.
19. Wang J, Suri S, Watts DJ (2012) Cooperation and assortativity with dynamic partner updating. *Proc Natl Acad Sci USA* 109:14363–14368.
20. Jordan JJ, Rand DG, Arbesman S, Fowler JH, Christakis NA (2013) Contagion of cooperation in static and fluid social networks. *PLoS One* 8:e66199.
21. Gallo E, Yan C (2015) The effects of reputational and social knowledge on cooperation. *Proc Natl Acad Sci USA* 112:3647–3652.
22. Shirado H, Fu F, Fowler JH, Christakis NA (2013) Quality versus quantity of social ties in experimental cooperative networks. *Nat Commun* 4:2814.
23. Barclay P (2013) Strategies for cooperation in biological markets, especially for humans. *Evol Hum Behav* 34:164–175.
24. Roberts G (2015) Partner choice drives the evolution of cooperation via indirect reciprocity. *PLoS One* 10:e0129442.
25. van Dolder D, Buskens V (2014) Individual choices in dynamic networks: An experiment on social preferences. *PLoS One* 9:e92276.
26. Coricelli G, Fehr D, Fellner G (2004) Partner selection in public goods experiments. *J Conflict Resolut* 48:356–378.
27. Greif A (2006) *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade* (Cambridge Univ Press, Cambridge, UK).
28. Melamed D, Simpson B, Harrell A (2017) Prosocial orientation alters network dynamics and fosters cooperation. *Sci Rep* 7:357.
29. Corten R, Rosenkranz S, Buskens V, Cook KS (2016) Reputation effects in social networks do not promote cooperation: An experimental test of the Raub & Weesie model. *PLoS One* 11:e0155703.
30. Melamed D, Simpson B (2016) Strong ties promote the evolution of cooperation in dynamic networks. *Soc Networks* 45:32–44.
31. Davidsen J, Ebel H, Bornholdt S (2002) Emergence of a small world from local interactions: Modeling acquaintance networks. *Phys Rev Lett* 88:128701.
32. Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393:440–442.
33. Yamagishi T, Hayashi N, Jin N (1994) Prisoner's dilemma networks: Selection strategy versus action strategy. *Social Dilemmas and Cooperation* (Springer, Berlin), pp 233–250.
34. Boyd R, Richerson PJ (1988) The evolution of reciprocity in sizable groups. *J Theor Biol* 132:337–356.
35. Sylwester K, Roberts G (2010) Cooperators benefit through reputation-based partner choice in economic games. *Biol Lett* 6:659–662.
36. Milinski M, Semmann D, Krambeck H-J (2002) Reputation helps solve the 'tragedy of the commons'. *Nature* 415:424–426.
37. Buhrmester M, Kwang T, Gosling SD (2011) Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspect Psychol Sci* 6:3–5.
38. Weinberg JD, Freese J, McElhattan D (2014) Comparing data characteristics and results of an online factorial survey between a population-based and a crowdsource-recruited sample. *Sociol Sci* 1:292–310.
39. Rand DG (2012) The promise of mechanical turk: How online labor markets can help theorists run behavioral experiments. *J Theor Biol* 299:172–179.