

Social Networks, Social Context, and Political Participation: Evidence from Uganda

Nicholas Eubank,^{*}
Guy Grossman,[†]
Melina Platas,[‡]
Jonathan Rodden[§]

September 17, 2017

Preliminary Draft

Comments Welcome

[Click here for most recent version](#)

Abstract

The tendency for citizens to mirror the behavior of their peers is well-documented, but our understanding of the *dynamics* of this influence is more limited. For example, there is ample evidence that voter turnout is shaped by peer effects, but in what settings does the choice of one person to vote cascade through a community and lead to high voter turnout? Scholars have made substantial theoretical inroads into this question in recent years (Fowler, 2005; Siegel, 2009; Rolfe, 2012; Larson et al., 2017; Sinclair, 2012), but direct empirical tests of these theories remain scarce. Using data on the full social networks of 15 villages in rural Uganda, this paper develops theoretical predictions about expected cross-village variation in turnout based on the network structure of each village, and demonstrates that these predictions are tightly linked with actual turnout in low-salience local elections, for which a lack of mass media coverage means that individuals must rely on their social networks to learn about the likelihood that others will vote. These results provide the first direct empirical validation of the “social context” theory long suggested by scholars like Fowler (2005), Siegel (2009), and Rolfe (2012), and introduce a finding of importance for future empirical network research: the salience of social networks may be conditional on the information environment.

^{*}Center for the Study of Democratic Institutions, Vanderbilt University, nick@nickeubank.com

[†]Department of Political Science, University of Pennsylvania and EGAP, ggros@sas.upenn.edu

[‡]Department of Political Science, New York University Abu Dhabi, mplatas@nyu.edu

[§]Department of Political Science, Stanford University, jrodden@stanford.edu

Individual characteristics are often inadequate predictors of behavior, including many forms of political participation (Huckfeldt and Sprague, 1995; Kenny, 1992; Leighley, 1990; McClurg, 2003). Rather, the cause of individuals' actions is often traced to the behavior of their peers (Ioannides, 2013). Social influence, or peer effects, does not have to be purposive and may even occur unconsciously (Cialdini, 2015; Ross and Nisbett, 2011). While the impact of social ties on individuals' behavior has been well-documented, empirical tests of the real-world *dynamics* of social influence are more limited. An important question remains unanswered: Under what circumstances will the actions of a small number of innately motivated individuals snowball, such that their actions lead to participation by their peers, leading to further participation by their peers, and so on? And when will such individual actions remain isolated or peter out before diffusing widely?

In recent years, a growing number of scholars have argued that the answer to this question depends on the structure of social networks (Siegel, 2009; Sinclair, 2012; Rolfe, 2012; Fowler, 2005; Larson et al., 2017). While our social *ties* define who is likely to directly influence our behavior, the structure of whole networks determines how changes in behavior in small groups may propagate to the broader population. Drawing on simulations based on simple and empirically grounded behavioral assumptions, scholars have shown that variations in networks structure can substantially impact the “contagiousness” of political participation that starts from a few individuals, holding other factors constant.

However, empirical tests of the above proposition—that the structure of social networks shapes levels of participation—remain limited. This is mostly due to the difficulty and cost associated with collecting the data on the structure of social networks required for rigorous empirical testing. As a result, most empirical work on the role of network structure has either relied on ego-network data; i.e., reports by survey respondents on their friends, with no linking across respondents to create a full network (Abrams et al., 2011), or indirect tests of the assumptions that underlie theories about the importance of networks. For example, Nickerson (2008) examined whether canvassing an individual increases the likelihood that other household members turn out to vote.

This study aims to bridge the divide between theory and empirics in the study of the network structure and voter turnout. We build on the work of Siegel (2009), Rolfe (2012) and Fowler (2005) to development measures of theoretically-predicted levels of political participation based on the structure of networks in a cross-section of villages in rural Uganda. We then correlate these predictive measures with actual voter turnout as reported in official electoral data. We find that the highest levels of turnout are indeed found in villages that have network structures that, according to theory, should support the highest levels of participation. We also find these same villages have substantially higher attendance at village meetings, and supportive evidence of higher contributions to village projects. Using lab-in-the-field behavioral games, we rule out the possibility that differences in village pro-sociality are driving both differences in network structure and political participation.

Consistent with theory, we find that social networks appear to be substantially more important in the lower-salience elections for district chairperson than in the high-salience presidential election. Media attention and campaign events are far more limited for local elections, and thus social networks are more important for learning about peer behavior. In addition to providing ‘real-world’ evidence that network structure is driving spatial variation in political behavior, our core findings also have important implications for future network research. Indeed, if the salience of networks is contingent on the availability of non-network-based information (like mass media coverage), then future studies must consider how the study context impacts the conclusions that can be drawn about the role of networks structure. Moreover, this result also suggests fruitful avenues of research into how increasing dependence on network-based social media rather than mass media may change the dynamics of political behavior.

1 Social Context Theory

Our analysis focuses on the “social context” model of social influence. Social context theory posits that an individual’s likelihood of political participation is determined by two components: a personal disposition to participate (affected by factors such as income, gender and education), and the level of participation among one’s peers (Fowler, 2005; Siegel, 2009; Rolfe, 2012). According to this model, while some people will always be inclined to participate politically (individuals labeled “unconditional” decision makers by Rolfe (2012) and “rabble-rousers” by Siegel (2009)), others will only participate if they observe high enough levels of participation among their peers.

The influence of social context has been extensively documented in psychology experiments (Ross and Nisbett, 2011), and it is worth emphasizing that while in some cases mirroring behavior may be the result of bayesian updating by rational agents about the desirability of a behavior or strategic social conformity (Goyal, 2012), research also suggests this dynamic may not be fully conscious. In one illustrative social context experiment cited by Rolfe (2012), a researcher watched individuals walking by a street performer to see if they would put money in the performer’s hat (Cialdini, 2015). After the subjects passed, the researcher would then ask the subjects why they had chosen to either give money to the performer or not. In some cases, the researcher would have a confederate pass in front the unknowing subject and throw in money. The researcher found that subjects were *eight times* more likely to contribute if a confederate contributed in front of them, and yet during interviews, almost none of the subjects recalled seeing someone contribute. In other words, while subjects were strongly driven to mirror the behavior of others, they seemed relatively unaware of this influence.

But while this mechanism of influence has been well-documented among individuals, the dynamics of how it may diffuse to larger populations has received less

empirical attention. This is due to the fact that social context models assume fundamental interdependencies in behavior that require the use of complete network analysis for studying macro social influence processes. If we wish to understand how the behavior of a few rabble-rousers may or may not propagate across a population, it is not enough to just look at individuals and their immediate peers. Rather, we must work with full networks so we can examine how the higher-order topological features of network structure shape not only who we interact with directly, but also how our influence may potentially spread beyond our immediate contacts to the the broader network.

One core theoretical result of network scholars is that there are no easy answers when it comes to predicting how social influence may spread through a network (Centola and Macy, 2007; Jackson and Yariv, 2010). Simple measures like average number of connections or average shortest paths cannot explain whether the actions of a few people in a network will have large effects or not. Rather, influence dynamics are shaped by numerous topological features of social networks (Centola, 2015), and simulation remains the primary method of determining how a given network will support diffusion processes.

To date, however, it has been difficult to corroborate the results of theory-driven simulations due to the paucity of real-world network data. To test diffusion theories requires not only data on *one* full network, but data on the full networks of *multiple communities* along with community-level measures of political participation to allow for cross-sectional analysis. And it is this gap in the existing literature this study aims to fill.

We accomplish this by estimating the theoretically-predicted level of political participation (TPP) based only on the structure of village networks using the theoretical insights of Siegel (2009) and Rolfe (2012). These values of TPP are then correlated with actual turnout for two types of elections that took place in Uganda in 2016.

2 Data

This analysis relies on two primary sources of data: network data collected as part of an original survey, and precinct-level data on turnout in Uganda’s 2016 Presidential elections and in elections for the chief executive (chairperson) of the district government, the highest subnational tier of government in Uganda below the central government.

2.1 Network Data

We collect data from 16 Ugandan villages that took part in a multi-year program called Governance, Accountability, Participation, and Performance (GAPP), which

was implemented by RTI International and funded by the United States Agency for International Development (USAID). Surveying all village households, we are able to construct 16 independent ‘whole’ networks by using a simple name generator technique (Knoke and Yang, 2008), eliciting information on respondents’ familial and friendship ties, as well as ties to village money lenders and more generally, local ‘problem solvers’. See Ferrali et al. (2017) for full survey details.

These network surveys are used to compute empirical networks: the *Friends and Family* network, which consists of all connections listed as “friends” or “family”, and the *Union Network*, which consists of the friends and family network plus ties reported as people the respondent “would go to if they had to borrow money” and people he or she “would go to in order to solve a problem regarding public services in the village.” All networks are undirected, unweighted, and do not require reciprocity of ties. Summary statistics for final empirical networks are presented in Table 1. Note that results are also consistent – and in fact stronger – when limiting attention to reciprocated ties among friends and family, although we feel allowing for non-reciprocated ties generates more meaningful networks (See Appendix E for further discussion).

Throughout this analysis, attention is restricted to 15 of the 16 villages originally included in the survey. This is because the 16th village is substantially smaller than any other village under consideration. While the 15 core villages have between 160 and 283 residents, the omitted village network has only 30 people. However, our core findings are robust to the inclusion of the 16th village, as shown in Appendix D.

Table 1: Network Summary Statistics

	Friends and Family	Union
Average Size	210.3	210.3
Average Num Connections	1,231.5	1,693.9
Average Degree	5.8	8.0
Min Size	160.0	160.0
Max Size	283.0	283.0

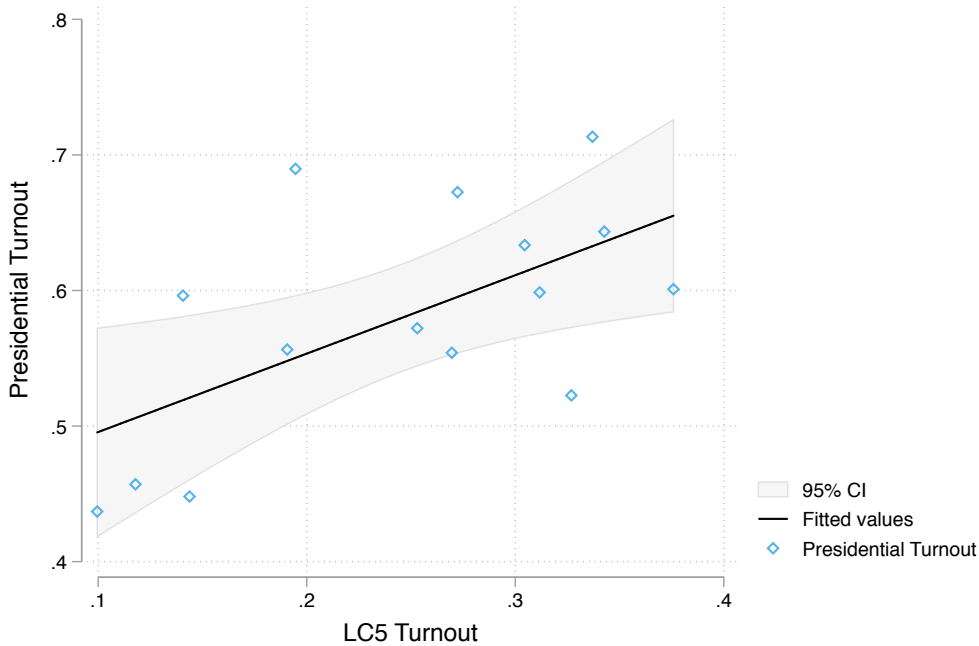
2.2 Turnout Data

Data on turnout comes from precinct-level electoral returns and the official voter register as compiled by the Electoral Commission of Uganda. Because precincts, or polling stations, do not correspond precisely to Ugandan villages, we extrapolate turnout using the voter registration data, which provides information on the precinct at which residents of each village are registered. In particular, our analysis below relies on the assumption that votes cast at each precinct were cast by residents of neighboring villages in proportion to each village’s share of voters registered at the precinct.

To illustrate, assume that at Precinct 1, 200 votes were cast. If 75% of the voters registered at Precinct 1 come from Village A and 25% come from Village B, we assume that Village A contributed 150 votes and Village B contributed 50 votes. Turnout for each village is then calculated as the sum of votes we infer to have been cast by its residents at all Precincts. A more formal presentation and summary statistics of the village-to-precinct mapping are presented in Appendix A. By assuming that at a given precinct all villages have the same turnout, this estimation should bias our analysis in favor of *not* finding differences in turnout across villages.

Figure 1 shows distribution of turnout as share of adult population for the Presidential and district (LC5) chairperson elections across villages. Average turnout for the Presidential election in our data was 60% (compared to 68% nationally); average turnout for the district chairperson was 25% (compared to 31% nationally). Turnout across the elections is correlated at 0.61%, suggesting the elections are distinct but that villages with high turnout tend to have high turnout across elections.

Figure 1: Presidential & LC5 Turnout



3 Simulating Influence Dynamics on a Network

Social context theory is premised on the assumption that individuals are more likely to participate politically if their peers participate politically. To understand the dynamics of how this assumption shapes behavior on different types of networks, we

use a simulation model closely related to the behavioral models of Siegel (2009) and Rolfe (2012).

The starting point of the simulation is that vertices $v \in V$ in a network choose whether or not to participate in a political activity. Initially, the simulation begins with all vertices endowed with some *individual* proclivity to participate. All vertices are assumed to begin in a state of non-participation at $t = 0$, but in the first stage, vertices (or nodes) with very high individual proclivities begin to participate. In each state of the simulation, vertices observe the behavior of their peers and then decide whether to also participate. If a vertex finds that all her peers are participating, she becomes more likely to participate herself. The simulation then continues this cycle of vertices observing their peers, updating their own behavior, then observing their peers once more until the network converges to a stable configuration in which behavior no longer changes between stages.

More specifically, the simulation proceeds as follows:

Stage 1: Model Initialization

- All vertices are assigned an individual propensity to participate $\beta_v \sim Normal(\beta_{mean}, \beta_{sd})$. Once randomly assigned, these values are fixed for the duration of the simulation.
- All vertices begin in a state of non-participation ($participation_{v,0} = 0 \forall v \in V$)

Stage 2: Social Influence Simulation

- At each step of the simulation $t \in T$, each vertex $v \in V$ calculates a *local participation rate* ($lpr_{v,t}$) which is equal to the share of the people connected to v at time t in the network who also plan to participate. These two factors can then be combined in a simple decision rule: $participation_{v,t+1} = 1$ if $\beta_v - (1 - lpr_{v,t}) > 0$.
- Calculate overall anticipated participation $OverallParticipation_t = \frac{\sum_v participation_v}{|V|}$.
- Repeat Stage 2 until the value of *OverallParticipation* converges.¹

Several aspects of this framework are worth noting. First, individuals with high values β (specifically, $\beta > 1$) will participate politically *even if none of their immediate neighbors plan to participate*. Similarly, individuals with very low values of β ($\beta < 0$) will never participate, even if all of their peers are participating. For anyone with a value of $\beta \in (0, 1)$, there is a threshold level of peer participation that will induce those individuals to be participating. For example, if $\beta_v = 0.5$, then v will participate if and only if at least 1/2 of her friends participate.

The second aspect of this model is that it is generally dynamic. We begin in a state of non-participation at time $t = 0$, then in the first period only people with $\beta > 1$ will participate. But as people with $\beta > 1$ announce they are participating, that changes

¹Additional technical details of simulation can be found in Appendix B.

the value of lpr for everyone connected to one of these rabble-rousers, potentially leading them to plan to participate as well. These spillovers may—but also may not, *depending on network structure*—cascade for a period of time before eventually the network stabilizes into an equilibrium level of political participation, which may occur at any level between no one participating and everyone participating.


The focus of our analysis is on the average level of political participation to which this model converges for given values of β_{mean} and β_{sd} . This is calculated by simulating this process of influence repeatedly on the network of each village until the simulation converges, then calculating the average level of participation at these convergent states. For a given pair of parameters β_{mean} and β_{sd} , if one village’s network converges to higher levels of participation than the network of another village, then our theoretical prediction is that the village that converged to higher levels of political participation should also have higher levels of actual observed turnout.

3.1 TPP and Network Structure


The use of simulations is motivated by the fact basic network properties (like average number of connections or even full degree distribution) do not map cleanly onto the ability of a network to diffuse social influence. This is illustrated in Figure 2, which shows two networks that have identical degree distributions (degree is the number of edges connected to a given vertex), but have very different equilibrium levels of participation, even given the same distributions of individual proclivities to participate. Indeed, as shown in Table 2, which plots simulation results on these two graphs, the more efficient network has on average 7.8% higher equilibrium participation, and as much as 16.1% higher participation under some parameter values.

Figure 2: Identical Degree Distributions, Different TPP

(a) Efficient Network



(b) Inefficient Network



Both networks have the same number of vertices (16), edges (64) and degree distributions (all vertices are of degree 4), but even given the same distribution of individual proclivities to participate, in the more efficient network equilibrium participation is 7.8% higher.

Table 2: Simulated TPP

β_{mean}	β_{sd}	Efficient	Inefficient	Percent Difference
0.50	0.50	0.36	0.35	0.01
0.60	0.50	0.51	0.50	0.02
0.60	0.25	0.26	0.23	0.12
0.70	0.50	0.70	0.65	0.08
0.70	0.25	0.69	0.59	0.16

This is not to say that different network properties may not be highly correlated. Indeed, in our data, the correlation between average degree and index of simulated equilibrium participation is 0.96 for the Friends and Family network and 0.95 for the Union network. Nevertheless, this example illustrates these will not always be the case, and cannot be taken for granted.

3.2 Simulation Result Summary

Following Siegel (2009), we focus on parameter values of $\beta_{mean} \in \{0.5, 0.6, 0.7\}$ and $\beta_{sd} \in \{0.25, 0.5\}$.² This is due to the fact that significantly higher values of β_{mean} tend to result in convergence to full participation, while substantially lower values lead to non-dynamic simulations (those with values of $\beta > 1$ participate, but they are rare and others tend to have very low proclivities to participate, as a result of which almost no vertices flip from non-participation to participation). Similarly, larger values of β_{sd} increase the share of individuals whose behavior is unaffected by the behavior of other so much that the simulations tend not to be dynamic. In these non-dynamic settings, all networks are essentially comparable, as participation ends up being roughly equal to the share of nodes with $\beta_{mean} > 1$, which is the same for all networks in expectation. Average TPP scores across villages for different parameter values are presented in Table 3.

In addition to giving rise to interesting dynamics, the values of β studied here also seem plausible given our study context. Though individual β values have no mechanical relationship with TPP (any effect is mediated by the network), note that an individual with a $\beta = 0.7$ will only participate if at least 30% of his peers are participating, while an individual with $\beta = 0.5$ will participate if at least 50% of his peers participate. Given observed equilibrium turnout rates of 60% (for the Presidential election) and 25% in the LC5 election, these values seem substantively reasonable.

Table 3: Average Theoretically-Predicted Participation (TPP)

β_{mean}	β_{sd}	Mean, Friends & Family	Mean, Union
0.50	0.50	0.41	0.42
0.60	0.50	0.64	0.66
0.60	0.25	0.48	0.45
0.70	0.50	0.81	0.83
0.70	0.25	0.98	0.98

Moreover, the inter-village correlation in TPP scores across this parameter space is quite high, as shown in Tables 4 and 5 in Appendix C. The overall average correlation for the Friends and Family network is 0.80 and for the Union network is 0.85, and the correlation across the two network types is 0.80.

²We exclude one parameter pair from those sets ($\beta_{mean} = 0.5, \beta_{sd} = 0.25$), as it generates almost no unconditional participators, and thus no dynamics.

Thus for ease of exposition, most following results will be presented using an index constructed as the first principle component of these statistics for each network type.

4 Social Context and Turnout

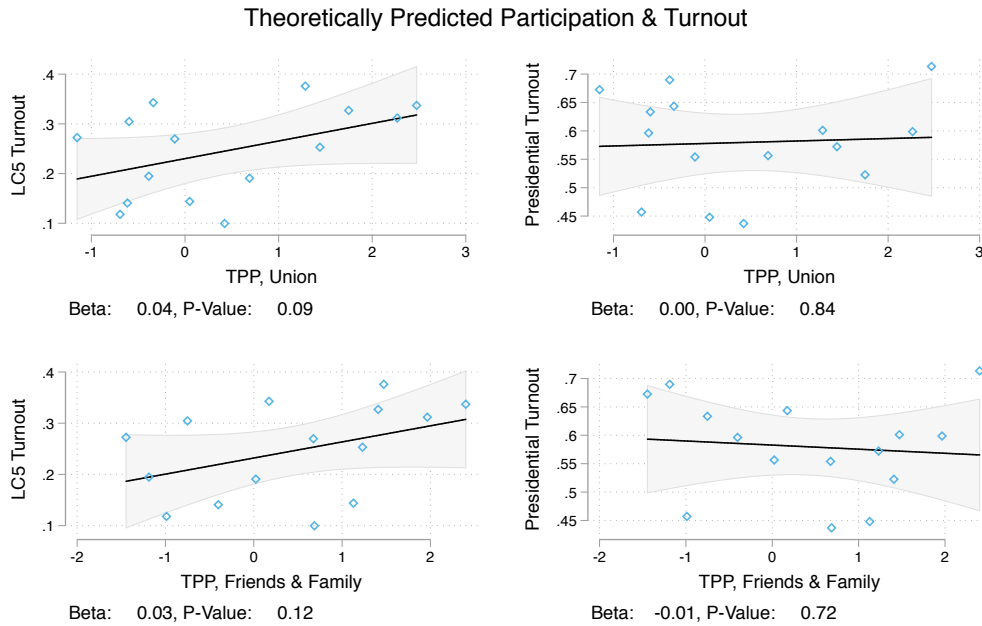
Figure 3 presents the bivariate correlation between the theoretically-predicted participation (TPP) and actual turnout in the Presidential and district (LC5) elections. Results are replicated for both the Friends and Family network and the Union network.

First, in both LC5 specifications, TPP is positively correlated with electoral turnout, as predicted by theories of how social influence spreads on networks. Second, despite a small sample size of 15 villages, results are significant using the Union network specification, and nearly significant with the Friends & Family network specification. Moreover, the correlation between these factors appears relatively uniform — results are not being driven by outliers, as is the risk in small-N studies — and the relative consistency across specifications provides further evidence of a genuine relationship.

These correlations are also robust to the introduction of a number of different controls and sample restrictions. As shown in Appendix D, results are consistent when controlling for ethno-linguistic fractionalization or the share of each village that has completed primary school, when restricting the sample to the set of villages for which our estimates of turnout are likely most accurate, and are only weakened somewhat by the inclusion of the exceptionally small 16th village discussed in Section 2.

We interpret the correlation in the LC5 elections as lending support for network structure theories of social influence. In Uganda, the presidential election is a much higher-salience election that garners substantially greater media attention and entails far more campaign efforts. It seems likely that voters are exposed to information about the likelihood that peers and non-peers will turn out from many non-network channels, such as election rallies. As a result, the specific topology of village networks should matter less for shaping the social contexts that influence voter turnout decisions (though the correlation is still in the predicted direction). In the lower salience LC5 election, by contrast, a larger share of the information voters receive about anticipated participation may come through their day-to-day interactions and conversations, which are largely dictated by their social networks, increasing the observed correlation between network structure and turnout.

Figure 3



TPP in standard deviations of PCA index across parameter values, turnout in shares.

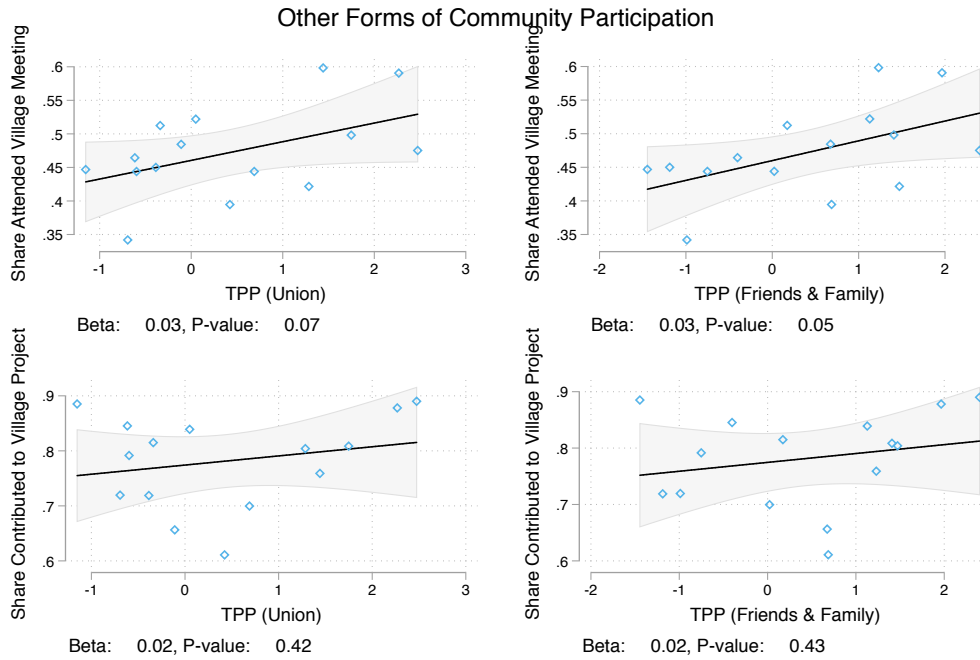
Correlation between index of equilibrium simulated participation rates and voter turnout.

4.1 Other Forms of Political Participation

While the focus of much research on social influence and networks has focused on voter turnout, the theory is agnostic about the specific form of political participation being fostered. With that in mind, we further examine the relationship between TPP and self-reported information on participation in village governance. In particular, we find that TPP correlates with (a) the share of villagers reporting having attended a village meeting and (b) the share who report having contributed (in time or labor) to a village project. These results are presented in Figure 4.

Consistent with theory, we find that across network specifications and outcomes, TPP and actual participation are positively correlated. In addition, in the case of meeting attendance, the correlation is quite strong and significant despite the sample size of 15 villages. The correlation between TPP and village project contributions is not nearly as significant, but the point estimate is positive, which provides another important “real-world” data point in support of theories of the relationship between network structure and social influence dynamics.

Figure 4



5 Alternative Explanations

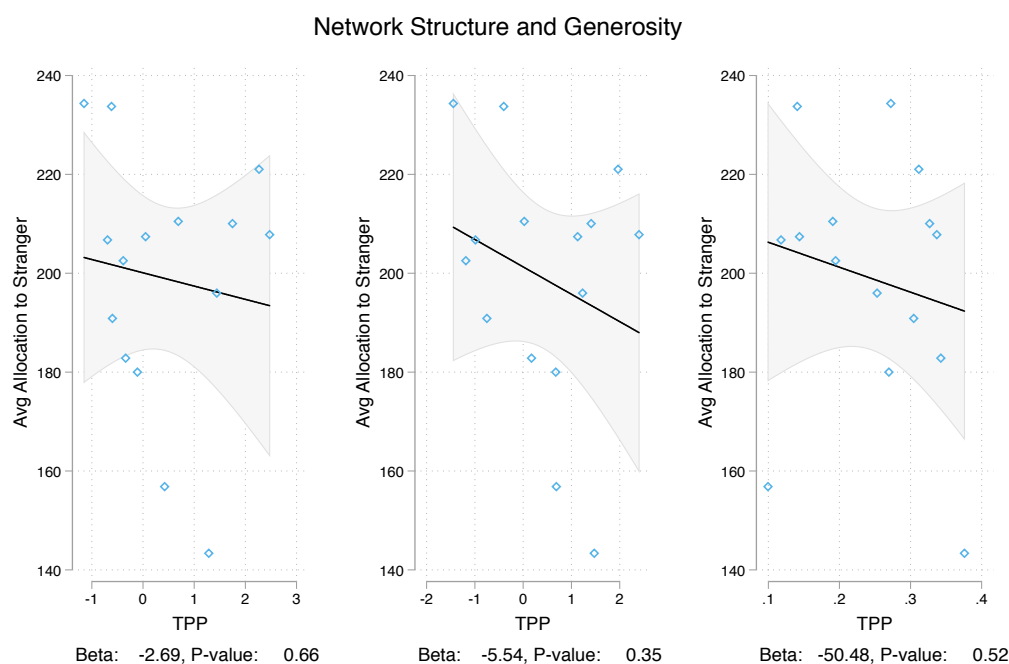
This section explores some alternative explanations for the observed correlation between equilibrium simulated participation and voter turnout. Specifically, we test for the possibility that both networks and turnout are shaped by a third factor like pro-social norms, and that differences in network structure are driven by the presence of “brokers” who may drive differences in mobilization. In addition, we test the possibility that differences are driven by the capacity of networks to diffuse information about the existence of an election (rather than information about participation). In all cases, we present the results of the “best available” tests, and conclude that theories of social context on a network remain the best fit for the available data.

5.1 Differences in Pro-Social Norms

The most common concern raised in any network study with observational data is the possibility that a third factor – for example, pro-social norms – are driving higher turnout and also the emergence of denser networks that, in simulations, appear to better spread influence.

We can test for this directly using the results of behavioral games conducted as part of a “lab-in-the-field” component of the survey from which this network data is drawn. In particular, we can test whether villages with higher equilibrium simulated

Figure 5



participation rates are also villages in which participants were more generous during the divide-the-dollar dictator game.³ If pro-sociality is driving both networks structure and turnout, generosity in the divide-the-dollar game should be positively correlated with equilibrium participation rates.

Figure 5 plots the correlations between average generosity of subjects in the divide-the-dollar game for each village and TPP and, in the third plot, actual turnout. The figure shows clearly that if anything, there is a *negative* correlation between pro-sociality among lab subjects and TPP.

In discussing threats to endogeneity, it is also worth noting an aspect of this analysis that departs from traditional network studies. In most network studies, researchers examine the correlation between behavior of an individual (the “ego”) and the behavior of the ego’s peers to determine the impact of peers on ego behavior. This inevitably raises the question of whether peer behavior is shaping ego behavior, or whether egos peers are peers precisely because they tend to be similar and thus behave similarly (McPherson et al., 2001). In studying the relationship between network structure as a whole and aggregate outcomes, however, we circumvent this traditional limitation of network studies; at no point does the *actual* behavior of individuals in our network or the *actual* behavior of their peers enter into the net-

³The divide-the-dollar game was organized as follows: first, subjects were given ten 100UGX coins. Subjects were then advised that they could split these coins between themselves and a stranger, who they were told will be “someone from Arua whom you do not know personally. We chose the stranger by randomly selecting someone living in Arua district from a long list.”

work measures used herein. In fact, we do not have individual-level information on turnout. Rather, our analysis focuses only on how the structure of a network shapes behavior on average across a wide range of average initial individual proclivities to participate politically, avoiding the concern of homophily entirely.

5.2 Role of Central Actors

Past research has often suggested an important role for “brokers” in political networks – individuals who occupy a central position in the network and are able to drive political behavior by virtue of their network position. Rojo et al. (2014), for example, suggest that slums in India with central brokers are more able to commit to turning out voters and thus get better public goods, and Cruz et al. (2017) finds a relationship between turnout and highly-central families. It is thus reasonable to ask whether the differences between networks found in this analysis may be strongly influenced by these types of central actors.

To test for this possibility, Figure 6 below plots the correlation between turnout and equilibrium simulated participation *after* dropping the 5 people with highest eigenvector centrality from each network. If results are indeed driven by “brokers,” then after dropping the most well-connected individuals, no relationship should exist between network structure and turnout. As the figures show, however, this is not the case. This shows that results are driven by general network connectedness and not a small number of facilitators.

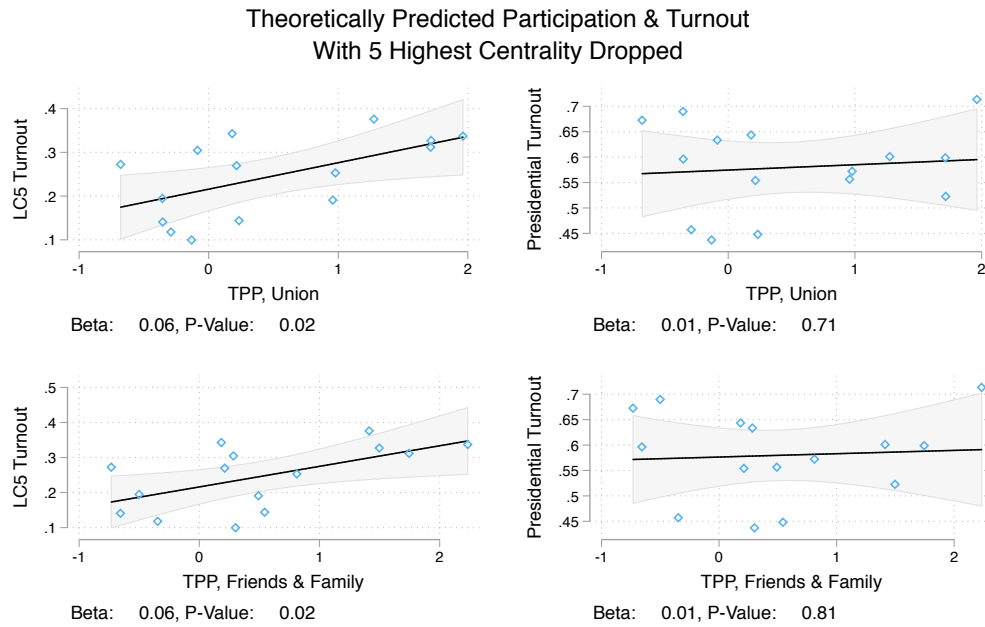
5.3 Information Diffusion

An additional concern is that networks that give rise to higher TPP may also be networks that better support the efficient diffusion of information, leading to greater turnout. We address this in two ways.

First, we correlate awareness of the UBridge program with turnout. UBridge was a novel program introduced by USAID to a number of individuals within each village. If the efficiency by which networks diffuse information is driving turnout, then UBridge awareness and turnout should be positively correlated. As shown in Figure 7, they are not.


Second, we run diffusion simulations on village networks to test which networks we think might better support the efficient spread of information, then correlate these theoretically-derived values with actual turnout. Figure 8 below plots actual turnout against an index of simulated information diffusion efficiency (details of the information diffusion simulation and index creation can be found in Appendix F). As before, we find no support for the information channel for explaining network influence on turnout.

Figure 6



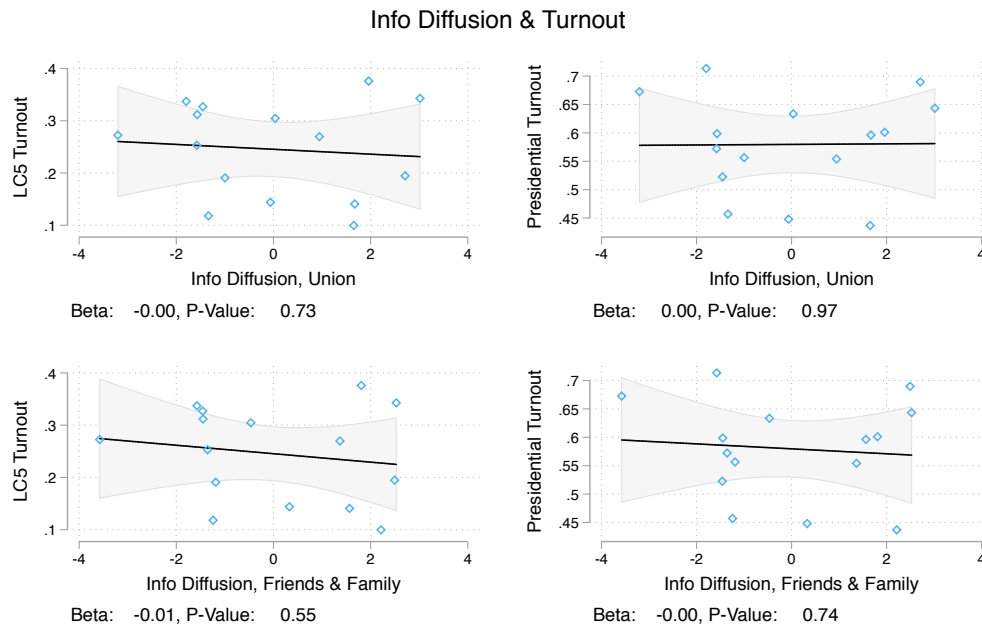
TPP in standard deviations of PCA index across parameter values, turnout in shares.

Figure 7: UBridge Awareness and Turnout



Awareness is share of village residents aware of UBridge, turnout in shares.

Figure 8: Simulated Information Diffusion and Turnout



Info Diffusion in standard deviations of PCA index across parameter values, turnout in shares.

6 Conclusions

This work provides, to the best of our knowledge, the first direct, empirical test of the theoretical predictions of social context theory that goes beyond simulations. It finds that at least among the 15 Ugandan villages examined herein, the theoretical predictions of Siegel (2009) and Rolfe (2012) on how network structure may impact political participation are borne out.

In addition to providing the first empirical test of these theories, our exercise offers several important lessons for network scholars. Most importantly, it shows that the salience of networks may be contingent on the environment being studied. When individuals are exposed to extensive messaging by extra-network mediums, the influence of network dynamics may be diminished. This finding has important implications not only for the generalizability of network studies in general, but also may provide a fruitful avenue for future research into how increasing dependence on network-based social media for all forms of political information may influence mass political behavior.

References

- Abrams, S., Iversen, T., and Soskice, D. (2011). Informal Social Networks and Rational Voting. *British Journal of Political Science*, 41(2):229–257–257.
- Centola, D. (2015). The social origins of networks and diffusion. *American Journal of Sociology*, 120(5):1295–1338.
- Centola, D. and Macy, M. (2007). Complex contagions and the weakness of long ties. *American journal of Sociology*, 113(3):702–734.
- Cialdini, R. B. (2015). Basic Social Influence Is Underestimated. *Psychological Inquiry*, 16(4):158–161.
- Cruz, C., Labonne, J., and Querubin, P. (2017). Village Social Network Structures and Electoral Competition. *Working Paper*.
- Ferrali, R., Grossman, G., Platas, M., and Rodden, J. (2017). Social Networks and Community Reporting. *Working Paper*.
- Fowler, J. H. (2005). Turnout in a Small World. In *Social Logic of Politics*, pages 269–287. Temple University Press, Philadelphia.
- Goyal, S. (2012). *Connections: an introduction to the economics of networks*. Princeton University Press.
- Huckfeldt, R. and Sprague, J. (1995). *Citizens, Politics and Social Communication: Information and Influence in an Election Campaign*. Cambridge University Press.
- Ioannides, Y. M. (2013). *From neighborhoods to nations: The economics of social interactions*. Princeton University Press.
- Jackson, M. O. and Yariv, L. (2010). Diffusion, Strategic Interaction, and Social Structure. In Benhabib, J., Bisin, A., and Jackson, M. O., editors, *Handbook of Social Economics*. Elsevier.
- Kenny, C. B. (1992). Political Participation and Effects from the Social Environment. *American Journal of Political Science*, 36(1):259.
- Knoke, D. and Yang, S. (2008). *Social Network Analysis*. Sage Publications, Los Angeles.
- Larson, J. M. and Lewis, J. I. (2017). Ethnic Networks. *Forthcoming in American Journal of Political Science*.
- Larson, J. M., Nagler, J., Ronen, J., and Tucker, J. A. (2017). Social Networks and Protest Participation: Evidence from 130 Million Twitter Users . *Working Paper*.
- Leighley, J. E. (1990). Social Interaction and Contextual Influences On Political Participation. *American Politics Quarterly*, 18(4):459–475.

- Mcclurg, S. D. (2003). Social Networks and Political Participation: The Role of Social Interaction in Explaining Political Participation. *Political Research Quarterly*, 56(4):449–464.
- McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444.
- Newman, M. *17 Epidemics on networks*. Networks: An Introduction. Oxford University Press.
- Nickerson, D. W. (2008). Is Voting Contagious? Evidence from Two Field Experiments. *American Political Science Review*, 102(1):49–57–57.
- Rojo, G., Jha, S., and Wibbels, E. (2014). Political Networks, Clientelism and Public Goods: Evidence from Slums in Udaipur, India.
- Rolfe, M. (2012). *Voter Turnout: A Social Theory of Political Participation*. Cambridge University Press, Cambridge.
- Ross, L. and Nisbett, R. E. (2011). *The person and the situation: Perspectives of social psychology*. Pinter & Martin Publishers.
- Siegel, D. A. (2009). Social Networks and Collective Action. *American Journal of Political Science*, 53(1):122–138.
- Sinclair, B. (2012). *The Social Citizen: Peer Networks and Political Behavior*. The University of Chicago Press, Chicago.

Appendix

A Turnout Estimation

As previously noted, this analysis assumes that votes at a precinct (polling place) are evenly distributed among voters registered at that precinct from different villages.

One constraint of this measure is that the accuracy of these estimates will be related to the correspondence of villages to precincts. If each village sends all residents to its own polling place (i.e. that polling place is only attended by residents of one village), inferences about village voting will be perfect. If, by contrast, all villages send their voters to a single polling place, effectively no information can be learned about how individual villages voted.

This correspondence can be summarized using a *concentration* statistic, where higher values mean the mapping from precincts to villages is more precise. For a village $v \in V$ and a polling place (precinct) $p \in P$, let $voters_v$ be the set of voters who live in v , let $voters_p$ be the set of voters who vote at polling place p , and let $voters_{v,p}$ be the set of voters who reside in v and voted at p . Then:

$$concentration_{v,p} \equiv \frac{\#voters_{v,p}}{\#voters_p} \quad (1)$$


In other words, $concentration_{v,p}$ is the share of voters at a precinct from village v .

Since each village sends people to multiple polling places, however, we must then calculate a weighted average of $concentration_{v,p}$ across polling places where the weight for each polling place is the share of voters from each village going to that polling place. Formally:

$$concentration_v \equiv \sum_{p \in P} concentration_{v,p} * \frac{\#voters_{v,p}}{\#voters_v} \quad (2)$$

The distribution of this statistic are presented in Figure 9:

Figure 9



B TPP Simulation Notes

B.1 Participation Updating

Updating of the $participation_{v,t+1}$ is accomplished by iterating through all vertices in the network in random order and having each vertex update its value of lpr and its participation $participation_{v,t+1}$ *sequentially* rather than simultaneously. This is the one departure from Siegel (2009). When all vertices update lpr simultaneously, it is possible to converge to a “flashing” state in which at time t a portion of the network is planning to vote while another portion is not planning to vote, at time $t + 1$, these two groups flip inclinations, and at time $t + 2$ they return to their initial state. This is caused by knife-edge simultaneity of updating, which seems unrealistic, since real updating is almost certainly sequential. Thus simulation uses sequential updating.

B.2 Convergence

The simulation is run until no more than 1% percent of the vertices in the network change participation status for at least 20 consecutive periods. Results below are averaged across 1000 runs for each set of parameter values.

C Social Context Simulation Validity

Table 4: Correlations across Parameter Values, Friends and Family Network

Variables	Mean 0.5, SD 0.5	Mean 0.6, SD 0.5	Mean 0.6, SD 0.25	Mean 0.7, SD 0.5	Mean 0.7, SD 0.25
Mean 0.5, SD 0.5	1.00				
Mean 0.6, SD 0.5	0.80	1.00			
Mean 0.6, SD 0.25	0.58	0.78	1.00		
Mean 0.7, SD 0.5	0.78	0.95	0.74	1.00	
Mean 0.7, SD 0.25	0.33	0.49	0.52	0.58	1.00

Table 5: Correlations across Parameter Values, Union Network

Variables	Mean 0.5, SD 0.5	Mean 0.6, SD 0.5	Mean 0.6, SD 0.25	Mean 0.7, SD 0.5	Mean 0.7, SD 0.25
Mean 0.5, SD 0.5	1.00				
Mean 0.6, SD 0.5	0.85	1.00			
Mean 0.6, SD 0.25	0.77	0.82	1.00		
Mean 0.7, SD 0.5	0.85	0.91	0.69	1.00	
Mean 0.7, SD 0.25	0.66	0.60	0.74	0.47	1.00

D Robustness Regression Tables

Table 6: Robustness for LC5 Election

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Basic	Basic	ELF	ELF	Educ	Educ	Concentrate	Concentrate	W/Oviva	W/Oviva
TPP (FAF)	0.0314 (1.64)		0.0310 (1.58)		0.0321 (1.59)		0.0407 (2.53)		0.0137 (1.27)	
TPP (U)		0.0355 (1.83)		0.0363 (1.84)		0.0355 (1.76)		0.0359 (1.82)		0.0143 (1.33)
ELF			-0.0867 (-0.62)	-0.108 (-0.80)						
Educ					0.0426 (0.21)	-0.0000690 (-0.00)				
Observations	15	15	15	15	15	15	8	8	16	16

t statistics in parentheses

Table 7: Robustness for Pres Election

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Basic	Basic	ELF	ELF	Educ	Educ	Concentrate	Concentrate	W/Oviva	W/Oviva
TPP (FAF)	-0.00722 (-0.36)		-0.00814 (-0.42)		-0.00815 (-0.39)		0.00319 (0.14)		-0.0297 (-2.58)	
TPP (U)		0.00437 (0.21)		0.00565 (0.28)		0.00423 (0.20)		0.00311 (0.13)		-0.0265 (-2.20)
ELF			-0.172 (-1.25)	-0.172 (-1.24)						
Educ					-0.0562 (-0.27)	-0.0410 (-0.20)				
Observations	15	15	15	15	15	15	8	8	16	16

t statistics in parentheses

E Considering Non-Recipical Ties

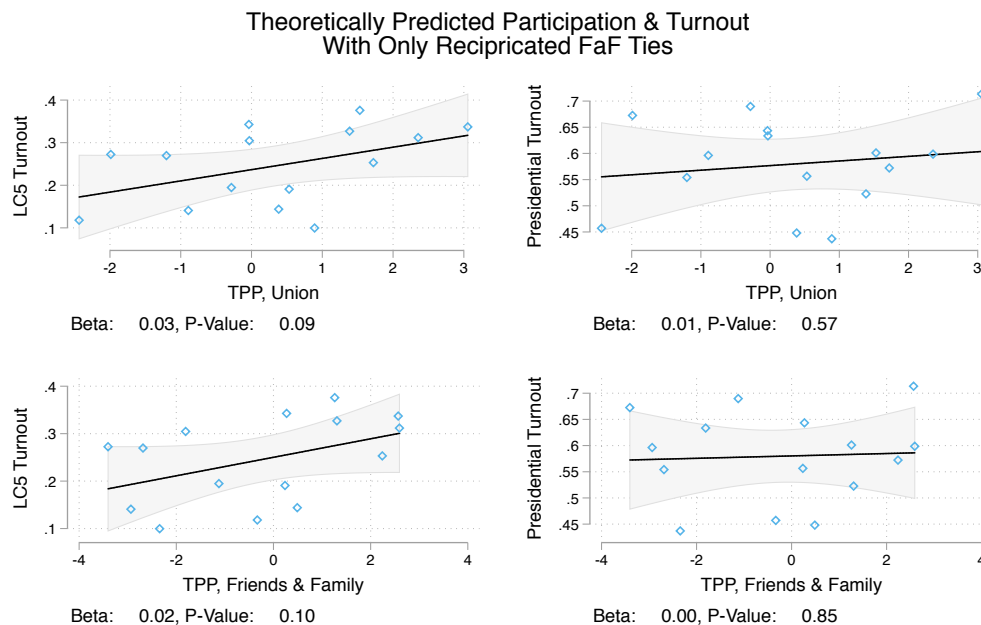
Figure 10 presents results when networks are created using only reciprocated ties to form Friends and Family ties (note the other two inputs into the Union network – the Lender / Solver networks – cannot be restricted in an analogous manner). The figures show results are, if anything, stronger under this restriction.

As shown in Table 8, however, it is not clear that these restrictions are reasonable given the low average degree they generate. This may be due censoring caused by the limited number of people individuals are allowed to list (5 family members and 5 friends), or failures to recall individuals.

Table 8: Network Summary Statistics: Including Only Reciprocated Friends and Family Ties

	Friends and Family	Union
Average Size	210.3	210.3
Average Num Connections	189.5	924.1
Average Degree	0.9	4.3
Min Size	160.0	160.0
Max Size	283.0	283.0

Figure 10



TPP in standard deviations of PCA index across parameter values, turnout in shares.

F Information Diffusion Simulation

There is no simple statistic – like average degree or average shortest path length – which reliably summarizes the ability of a network to efficiently diffusion information in this type of stochastic manner. (Newman, p. 19-35) Instead, the ability of a network to support information diffusion is estimated by simulating the diffusion process described above on empirical village networks and the examining the average speed with which information spreads for each village.

More precisely, the simulation proceeds as follows:

1. At $t = 0$, one vertex v_0 in the network (selected with uniform probability) is endowed with a unique piece of knowledge. It is thus “informed” ($I(v_0) = 1$). All other vertices are assumed to be ignorant of this knowledge ($I(v_j) = 0 \forall j \in V \setminus \{0\}$).
2. At $t = 1$, information spreads from v_0 to each of the neighbors of v_0 , denoted $N(v_0)$ with i.i.d. probability $\frac{p}{|N(v_0)|} \in (0, 1)$.
3. Step 2 is then repeated indefinitely, where at each stage all “informed” vertices spread their knowledge to neighbors with i.i.d. probability p .

The ability of the network to support diffusion can then be specified as the number of people in the network that have become “informed” after s steps of the diffusion model. The larger the number of people “informed” for a given number of steps s , the more efficient a village’s network.

Note that the probability of information diffusion from a vertex to her neighbors is normalized by the number of neighbors. This can be thought of as approximating the idea that individuals can only have so many interactions in a given period of time. This normalization more closely approximates the idea that all individuals have the same probability of interacting and sharing information with at least a friend in a given period, a dynamic suggested by recent work on information diffusion elsewhere in Uganda (Larson and Lewis, 2017). With that said, results look similar without the normalization.

F.1 Information Diffusion Summary Statistics

Table 9 below shows the correlation in the share of individuals in each village informed at different step thresholds, with different spread probabilities, and with different network specifications. As the table shows, inter-parameter correlations are quite high, and so an index is created for expositional ease by taking the first component of a PCA index for each network specification. It is this index that is used in Figure 8.

Table 9: Diffusion Correlations across Parameter Values

Variables	p 0.60, 10 steps, U	p 0.60, 20 steps, U	p 0.35, 10 steps, U	p 0.35, 20 steps, U	p 0.60, 10 steps, FaF	p 0.60, 20 steps, FaF
p 0.60, 10 steps, U	1.00					
p 0.60, 20 steps, U	0.75	1.00				
p 0.35, 10 steps, U	0.98	0.66	1.00			
p 0.35, 20 steps, U	0.96	0.84	0.94	1.00		
p 0.60, 10 steps, FaF	0.97	0.76	0.96	0.97	1.00	
p 0.60, 20 steps, FaF	0.70	0.95	0.62	0.83	0.77	1.00
p 0.35, 10 steps, FaF	0.98	0.66	0.99	0.92	0.96	0.61
p 0.35, 20 steps, FaF	0.95	0.85	0.91	0.98	0.97	0.86