

Title:

Channels for Influence or Maps of Behavior? A Field Experiment on Social Networks and Cooperation

Short title runner:

Channels for Influence or Maps of Behavior?

Keywords:

Social networks; local public goods; collective action; cooperation

Authors:

1. Paul Atwell

Phd Candidate

Department of Political Science

5700 Haven Hall

505 S. State Street

Ann Arbor, MI 48109

patwell@umich.edu

2. Noah L. Nathan

Assistant Professor

Department of Political Science

5700 Haven Hall

505 S. State Street

Ann Arbor, MI 48109

nl Nathan@umich.edu

Acknowledgements footnote:

We thank Jon Atwell, Alex Coppock, Mai Hassan, Allen Hicken, Walter Mebane, Michael Weaver, the anonymous reviewers, and audience members at MPSA 2019, the University of Michigan, and the 2019 Sociology of Development Conference at Notre Dame for comments and suggestions. We also thank Paul Osei-Kuffour, Mildred Adzraku, and the staff of CDD-Ghana for their generous support, as well as Anthony Bilandam, Sulley Daniel, Mathew Dome, Yahaya Halik, Paul Larri, Mohammed Osman, Rose Sandow, Mohammed Seidu, and Nagob Solomon for their research assistance. Funding was provided by the Center for Political Studies and the International Policy Center at the University of Michigan. This research was declared exempt from human subjects review by the University of Michigan Health and Behavioral Sciences Institutional Review Board (HUM00147575, 4 June 2018). Our study pre-registration can be found at: <http://egap.org/registration/4750>.

Channels for Influence or Maps of Behavior? A Field Experiment on Social Networks and Cooperation

Paul Atwell* and Noah L. Nathan†

Abstract

Communities in developing countries often must cooperate to self-provide or co-produce local public goods. Many expect that community social networks facilitate this cooperation, but few studies directly observe real-life networks in these settings. We collect detailed social network data in rural Northern Ghana to explore how social positions and proximity to community leaders predict donations to a local public good. We then implement a field experiment manipulating participants' opportunity to communicate and apply social pressure before donating. We find clear evidence that locations in community social networks predict cooperative behavior, but no evidence that communication improves coordination or cooperation, in contrast to common theoretical expectations and laboratory findings. Our results show that evolved, real-life social networks serve as a mapping of community members' already-engrained behaviors, not only as an active technology through which social influence propagates to solve collective action problems.

Replication Materials: The data, code, and any additional materials required to replicate all analyses in this article are available on the American Journal of Political Science Dataverse within the Harvard Dataverse Network, at: <https://doi.org/10.7910/DVN/O059HN>

Word count (including all notes and references): 9,996

*PhD Candidate, Department of Political Science and Ford School of Public Policy, University of Michigan. Email: patwell@umich.edu.

†Assistant Professor, Department of Political Science, University of Michigan. Email: nlnathan@umich.edu

Introduction

Communities in the developing world often complement limited state provision by self-providing or co-producing local public goods (Miguel and Gugerty 2005, Habyarimana et al. 2007, 2009, Baldwin 2015). Many mechanisms for overcoming the collective action problems inherent in self-provision are expected to operate through the social ties that connect community members to their peers. Social networks facilitate cooperation by allowing community members to communicate and coordinate their behaviors (Ostrom et al. 1992, Balliet 2010), as well as by providing avenues to observe and deter free-riding (Ostrom 1990, Habyarimana et al. 2007, 2009). Similarly, the ability of community leaders – such as traditional chiefs – to serve as focal points who induce cooperation is said to depend on the nature of their social ties to community members (Baldassarri and Grossman 2011, Baldwin 2015, Gottlieb 2017).

Links between social networks and cooperation have been documented in laboratory settings (e.g., Rand et al. 2011). But few studies examine how real-life social network structures mediate the self-provision of local public goods, especially in the developing world.¹ More importantly, key questions remain about the mechanisms through which social networks affect cooperation. Many political science theories assume that social relationships matter primarily because of the information that actively flows through ties to facilitate social influence (Ostrom et al. 1992, Habyarimana et al. 2007, 2009, Siegel 2009, Rolfe 2012, Larson and Lewis 2017). Yet social networks may also explain cooperative behavior even when no information passes among ties at all.

Studies of evolutionary network formation suggest that real-life network positions will often be outcomes of past behavior. Current networks may be endogenously sorted by individuals' innate types or traits (e.g., cooperative or non-cooperative) or existing internalization of cooperative norms (Eguiluz et al. 2005, Wang et al. 2012, Gallo and Yan 2015). Moreover, a long tradition of

¹More common are proxy measures of social ties rather than full networks (Habyarimana et al. 2007, 2009, Dionne 2015). Research mapping networks in developing countries primarily focuses on other outcomes, such as technology adoption, information diffusion, or political behavior (Alatas et al. 2016, Cruz et al. 2017, Larson and Lewis 2017, Eubank et al. 2019, Ferrali et al. 2019).

scholarship in sociology posits that individuals act based on internalized norms of behavior tied to their pre-existing positions within social structures (White et al. 1976, Burt 1982). Each of these theories suggests that social network positions may be correlated with behavior in a new situation even absent any contemporaneous node-to-node influence (Shalizi and Thomas 2011).

Observing this second category of explanations – which views a network as a map of likely behaviors rather than a technology for active information transmission and social influence – requires studying cooperation in the field. Unlike in laboratory games among strangers, real social networks shaping local public goods provision in rural communities are evolutionary, accumulated outcomes of lifetimes of repeated interaction that are likely already sorted by past behavior. Moreover, theories of networks as technologies for information transmission and networks as maps make observationally equivalent predictions in many applications. For example, if scholars observe that more network central – or findable (e.g., Habyarimana et al. 2007, 2009) – individuals cooperate more, they often will have no means to distinguish whether this is due to greater information transmission into and out of more central nodes, or to underlying differences in traits or internalized norms already correlated with network centrality. But these two sets of theories have very different implications for understanding local public goods provision: to the extent that the assortative, evolutionary selection of network ties confounds estimates of the effect of network position on behavior, we may significantly over-attribute real-world cooperation to active social influence within communities.

We implement an original field experiment that distinguishes between these two categories of explanations by randomly varying the opportunity for information to flow through nodes in real-life community social networks as community members coordinate in the provision of a real local public good. After mapping household-level social networks in five ethnically homogenous communities in Northern Ghana, household heads in the control group could donate some or all of a fixed endowment to a community leader to fund a local public good, with contributions doubled by the research team. The leader and good were selected by plurality vote among the donors.

Participants in the control group made anonymous donations immediately after this activity was introduced to them, with no opportunity for information or influence to flow between them and peers or community leaders.

For treatment households, enumerators instead waited for one to two days after the baseline survey to collect donations. This created an opportunity for peers to communicate, as well as for community leaders to steer participants towards a preferred donation strategy. We observe whether communication affects intra-community and dyadic *coordination* – the degree of clustering on the same behaviors – and *cooperation* – the amount donated. We are especially interested in how the locations of participants in the social network – relative to each other and to community leaders – moderate these effects.

Overall, many participants cooperated (donated) and there is clear evidence that social network positions predict behavior. More network central participants donated significantly more to the local public good. “Conditional cooperators” – those who donate more when they expect others to also donate more (Ostrom 2000) – were also heavily concentrated at the center of each network, but absent on the periphery, where free-riding was more likely. In addition, participants who were more closely linked were more likely to vote for the same projects and select the same leaders to manage them. Crucially, each of these relationships exists within the overall sample and the control group alone, suggesting that social positions explain cooperative behavior even absent active information or influence propagating over the network.

But we find no support for any of our hypotheses about the “delayed collection” treatment.² Treated participants did not coordinate – converge on donation amounts – during the waiting period, even with their closest social alters. Donation amounts and leader and project preferences also did not differ systematically by participants’ social proximity to community leaders. Despite no change in coordination (clustering), cooperation (donation amounts) actually declined slightly among participants who could communicate, not increased, as existing studies predict (Ostrom et

²Our hypotheses and pre-analysis plan were pre-registered prior to implementation.

al. 1992, Bochet et al. 2006, Vanberg 2008, Balliet 2010).

These null results cannot be explained by treated participants consuming their endowments, being influenced by social ties in the control group, or by the stakes being too low. Instead, significant communication occurred, with most treated participants speaking with others about the activity and many attempting to convince each other to donate more. But this social influence mostly failed. Meanwhile, with the exception of attitudes about which project to select (e.g., water vs. electricity), network positions still predict similar behaviors among treated participants as in the control group, suggesting that cooperative behavior may already be encoded in network positions.

External validity is necessarily constrained by focusing on a particular collective action situation in a small set of communities. In addition, by keeping all donations anonymous, our design forecloses on monitoring and sanctioning for free-riding – an additional mechanism by which networks are hypothesized to actively influence behavior. We keep donations anonymous for ethical reasons, but also to ensure that behavior in the control group could not be shaped by anticipated social pressure: control participants knew free-riding would go unpunished, yet cooperation was correlated with network positions all the same. Because laboratory studies regularly find that communication alone increases cooperation even absent opportunities to monitor or punish free-riding (Ostrom et al. 1992, Bochet et al. 2006), finding that it fails to do so in the field still suggests a need to broaden standard expectations about how networks facilitate cooperation.

This paper makes several contributions. First, we demonstrate that scholars of cooperative local public goods provision may be focusing too much on active and direct pressure from peers and leaders at the expense of studying the ways in which cooperative behaviors may already have become embedded in local social structures. The two broad views of social networks we compare are not mutually exclusive; our single study by no means implies that information does not flow through networks in other contexts in ways that improve cooperation, especially where sanctioning is possible. But showing that networks explain cooperative behavior even *without* any information or influence being transmitted challenges scholars to think more carefully about the possible mech-

anisms linking social ties to real-life cooperation.

Second, we join a small set of studies in political science demonstrating that real-life locations in social networks predict political behavior in the developing world (Cruz et al. 2017, Larson and Lewis 2017, Eubank et al. 2019, Ferrali et al. 2019). We advance this literature by being among the first, alongside Eubank et al. (2019), to use networks to explore how communities solve real collective action problems.

Third, null findings for our experimental hypotheses about the influence of traditional chiefs and peer-to-peer interactions complicate our understanding of why collective action problems appear simplest to resolve in small, ethnically homogenous communities where community members have strong social ties (Olson 1965, Miguel and Gugerty 2005, Habyarimana et al. 2007, 2009). Even in what approximates a “most-likely” setting for social influence and leadership in encouraging cooperation – key mechanisms used to explain cooperation in these settings – we find no evidence that peers or leaders influenced contributions. Instead, the already high degree of cooperation in our control group suggests a need to look elsewhere – such as to processes of evolutionary network formation or the imprinting of social roles – to better understand how small communities solve collective action problems.

Two views of social networks

In contexts of low state capacity, communities often must overcome collective action dilemmas to co-produce or self-provide some local public goods. Many theories assume that social ties play an important role in preventing free-riding. In this first view of social networks, they are a technology through which contemporaneous information spreads to facilitate social influence (Siegel 2009, Larson and Lewis 2017).

A key means by which networks serve as a technology for influence is by facilitating communication among nodes. Despite being “cheap talk,” communication regularly reduces free-riding in the laboratory – without a threat of *ex post* sanctioning – by allowing peers to make promises

of cooperation they find hard to break (Ostrom et al. 1992, Ostrom 2000, Bochet et al. 2006, Balliet 2010). Communication causes players to feel empathy, sparking altruism (Andreoni and Rao 2011), and many players prefer not to break cooperative promises due to feelings of “guilt aversion” (Charness and Dufwenberg 2006, Vanberg 2008). Similarly, respected community leaders, such as chiefs, are claimed to facilitate cooperation by communicating which behaviors are in their community’s best interest (Baldwin 2015, Gottlieb 2017). Closer alters in the network are more likely to communicate or be able to do so effectively (Habyarimana et al. 2007, Larsen and Lewis 2017). Network locations – such as centrality or proximity to peers or leaders – should then affect how much social influence individuals are subject to, and in turn, how much they cooperate and coordinate.

The broader literature linking social networks to other types of political behavior similarly views networks as a technology through which actors learn contemporaneous information about their peers’ (intended) behavior and adjust their own (Fowler 2005, Sinclair 2012, Rolfe 2012, Eubank et al. 2019). Importantly, just as in laboratory experiments on cooperation, social influence through networks is commonly argued to occur even without direct sanctioning by peers (Rolfe 2012). Communication about each other’s behaviors can produce contagion, learning, updated priors about a behavior’s risks, preferences for conformity, or related mechanisms of social influence (Granovetter 1978, Fowler 2005, Siegel 2009).

Yet there is a second possible relationship between networks and cooperation, emerging from observations that cooperation does not *require* contemporaneous exchanges of information. In laboratory public goods experiments, full free-riding is rare, even if behavior is anonymous and players cannot communicate (Ostrom 2000). Instead, many players behave as “conditional cooperators,” cooperating as long as they expect others will do so as well (Fischbacher and Gächter 2010). Network positions – such as centrality or the distance between pairs of nodes – may *already* map to these types of behaviors. This could be either (a) due to the endogenous selection of network ties or (b) because pre-existing locations in social structures have already caused participants to

adopt different behavioral rules.

First, individuals in an evolved, real-life social network may be sorted by their innate or learned propensity to cooperate, with current network locations selected on prior behavior. Laboratory and simulation studies of network formation demonstrate that, through repeated interaction, cooperative players are often sorted into the centers of networks because they develop and keep more social ties (Gallo and Yan 2015). Similarly, individuals who act as conditional cooperators often cluster in central social positions because others are more likely to maintain social relationships with them (Wang et al. 2012). Sorting can also affect social proximity among pairs, or dyads, of nodes: homophily – the assortative clustering of individuals with similar characteristics – is an extremely common feature of networks (McPherson et al. 2001).

Endogenous sorting is possible regardless of whether we view individuals' underlying propensity to cooperate as emerging from their fixed "types" – for example, from innate personality traits that affect cooperative dispositions (Volk et al. 2012, Balliet and Van Lange 2013) – or from pre-existing differences in internalized cooperative norms and behavioral rules.³ Either way, the process through which ties are selected is a major confounder for social influence: proximate ties in the network may be linked because they usually behave similarly, not behave similarly because they are linked (Shalizi and Thomas 2011).

Second, individuals' positions within evolved local social structures may have already shaped their beliefs about the appropriate social roles and behavioral rules that apply to them. Sociologists of "role theory" have long posited that different social positions and their normative construction cause individuals to internalize different behavioral norms (White et al. 1976, Burt 1982). For example, socially central individuals may see themselves as prominent community members with responsibility to ensure their community is provided for, cooperating regardless of whether they communicate with others or are pressured to do so. Such a dynamic has similar implications as

³There is broad variation in whether an acquired or innate behavioral rule such as conditional cooperation is labeled a "norm" (Elster 1989, Bicchieri 2017) or a player's "type" or "trait" (Ostrom 2000, Rolfe 2012).

sorting: existing network locations will *already* be correlated with likely behavior. Under either mechanism, cooperation in a new situation may be highly correlated with network positions even without further social influence or information transmission.

In what follows, we are agnostic about whether endogenous sorting or the internalization of social roles is most responsible for mapping network positions to behavior. But we attempt to distinguish this second broad view of social networks from more standard expectations that networks matter for cooperation instead because they facilitate the *active, contemporaneous* transmission of information and influence. We do so by experimentally manipulating opportunities for communication and influence over already evolved, real-world networks with real local public goods at stake.

The design below does not speak to the effects of sanctioning, however. Sanctioning is an additional mechanism falling under the first view of networks as a technology for influence: via monitoring, information can also pass among nodes about who is free-riding and be used to punish defectors (Ostrom 1990, Miguel and Gugerty 2005, Habyarimana et al. 2007). Showing that networks already map to cooperative behavior suggests a plausible alternative explanation for common claims about sanctioning, such as that more central, or findable, individuals cooperate more because they are easier to sanction. But our design cannot rule out that networks could still be technologies for influence in situations in which sanctioning is possible.

The case: rural Northern Ghana

We focus on Northern Ghana as an illustrative case of the collective action dilemmas faced by rural communities in which the state is weak. Communities in this peripheral region frequently must bridge gaps in state service provision, at least in part, through the co-production of basic local public goods. For example, faced with limited state funding for school construction, it has been common for communities to pool resources and labor to put up simple school structures on their

own initiative, with the hope their informal school is eventually absorbed by the state system.⁴ In another example, the financially-constrained state electrical utility has a long-standing cost-sharing program that prioritizes connecting rural communities to the grid if they first pool their own resources to buy the electrical poles and other supplies (Ministry of Energy 2010). Related forms of self-provision or co-production have been studied in Kenya (Miguel and Gugerty 2005), Malawi (Dionne 2015), Uganda (Habyarimana et al. 2009), and Zambia (Baldwin 2015), among other cases.

While Northern Ghana is home to 30 distinct ethnic groups, most rural communities are homogenous. As a result, Northern Ghana provides archetypical examples of the small, ethnically homogenous communities in which social relationships among peers are thought to most easily facilitate cooperation (Miguel and Gugerty 2005, Habyarimana et al. 2007, 2009). In addition to ties among community members, we also focus on links to three main community leaders: chiefs, elected District Assembly members (non-partisan local government representatives), and religious leaders (pastors or imams, depending on ethnic group). In practice, chiefs are often most important; as elsewhere across much of rural Africa, they play a particularly central role in coordinating self-provision and co-production (Baldwin 2015), although Northern ethnic groups vary in whether their traditional leaders are legally recognized.

Data and design

Sample selection

Our initial sample included six rural communities, chosen in matched pairs by ethnic group. This includes two Mamprusi communities from the West Mamprusi District of the North East Region, two Builsa communities in the Builsa District of the Upper East Region, and two Konkomba communities in the Kpandai District of the Northern Region. All are ethnically homogenous. These

⁴These are informally called “wing schools” in Ghana. Interview with community elders in Mion District, 16 May 2019; interviews with community elders in Nanumba North District, 4-5 July 2019.

communities were chosen non-randomly to both (a) select typical rural villages in Northern Ghana and (b) to examine the most similar possible sets of communities – within and across ethnic groups – on basic demographic covariates. The community selection procedure is described in the Supporting Information (pg. SI.2). Because of enumerator errors, we were not able to compile network data for one Mamprusi community. It is dropped, leaving a final sample of five communities (pg. SI.3). On average, these communities had 725 residents across 89 households as of the 2010 census, with 95% of households engaged in farming. Demographic information is in the SI (pg. SI.3).

Data collection occurred across two survey waves in July 2018. In the first, we interviewed the household head (male or female, as applicable) of *all* households in each community. There were 431 households total, ranging from 73 in the smallest community (Builsa #2) to 97 in the largest (Konkomba #1). In the second wave, we attempted to re-interview the 229 (53%) respondents who had been randomly assigned to treatment (see below). Enumerators were instructed to wait at least one full day after finishing baseline interviews before returning to conduct the second wave. We successfully re-interviewed 215 (94%) treated respondents.

Network measurement

We use the first wave to map community social networks. We define each node as a household.⁵ Respondents were asked to complete a census of household residents, including full name, nicknames, gender, and age. Respondents then completed four prompts in Table 1 representing different domains of social relationships for which respondents could nominate five people each, derived from Eubank et al. (2019). The domains capture (1) ties to friends, (2) ties to individuals from which household heads seek help, (3) ties to individuals with whom household heads most frequently discuss politics, and (4) ties to family members (outside the immediate household).

⁵While interviewing every adult resident would produce even higher quality data (Larson and Lewis 2019), this would have increased the survey size by nearly 4.5x, requiring a prohibitively large budget for endowments in the game.

Respondents could repeat names across domains. Because we are focused on intra-community collective action, we restricted respondents to naming individuals *currently* living in the same community, but outside their household.

[Table 1 around here]

For our main analyses, we generate an undirected and “closed” network by matching names from each set of reported ties to names in the household rosters. Throughout the paper we use the union of the four network domains, though we also show that our results do not differ very substantially across types of ties (pg. SI.7). Relying on both automated and manual matching, we code two households as linked if either household head lists an adult member of the other household for any of the four questions in Table 1. If a household head names a person who is not listed on the roster of any other household, this link is dropped (i.e., a “closed” network). All links between households are given equal weight, although we also report specifications that instead weight ties by the sum of the four domains (pg. SI.8). Graphs for each of the five community networks are in Figure 1.

[Figure 1 about here]

Overall, we matched 62% of named alters to households in each community.⁶ For comparison, this is a significantly better match rate than Larson and Lewis (2017), who report in their appendix that only 15% of alters listed by their rural Ugandan respondents were matched, and is in the same range as Ferrali et al. (2019), who report an approximately 70% match rate, also in Uganda. Moreover, additional field work in October 2018 that attempted to locate the remaining unmatched alters in the Mamprusi community indicated that many were rural-urban or seasonal migrants living elsewhere in Ghana at the time of the experiment, or instead from nearby communities. Because we seek to measure the social network of community members *at the time of the study*,

⁶This ranges from 78% matched in the two Konkomba communities to only 42% matched in the Mamprusi community.

these unmatched alters fall outside the population of interest; it is appropriate to drop them from the network. However, at least some unmatched alters could still be residents not matched due to measurement error (e.g., inconsistencies in transliterating names).⁷ But we have no reason to believe that these matching errors are systematically worse than those in other recent studies employing similar data.⁸

Experimental design

Our experiment occurs in the context of a standard public goods game. Respondents in the first wave survey were randomly assigned with equal probability to treatment and control.⁹ Balance statistics are in the SI (pg. SI.12). At the conclusion of the first interview, control respondents were given 5 Ghana cedis (GHS), approximately \$1.25 USD, in 1 GHS notes. They were informed they could keep it or donate any integer increment towards a locally non-excludable public good to be used by the community, such as the purchase of school materials or electrical poles. Any donation was matched (doubled) by the research team.¹⁰ This interaction occurred in private. Participants were assured donations would be kept anonymous. The details of this activity were not previously explained prior to this survey interview.

⁷We are less concerned that we systematically missed entire households, however, as the total number of households interviewed in each community very closely match pre-existing census data. Moreover, the follow-up interviews did not indicate any households had been missed.

⁸In addition, we also introduce an alternative method of calculating networks that incorporates information lost when dropping unmatched alters. We create weighted networks in which direct links have double the weight of indirect links that pass through a shared unmatched alter. See pg. SI.10.

⁹We also separately randomly assigned an unrelated experimental condition that primed evaluations of land tenure security. We evaluate this elsewhere, but control for an indicator for the second treatment in all analyses here.

¹⁰The prompt was (in translation): “*We will be working with community members to fund a project that is in the interest of the community. This could be any local project that the community needs. We are soliciting donations to fund this project from community members like you. Anyone who donates will get to vote on who should be responsible for leading the project and spending the money that we raise. Once we finish interviewing everyone in this community we will present all of the funds that have been raised to the leader the community selects. We will also solicit feedback from everyone who makes a donation about what the money we raise should be spent on. We will give this information on your preferences to the leader that you collectively select... If you would like, you can keep the 5 Ghana Cedi we have just given you. But we wanted to give you this opportunity to donate some of this money to help fund the activity. If you make a donation, our team will add the exact same amount that you donate, up to 5 Ghana Cedi. This means we will double your donation with our own contribution to the community project and will do so with all contributions from the community.*”

We also elicit two accompanying preferences after a donation is made: (1) the type of project for which the participant preferred the money be used,¹¹ and (2) which leader the participant wanted to manage the accumulated funds among the chief/headman, pastor/imam of the largest congregation, and assemblymember. Participants were also asked about their beliefs as to how much other community members were likely to donate, and which leaders and projects they believed other community members would select. Participants were informed before donating that the plurality leader chosen by donating players would receive the funds with information on the project selected by a plurality of donors.¹² The research team presented the accumulated donations and information to these leaders in public events after the second wave survey.

For treated participants, enumerators still explained the activity, secured informed consent, and gave the 5 GHS gift at the end of the first interview. But treated participants were informed that the enumerator would return to collect any donations several days later.¹³ During this second interaction, a brief interview was conducted and donations and project/leader votes were collected. Treated participants were asked what had happened in the community in the interim and probed again about their expectations of others' behavior.

Ethics

We took several measures to ensure that the experiment was conducted as ethically as possible. We obtained informed consent from community leaders and individual participation was voluntary. At the end of the first interview, participants were asked if they would like to continue on to the public goods activity, which was explained to them, and told that they would still receive their gift regardless of whether they agreed. Only one opted out. We used 5 GHS as an amount that would fairly compensate for the interview time relative to local wages (the median interview was 20 minutes),

¹¹This question was open-ended, and subsequently coded into categories such as “education”, “water”, and “electricity.”

¹²Participants chose the traditional chief in four communities. The first Konkomba community selected the assemblyman.

¹³On average, treated respondents were re-interviewed 33.9 hours after their initial interview.

but also not coerce participation.¹⁴ It was not within our power to control what the selected leaders did with the funds after we left each community. By providing a choice over leaders, we instead allowed donors to select the leader they trusted most to use the funds appropriately. They could also choose not to donate at all if they did not trust how the money would be spent.¹⁵ Finally, all donations were collected in private to ensure that there could be no direct harm as a result of respondents' choices. This meant free-riding would not be observed or punished.

Pre-registered hypotheses

We pre-registered two sets of hypotheses. Our first three hypotheses (Table 2) focus on the ways in which social networks may predict underlying behavior in the public goods game. Either view of networks above should expect that more socially central individuals cooperate (donate) more (*H1*), that individuals more socially proximate to the leaders expected to manage the public good will cooperate more (*H2*), and that participants who are more proximate in the network behave more similarly (*H3*).

[Table 2 about here]

Second, we have three experimental hypotheses about the effects of communication (Table 3). Existing theory based on laboratory studies makes the strong prediction that communication should increase cooperation (donation amounts), even without monitoring and sanctioning (Ostrom et al. 1992, Sally 1995, Ostrom 2000, Bochet et al. 2006, Vanberg 2008, Balliet 2010). However, we initially pre-registered only the more modest hypothesis (*H4*) that communication would improve coordination (clustering in behavior). We test both hypotheses below: that treatment causes an increase in donations (higher amounts) or an increase in coordination (lower variance).

¹⁴By our best back-of-the-envelope calculation, this is one quarter to one third of a median daily wage in rural Northern Ghana.

¹⁵Participants were only told that donations would be turned over with information about donors' preferences.

H5 and *H6* are instead about interactions of treatment with pre-treatment network characteristics to test for the presence of influence traveling through the social network. This includes influence from participants on each other (*H6*), as well as the influence that community leaders may wield over participants to whom they are more closely tied (*H5*).

[Table 3 about here]

Non-experimental analyses

Before investigating the effects of treatment, we document overall behavior across the combined sample (*H1-H3c*) and then explore behavior in the control group only, where there were, by design, no opportunities for active information transmission or social influence.

Descriptive results in the full sample

In the first set of analyses, we find support for *H1*, *H3b*, and *H3c* in the full sample, with some social network features clearly correlated with behavior in the public goods activity. These results are in Table 4.

[Table 4 about here]

We test *H1* by estimating an OLS regression of donation amount on the eigenvector centrality of each participant, controlling for an indicator of treatment status, as well as community fixed effects and individual-level demographic controls.¹⁶ Eigenvector centrality captures a node's relative influence – the degree to which each household has links to other households who also have many links.¹⁷ In clear support of *H1*, participants who are more network central make significantly

¹⁶The fixed effects ensure we only make relative comparisons among individuals within the same network – eigenvector centrality's scale is network specific. The individual-level controls in all models are: age, household size, participant's number of children, whether the participant has a middle school education or greater, an assets index (pg. SI.13), whether the participant regularly travels outside the community, and works in a skilled trade (as opposed to farming). The results are also robust to foregoing controls (pg. SI.14).

¹⁷The SI (pg. SI.16) shows robustness to alternative measures of network centrality.

larger donations, while less central participants exhibit greater free-riding ($p < 0.001$). Substantively, participants at the 90th percentile of eigenvector centrality donated 0.46 GHS more than participants at the 10th percentile.

We instead find no support for *H2*, which we test by regressing donation amounts on participants' proximity to the community leader they expected to be selected. We measure network proximity in two ways. First, we use a count of "shared nearest neighbors" (SNN), which is the number of common households to which two nodes are directly linked. Second, we use "minimum path distance" (MPD), defined as the shortest whole number of steps on the graph between nodes.

¹⁸ We include the same controls and fixed effects as for *H1*, and also now control for eigenvector centrality.

For *H3*, the unit of analysis switches to all possible dyads of participants in each community. For *H3a*, the outcome is the absolute value of the difference between the amount donated by each participant in the pair, controlling for indicators for the joint treatment assignment of each pair, the average values of the demographic controls within the pair, and community fixed effects. We estimate separate OLS regressions with the two different measures of proximity: shared nearest neighbors or minimum path distance. For *H3b* and *H3c*, the outcomes are instead binary indicators for whether the pair chose the same leader or same project type, respectively.¹⁹ Because participants appear in multiple dyads, we cluster standard errors as in Aronow et al. (2015).

While we find no support for *H3a*, more socially proximate participants are significantly more likely to choose the same leader to manage the public good (*H3b*). Moreover, when proximity is defined as shared nearest neighbors, we find that more proximate participants are also significantly more likely to pick the same type of project (*H3c*). Overall, participants who are socially closer

¹⁸The chief was not interviewed in the first Builsa community, so it is dropped for *H2* (and *H5* below). In addition, if there is no path between a participant and the community leader (see Figure 1), we set the minimum path distance to one link greater than the furthest minimum path distance to the leader among others in the community. Leaving path distance as "NA" instead systematically drops the most distant participants.

¹⁹These latter two outcomes are only asked of participants who donate to the public good. Outcomes in dyads in which one or both respondents did not donate are coded as 0 (non-match) to avoid inappropriately conditioning on an intermediate outcome (Coppock 2018).

have more similar preferences for how they would like community contributions to local public goods to be managed and spent.

Behavior without communication

Crucially, we also show that these results hold among the *control group only*, when the opportunity for any information or social influence to propagate among peers was limited. Control participants played the public goods game immediately when it was first explained to them in the context of a private interview, after being explicitly informed their donation would be anonymous.²⁰ Table 5 repeats the estimates of *H1*, *H3b*, and *H3c* above, subsetting to the control group only. All three results from Table 4 are present even without opportunities for communication, and *H3c* now holds using either measure of proximity. More central control participants cooperate more, and more proximate participants pick the same leader and project preferences.

[Table 5 about here]

Additional non-pre-registered analyses suggest that behavior in the control group instead reflects existing propensities to cooperate. Most participants in laboratories initially partially cooperate in public goods games, donating approximately 40% to 60% of their endowments in the first (or only) round (Ostrom 2000, Fischbacher and Gächter 2010). This is exactly what we find: the mean donation in our control group was 2.13 GHS, 43% of the 5 GHS endowment; only one control participant kept it all. Ostrom (2000) explains that this occurs because many players begin by applying a behavioral rule of “conditional cooperation,” donating if they expect that others will do so. Each control group participant donated 0.32 additional cedis (GHS) for each additional cedi

²⁰The only way influence could still have propagated among control participants is if the first respondents in a community immediately and systematically sought out later respondents and explained the activity to them while the enumerators were still working from house to house. This is unlikely. Interview order among control participants is completely uncorrelated with both donation amounts and expectations of others' behavior, inconsistent with later respondents updating from interactions with earlier respondents (pg. SI.16).

she expected others to donate ($p < 0.001$, 95% CI: 0.14, 0.50).²¹

Moreover, conditional cooperators are heavily concentrated near the center of each community's social network, but absent at the social periphery.²² Among control participants, an interaction between eigenvector centrality and what each participant expects others to donate predicts donation amounts ($p < 0.01$, see pg. SI.18). Substantively, control participants with eigenvector centrality above the community median donate 0.75 GHS more for each additional GHS they expect others to donate ($p < 0.001$, 95% CI: 0.46, 1.05). But the donations of control participants with below-median centrality in each community are not correlated with expectations of what others will donate ($\beta = 0.12$, $p = 0.35$, 95% CI: -0.14, 0.37). These results appear consistent with the second view of networks above: that social networks may already be sorted by players' likely behaviors.

Experimental analyses: behavior with communication

Next, we examine our experimental hypotheses ($H4-H6c$) to test the effect of communication. We find null results for each of $H4-H6c$, and also show that communication did not increase cooperation (donation amounts).

Estimation approach

Our analysis is complicated by the fact that treated participants could – by design – influence each other. This violates the SUTVA assumption of no interference (Rosenbaum 2007). Fortunately, the most problematic interference for estimating treatment effects – spillover from treated to control units (Bowers et al. 2013) – was not possible; control participants donated when the activity was first introduced to them, before interacting with treated participants. But dependency among

²¹This estimate is from an OLS regression among control participants of donation amount on the expected donation amount of other players, with community fixed effects and individual-level controls.

²²Control respondents' priors on how much others would cooperate were mostly accurate, suggesting they could reasonably apply a behavioral rule of conditional cooperation without actively communicating (pg. SI.18).

treated units who interact with each other during the waiting period means that conventional methods for calculating uncertainty (standard errors) will be inaccurate.

In Table 6, we estimate effects of the treatment and its interaction with pre-treatment covariates using standard OLS regressions, while relying on a randomization inference approach to calculate p-values for hypothesis testing in the absence of well-defined standard errors. Randomization inference compares observed effect sizes against the simulated distribution of possible random assignments under sharp null hypotheses of no effect (or no interacted effect) for any unit (Gerber and Green 2012). This allows for valid hypothesis testing even in situations where conventional standard errors are invalid due to complex potential dependencies and interference among units (Rosenbaum 2007, Bowers et al. 2013).

P-values for the main treatment effect are calculated via randomization inference for the sharp null hypothesis of no effect for any unit, with the treatment effect estimated via OLS as the test statistic. P-values for the interaction of treatment with pre-treatment moderators (network locations) are calculated via randomization inference for a sharp null hypothesis of a constant effect for every unit at the overall treatment effect estimated via OLS without the interaction, and the coefficient on the interaction term as the test statistic. All models include the same control variables as described above.

Based on their pre-existing network positions, treated participants vary in their probability of exposure to other treated participants. We instead pre-registered an alternative approach for estimating treatment effects, drawing on Aronow and Samii (2017). This separately estimates (a) the combined direct and indirect effects of treatment among treated participants subject to spillovers from treated alters from (b) the isolated direct effect of treatment for treated participants who had no treated alters. By contrast, estimates from a simple OLS regression are an average over these possible direct and indirect effects. However, only 5% of our sample received the isolated direct exposure due to the relatively high network degree of most household heads in these communities. This limits the usefulness of comparisons between the direct versus combined

effects facilitated by our pre-registered approach. In practice, both OLS and the more complex pre-registered estimation procedure return substantively identical results. We default to OLS in the main text for expositional simplicity. Our pre-registered estimates are in the SI (pg. SI.4).

Experimental results

In Table 6, we find that the opportunity for communication did not increase cooperation (donations), in contrast to laboratory literature (e.g., Ostrom et al. 1992, Bochet et al. 2006). Instead, treated participants actually donated 0.23 GHS *less* than control participants (two-sided $p = 0.02$). For $H4$, we also do not find greater clustering (lower variance) in donations in the treatment condition (one-sided $p = 0.48$), indicating that donation amounts did not become any more coordinated when participants had an opportunity to interact.²³ Substantively, the change in the variance of donations in the treatment group (-0.01) represents less than 1% of the variance in the control group (1.22).

[Table 6 around here]

$H5a$ and $H5b$ instead test whether community leaders used the waiting period to encourage greater cooperation or push for their preferred projects. For $H5a$, Table 6 lists the coefficient on the interaction between the treatment and the proximity of each participant to the leader picked by the community. We provide separate estimates using each of the proximity measures. We find no evidence that treated participants donate more when they are closer to the leader. For $H5b$, the outcome is now an indicator for choosing the same project (e.g., water) as the leader's stated preference. We again find no evidence that treated participants become more likely to vote for the community leader's preferred project when they are more socially proximate to the leader. The estimated effect sizes are again substantively small. The largest coefficient on any of these

²³We test $H4$ non-parametrically as the difference in variances of the amount donated in the treatment and control groups, with p-values calculated via randomization inference under the assumption of a sharp null hypothesis of no effect for any unit. This null holds separately for each community (pg. SI.19).

interaction terms (*H5a* with shared nearest neighbors) represents a shift of just 5% of the baseline donation amount in the control group.

For *H6a-H6c*, the unit of analysis switches back to all dyads in each community. We are now interested in the effect of *both* members of the pair being treated, and thus having an opportunity to coordinate with each other, compared to other combinations in which at least one member of the pair would have played before having any chance to communicate. The quantity of interest is now the interaction between the indicator for both units being treated and the proximity of the pair. We again report results using both proximity measures. The outcomes remain as for *H3a-H3c*.

For all three hypotheses we find no evidence – using either measure – that participants who are more socially proximate become more likely to donate the same amount (*H6a*), vote for the same leader (*H6b*), or vote for the same project (*H6c*) when they are both treated and have a direct opportunity to coordinate. The substantive effect sizes remain small. Ultimately, even as social network locations predict behavior in the control group, our experimental results are inconsistent with communication among nodes increasing coordination or cooperation.

Exploratory analyses: what explains the null results?

Ruling out alternatives

In exploratory (non-pre-registered) analyses, we rule out two sets of alternative explanations. First, although we cannot observe evolutionary sorting of network ties or explicitly measure personality traits or norms, we can rule out more mundane mechanisms that could also explain why network locations are already correlated with behavior in the control group. There could be concern, for example, that wealthier or better educated people cluster in the center of each network and can simply afford to donate more. However, a wealth effect is unlikely as all estimates already control for assets, education, and employment. We also show robustness to controlling for alternative measures of wealth (pg. SI.19). In addition, there could be concern that central participants rationally donated more because they stand to benefit more from a local public good. But the SI demonstrates

in several ways that this is implausible (pg. SI.20).

Second, we rule out separate alternative explanations for the lack of expected treatment effects. The treatment would have led to an overall decline in donations if treated participants consumed some of the endowment while waiting. But the SI (pg. SI.21) provides strong reasons to doubt this can account for our results. It is also possible treated participants donated less because they were somehow influenced by the control group's prior behavior. For example, treated participants may have been less willing to contribute if they heard from their control alters that they either did not donate, or had already donated a lot – raising prospects to be a “sucker” or free-rider, respectively. But in column 1 of Table 7, we find no correlation between the donations of treated participants and the mean donation of their direct control alters – those most likely to tell treated participants how they had already behaved. Column 2 also finds no relationship between treated participants' donations and those of control ties when control ties sent a clearer (lower variance) signal in their own behavior. Columns 3 and 4 similarly show that treated participants did not update expectations of how much others would donate based on their control ties' behavior.

[Table 7 about here]

In addition, perhaps the amount of money at stake was so low that participants did not have sufficient incentives to try to coordinate. But significant communication occurred: the majority (55%) of treated participants reported speaking directly with members of other households about how they should behave; 44% of treated participants reported that they spoke with “many” (as opposed to “a few” or no) other households. Of those who spoke to at least one other household, 67% reported that they were encouraged to donate more. Only 6% reported that they were encouraged to donate less. The stakes were still sufficiently high that treated participants were induced to talk with each other – often across multiple different interactions – and that many attempted to coordinate on larger donations. This is less communication than in laboratory studies that explicitly prompt players to communicate (e.g., Bochet et al. 2006). But it is still a substantial amount of

interaction, and a more realistic approximation of the possible effects of communication outside the laboratory.²⁴

Community leaders were relatively less active during the waiting period. Only 15% of treated participants reported they spoke with the traditional chief about how they would behave; 10% and 6% reported similar conversations with the assembly member and religious leader, respectively. But chiefs did still try to influence their closest ties, speaking more to treated participants to whom they had a shorter path distance (pg. SI.22). Their message was to cooperate: 71% of treated participants who spoke with the chief reported he pushed them to make larger donations; only one reported the chief discouraged cooperation.

[Table 8 around here]

Yet it appears community members were unable to induce each other to cooperate. In Table 8 we examine associations between whether treated participants spoke with other households, or the chief, and their behavior. Columns 1 and 2 show that having conversations with other community members or the chief is not associated with amount donated. However, Columns 3 and 4 of Table 8 imply that some persuasion may still have occurred around choices of project type. In column 3, having spoken to the chief is associated with being more likely to have changed project preferences in the second interview from the first. Column 4 shows that speaking to other households predicts being more likely to vote for the same project as the modal preference in the community. Indeed, 90% of the conversations treated participants held during the waiting period touched on which project to choose. These conversations may thus have helped participants coordinate on which project to support, even as they failed to coordinate on donation amounts.

A final means to show that communication was mostly non-impactful is to estimate the effect of treatment on participants' expectations of others' behavior. We do not find statistically significant effects of treatment on the accuracy of participants' expectations of what others would donate,

²⁴Future research could valuably explore whether more communication would have occurred, or been more impactful, if the waiting period was longer than our 1-2 days.

the probability of a participant reporting they “don’t know” what others would donate, or the overall amount participants expected others will donate (pg. SI.24).²⁵ This rules out additional alternative explanations for the results. For example, it is not the case that treated participants donated marginally less overall because they had learned through communication that their peers were likely to shirk or instead had gained a more accurate view of how much the community as a whole was contributing.

Behavior already encoded in networks?

Instead, communication may have mostly failed to update participants’ behavior and expectations because it could not displace more deeply-engrained, pre-existing behaviors already encoded in network positions. With the exception of preferences about which project to pick (*H3c*) – the single outcome for which Table 8 suggests that communication may have facilitated coordination – each of the results in Table 5 for the control group also carried over to the treatment group.

Table 9 repeats Table 5, subset to the treatment group instead. More central treated participants still donated more (*H1*), just as when there was no communication. The coefficient on an interaction between eigenvector centrality and treatment is small and insignificant (pg. SI.23). Participants who were more proximate to each other were also still more likely to choose the same leaders to manage the donations (*H3b*). The magnitudes of these relationships are virtually identical to in the control group (Table 5). Similarly, evidence of conditional cooperation is still present among treated participants (pg. SI.23).

[Table 9 about here]

This is not definitive evidence that network locations are already sorted by or encode players’ types, traits, or cooperative norms. To observe this more systematically, we would need to be able

²⁵However, when instead using the pre-registered estimation procedure (pg. SI.4), there is a small positive effect of treatment on reporting you “don’t know” how much others will donate at the $p < 0.1$ level, suggesting communication might have complicated, rather than clarified participants’ expectations. This is the only result not robust to the choice of estimation procedure.

track the dynamic sorting of network ties over time, examining the process by which different behaviors become associated with different network positions. This is not possible in a single-shot study such as ours. However, these patterns are more consistent with this category of expectations than any of the alternatives examined here.

Conclusion

Scholars increasingly collect social network data in developing countries under the assumption that the social relationships in which individuals are enmeshed are influencing their behavior (Larson and Lewis 2017, Ferrali et al. 2018). Indeed, social relationships among community members, and between community members and their leaders, are central to many theories of cooperative local public goods provision in the developing world. But when we observe that social positioning of community members explains their behavior, there is cause for significant caution before assuming this is a result of active influence that passes among social ties.

We show that social centrality predicts cooperation in the co-production of a local public good and find that strategies of “conditional cooperation” commonly observed in the laboratory extend to the field, especially among the most central individuals within local social networks. But these results exist even absent opportunities for active social influence from peers or leaders. We do not dispute that the transmission of information through social network ties still may improve cooperation in other situations, especially those with real opportunities for sanctioning. Yet we show that a reasonably high amount of cooperation is *already* possible even without any active social influence at all, with social network locations *already* closely mapping to cooperative behavior.

Our experiment suggests several promising avenues for future research. First, we discuss, but cannot directly test, several mechanisms for why real-life network locations are correlated with cooperative behavior absent active social influence: sorting on cooperative types or traits, sorting on previously internalized norms, or that pre-existing network locations cause players to internalize different behavioral rules and expectations. By better examining the dynamic processes through

which real networks form, future work can valuably explore how cooperation and other political behaviors become embedded in social structures.

Second, while the two views of networks described above are not mutually exclusive outside the context of our experiment, all three proposed mechanisms suggest that our alternative view of networks as maps of behavior will be comparatively more important in real-world networks that have already evolved over many prior rounds of interaction. The marginal effect of additional social influence in a particular cooperative situation may be smaller when the actors have a long history of interaction than when interacting for the first time. Yet many of our core theories of cooperation are still primarily evaluated through laboratory or lab-in-the-field studies, in which the participants are usually strangers or anonymous to each other. Understanding the complementarities between these theories of networks requires continuing to shift towards the study of real cooperation in real communities with fully-evolved social networks.

References

- Alatas, Vivi, Abhijit Banerjee, Arun G Chandrasekhar, Rema Hanna and Benjamin A Olken. 2016. “Network structure and the aggregation of information: Theory and evidence from Indonesia.” *American Economic Review* 106(7):1663–1704.
- Andreoni, James and Justin M. Rao. 2011. “The Power of Asking: How Communication Affects Selfishness, Empathy, and Altruism.” *Journal of Public Economics* 95:513–520.
- Aronow, Peter M. and Cyrus Samii. 2017. “Estimating Average Causal Effects Under General Interference, with Application to a Social Network Experiment.” *The Annals of Applied Statistics* 11(4):1912–1947.
- Aronow, Peter M., Cyrus Samii and Valentina A. Assenova. 2015. “Cluster-Robust Variance Estimation for Dyadic Data.” *Political Analysis* 23(4):564–577.
- Baldassarri, Delia and Guy Grossman. 2011. “Centralized Sanctioning and Legitimate Authority Promote Cooperation in Humans.” *PNAS* 108(27):11023–11027.
- Baldwin, Kate. 2015. *The Paradox of Traditional Chiefs in Democratic Africa*. New York: Cambridge University Press.
- Balliet, Daniel. 2010. “Communication and Cooperation in Social Dilemmas: A Meta-Analytic Review.” *Journal of Conflict Resolution* 54(1):39–57.
- Balliet, Daniel and Paul A.M. Van Lange. 2013. “Trust, Conflict, and Cooperation: A Meta-Analysis.” *Psychological Bulletin* 139(5):1090–1112.
- Bicchieri, Christina. 2017. *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. New York: Oxford University Press.
- Bochet, Olivier, Talbot Page and Louis Putterman. 2006. “Communication and Punishment in Voluntary Contribution Experiments.” *Journal of Economic Behavior and Organization* 60(1):11–26.
- Bowers, Jake, Mark M. Fredrickson and Costas Panagopoulos. 2013. “Reasoning about Interference Between Units: A General Framework.” *Political Analysis* 21:97–124.
- Burt, Ronald S. 1982. *Toward a Structural Theory of Action: Network Models of Social Structure, Perception and Action*. Academic Press.
- Charness, Gary and Martin Dufwenberg. 2006. “Promises and Partnership.” *Econometrica* 74(6):1579–1601.
- Coppock, Alexander. 2018. “Avoid Post-Treatment Bias in Audit Experiments.” Forthcoming, *Journal of Experimental Political Science*.

- Cruz, Cesi, Julien Labonne and Pablo Querubin. 2017. "Politician Family Networks and Electoral Outcomes: Evidence from the Philippines." *American Economic Review* 107(10):3006–37.
- Dionne, Kim Yi. 2015. "Social networks, ethnic diversity, and cooperative behavior in rural Malawi." *Journal of Theoretical Politics* 27(4):522–543.
- Eguiluz, Victor M., Martin G. Zimmermann, Camilo J. Cela-Conde and Maxi San Miguel. 2005. "Cooperation and the Emergence of Role Differentiation in the Dynamics of Social Networks." *American Journal of Sociology* 110(4):977–1008.
- Elster, Jon. 1989. "Social Norms and Economic Theory." *Journal of Economic Perspectives* 3(4):99–117.
- Eubank, Nicholas, Guy Grossman, Melina Platas and Jonathan Rodden. 2019. "Viral Voting: Social Networks and Political Participation." Forthcoming, *Quarterly Journal of Political Science*.
- Ferrali, Romain, Guy Grossman, Melina Platas and Jonathan Rodden. 2019. "It Takes a Village: Peer Effects and Externalities in Technology Adoption." Forthcoming, *American Journal of Political Science*.
- Fischbacher, Urs and Simon Gächter. 2010. "Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments." *American Economic Review* 100(1):541–556.
- Fowler, James H. 2005. Turnout in a Small World. In *Social Logic of Politics*, ed. Alan Zuckerman. Philadelphia: Temple University Press pp. 269–287.
- Gallo, Edoardo and Chang Yan. 2015. "The Effects of Reputational and Social Knowledge on Cooperation." *PNAS* 112(12):3647–3652.
- Gerber, Alan S. and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation*. New York: W.W. Norton.
- Gottlieb, Jessica. 2017. "Explaining Variation in Broker Strategies: A Lab-in-the-Field Experiment in Senegal." *Comparative Political Studies* 50(11):1556–1592.
- Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83(6):1420–1443.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2007. "Why Does Ethnic Diversity Undermine Public Goods Provision?" *American Political Science Review* 101(4):709–725.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the Dilemmas of Collective Action*. Russell Sage Foundation.
- Larson, Jennifer M. and Janet I. Lewis. 2017. "Ethnic Networks." *American Journal of Political Science* 61(2):350–364.

- Larson, Jennifer M. and Janet I. Lewis. 2019. "Measuring Networks in the Field." Forthcoming, *Political Science Research and Methods*.
- McPherson, Miller, Lynn Smith-Lovin and James M Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27(1):415–444.
- Miguel, Edward and Mary Kay Gugerty. 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics* 89:2325–2368.
- Ministry of Energy. 2010. "National Electrification Scheme (NES) Master Plan Review (2011-2020)." https://www.mida.gov.gh/pages/view/111/NES_Master_Plan_Review_Executive_Summary_Main_Report.pdf Accessed 10 March 2020.
- Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press.
- Ostrom, Elinor. 2000. "Collective Action and the Evolution of Social Norms." *Journal of Economic Perspectives* 14(3):137–158.
- Ostrom, Elinor, James Walker and Roy Gardner. 1992. "Covenants With and Without A Sword: Self-Governance is Possible." *American Political Science Review* 86(2):404–417.
- Rand, David G., Samuel Arbesman and Nicholas Christakis. 2011. "Dynamic Social Networks Promote Cooperation in Experiments with Humans." *PNAS* 108(48):19193–19198.
- Rolfe, Meredith. 2012. *Voter turnout: A social theory of political participation*. Cambridge University Press.
- Rosenbaum, Paul R. 2007. "Interference Between Units in Randomized Experiments." *Journal of the American Statistical Association* 102(477):191–200.
- Sally, David. 1995. "Conversation and Cooperation in Social Dilemmas: A Meta-Analysis of Experiments from 1958 to 1992." *Rationality and Society* 7(1):58–92.
- Shalizi, Cosma Rohilla and Andrew C. Thomas. 2011. "Homophily and Contagion Are Generically Confounded in Observational Network Studies." *Sociological Methods and Research* 40(2):211–239.
- Siegel, David A. 2009. "Social Networks and Collective Action." *American Journal of Political Science* 53(1):122–138.
- Sinclair, Betsy. 2012. *The Social Citizen: Peer Networks and Political Behavior*. Chicago: University of Chicago Press.

- Vanberg, Christoph. 2008. “Why Do People Keep Their Promises? An Experimental Test of Two Explanations.” *Econometrica* 76(6):1467–1480.
- Volk, Stefan, Christian Thoni and Winfried Ruigrok. 2012. “Temporal Stability and Psychological Foundations of Cooperation Preferences.” *Journal of Economic Behavior and Organization* 81:664–676.
- Wang, Jing, Siddharth Suri and Duncan J. Watts. 2012. “Cooperation and Assortativity with Dynamic Partner Updating.” *PNAS* 109(36):14363–14368.
- White, Harrison C, Scott A Boorman and Ronald L Breiger. 1976. “Social structure from multiple networks: Blockmodels of roles and positions.” *American Journal of Sociology* 81(4):730–780.

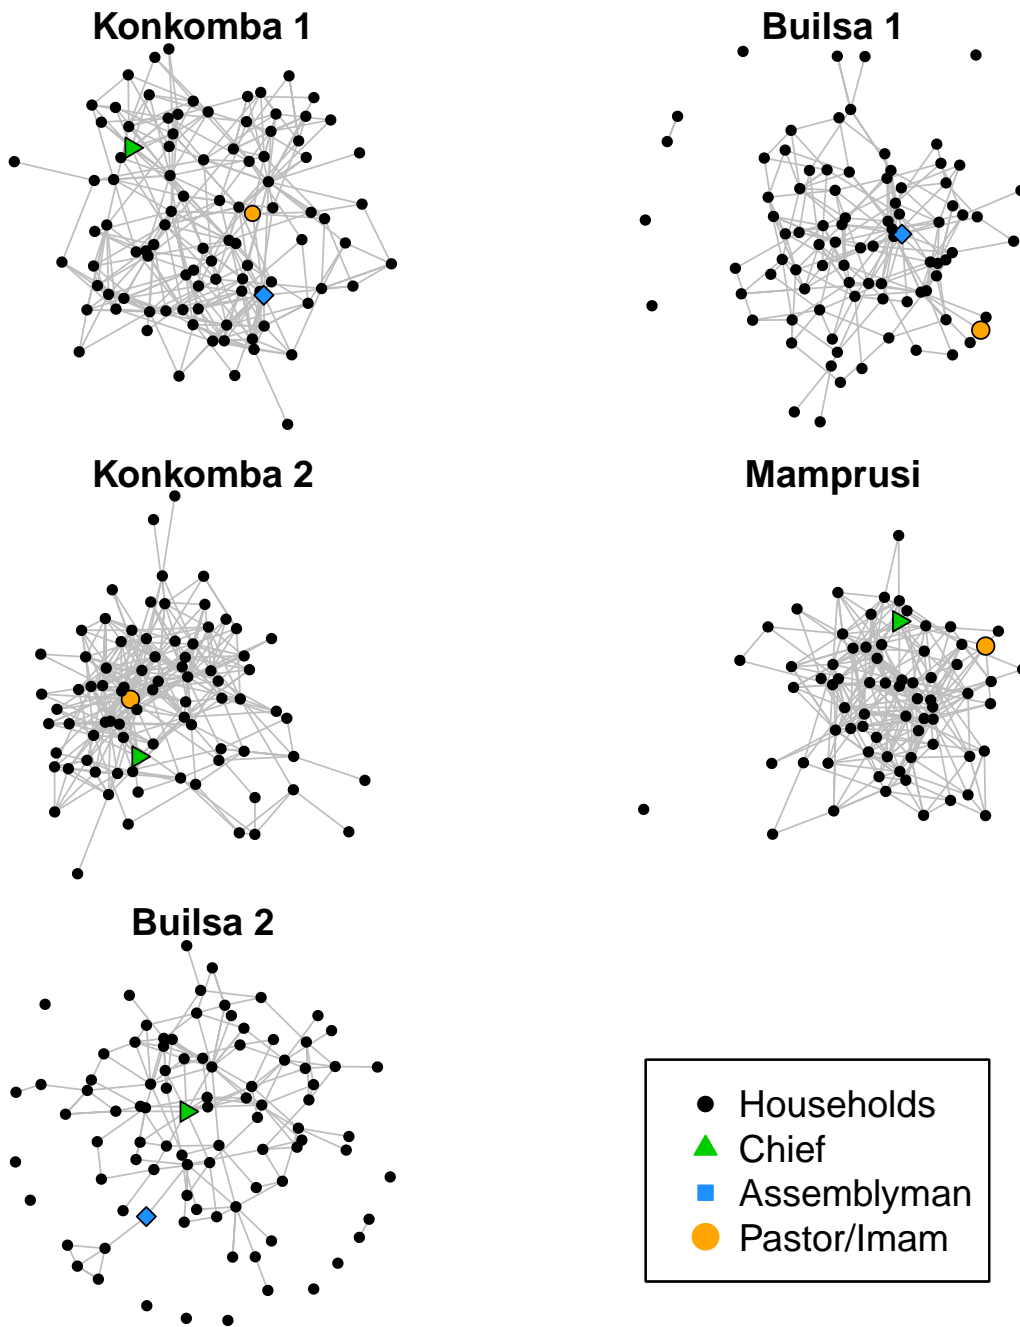


Figure 1: *Network diagrams of study communities*: the nodes are household heads and all ties are undirected. Community leaders' locations are highlighted.

Table 1: Name generator prompts

Domain	Prompt
Friends	Please name up to five persons who you consider to be your closest friends in this community:
Problem solving	Please name up to five persons in this community that you would contact first if you had a household problem or emergency that needed to be resolved:
Political conversation	Please name up to five persons in this community with whom you regularly speak the most about politics:
Family ties	Please name up to five other persons in this community to whom you are closely related:

Table 2: Pre-registered non-experimental hypotheses

H1	<i>All else equal, participants who are more central in their community's social network will donate more.</i>
H2	<i>All else equal, participants who are more proximate in the community social network to the community leader (i.e., chief, assembly member, or religious leader) that they believe most other community members will pick to manage the public good will donate more.</i>
H3a	<i>All else equal, participants who are more proximate to each other in the community social network will behave more similarly in donation amount.</i>
H3b	<i>All else equal, participants who are more proximate to each other in the community social network are more likely to choose the same leader to manage the public good.</i>
H3c	<i>All else equal, participants who are more proximate to each other in the community social network are more likely to request the same project.</i>

Table 3: Pre-registered experimental hypotheses

H4	<i>The variance in donation amounts among participants assigned to wait before deciding whether to donate is smaller than the variance among participants asked to donate immediately.</i>
H5a	<i>Participants who are assigned to wait before deciding whether to donate make larger donations than respondents asked to donate immediately when participants are more proximate in the social network to the modal community leader picked by other participants in the community to manage the public good.</i>
H5b	<i>Participants who are assigned to wait before deciding whether to donate are more likely than respondents asked to donate immediately to choose the project preferred by the modal community leader picked by other participants in the community to manage the public good, especially when these participants are more proximate to that leader in the community social network.</i>
H6a	<i>Dyads of participants in the same community make more similar donation amounts to each other when they are both assigned to wait before deciding whether to donate than dyads assigned to other combinations of treatment conditions, especially when these dyads are more proximate to each other in the social network.</i>
H6b	<i>Dyads of participants in the same community are more likely to choose the same leader to manage the public good when they are both assigned to wait before deciding whether to donate than dyads assigned to other combinations of treatment conditions, especially when these dyads are more proximate to each other in the social network.</i>
H6c	<i>Dyads of participants in the same community are more likely to request the same project when they are both assigned to wait before deciding whether to donate than dyads assigned to other combinations of treatment conditions, especially when these dyads are more proximate to each other in the social network.</i>

Table 4: Results for non-experimental hypotheses

	<i>Outcome variable</i>	<i>Explanatory variable</i>	β	p	95% CI	N
<i>H1</i>	Donation (0-5 GHS)	Eigenvector centrality	2.82	<0.01	(1.27, 4.37)	402
<i>H2</i>	Donation (0-5 GHS)	Proximity (SNN) to leader	0.08	0.21	(-0.05, 0.21)	318
		Proximity (MPD) to leader	0.07	0.24	(-0.05, 0.20)	319
<i>H3a</i>	Abs. value of diff. in donations (0-5 GHS)	Proximity (SNN) of dyad	-0.01	0.80	(-0.05, 0.04)	16,162
		Proximity (MPD) of dyad	0.05	0.12	(-0.01, 0.12)	16,162
<i>H3b</i>	Same leader preference in dyad (0,1)	Proximity (SNN) of dyad	0.03	<0.01	(0.01, 0.05)	16,178
		Proximity (MPD) of dyad	-0.03	<0.01	(-0.05, -0.01)	16,178
<i>H3c</i>	Same project preference in dyad (0,1)	Proximity (SNN) of dyad	0.03	0.02	(0.00, 0.05)	15,965
		Proximity (MPD) of dyad	-0.01	0.23	(-0.04, 0.01)	15,965

OLS estimates, as described in the text. Standard errors adjusted for dyadic data for *H3a-H3b* (Aronow et al. 2015).

Table 5: Network locations and behavior: control group only

	<i>Outcome variable</i>	<i>Explanatory variable</i>	β	p	95% CI	N
<i>H1</i>	Donation (0-5 GHS)	Eigenvector centrality	2.62	0.04	(0.14, 5.10)	195
<i>H3b</i>	Same leader preference in dyad (0,1)	Proximity (SNN) of dyad	0.03	0.04	(0.00, 0.06)	3,771
		Proximity (MPD) of dyad	-0.04	0.05	(-0.07, 0.00)	3,771
<i>H3c</i>	Same project preference in dyad (0,1)	Proximity (SNN) of dyad	0.05	<0.01	(0.02, 0.08)	3,690
		Proximity (MPD) of dyad	-0.05	0.01	(-0.09, -0.01)	3,690

All estimates calculated as in Table 4 above, subset to control participants only.

Table 6: Results for experimental hypotheses

	<i>Outcome variable</i>	<i>Explanatory variable</i>	β	p	N
–	Donation (0-5 GHS)	Treatment status (0,1)	-0.23	0.02 ^a	402
<i>H4</i>	Variance in donation amount (0-5 GHS)	Treatment status (0,1)	-0.01	0.48	402
<i>H5a</i>	Donation (0-5 GHS)	Proximity (SNN) to leader * treatment status (0,1)	-0.11	0.85	318
		Proximity (MPD) to leader * treatment status (0,1)	-0.09	0.24	319
<i>H5b</i>	Same project pref. as leader (0,1)	Proximity (SNN) to leader * treatment status (0,1)	0.03	0.23	335
		Proximity (MPD) to leader * treatment status (0,1)	0.02	0.71	335
<i>H6a</i>	Abs. value of diff. in donation (0-5 GHS)	Proximity (SNN) of dyad * treatment status (0,1)	0.02	0.68	16,162
		Proximity (MPD) of dyad * treatment status (0,1)	-0.05	0.87	16,162
<i>H6b</i>	Same leader pref. in dyad (0,1)	Proximity (SNN) of dyad * treatment status (0,1)	-0.00	0.51	16,178
		Proximity (MPD) of dyad * treatment status (0,1)	-0.01	0.27	16,178
<i>H6c</i>	Same proj. pref. in dyad (0,1)	Proximity (SNN) of dyad * treatment status (0,1)	-0.01	0.63	15,965
		Proximity (MPD) of dyad * treatment status (0,1)	0.01	0.74	15,965

P-values calculated via randomization inference, as described in the text. *a*: The first test is a two-sided p-value. All other tests are one-sided, given the directional predictions in the hypotheses.

Table 7: Donations of treated participants by donations of direct control group ties

	1	2	3	4
<i>Outcome:</i>	Donation amount (GHS)	Donation amount (GHS)	Expected donation by others (GHS)	Expected donation by others (GHS)
Average donation (GHS) of control alters	0.13 (0.09)	0.13 (0.10)	0.07 (0.09)	0.07 (0.10)
Variance in donations (GHS) of control alters		-0.14 (0.34)		-0.14 (0.32)
Avg of control alters * variance of control alters		0.04 (0.11)		0.04 (0.11)
Community fixed effects	Y	Y	Y	Y
Demographic controls	Y	Y	Y	Y
<i>N</i>	153	153	104	104
adj. R^2	0.36	0.35	0.20	0.19

Significant at: * $p < .05$; ** $p < .01$; *** $p < .001$. Participants with no control alters are dropped. Participants who “don’t know” how much others will donate are dropped in columns 3 and 4. All models are OLS.

Table 8: Game behavior by conversations during the waiting period, treated participants only

<i>Outcome:</i>	1	2	3	4
	Donation amount (GHS)	Deviation from median donation (GHS)	Change in project preference (0,1)	Selects modal project preference (0,1)
Spoke to other households (0,1)	0.04 (0.14)	0.04 (0.14)	0.04 (0.06)	0.15* (0.07)
Spoke to chief (0,1)	-0.02 (0.19)	-0.02 (0.19)	0.18* (0.09)	0.08 (0.10)
Comm. fixed effects	Y	Y	Y	Y
Demographic controls	Y	Y	Y	Y
<i>N</i>	207	207	193	198
adj. <i>R</i> ²	0.33	0.11	0.05	0.08

Significant at: **p* < .05; ***p* < .01; ****p* < .001. All models are OLS.

Table 9: Network locations and behavior: treated group only

	<i>Outcome variable</i>	<i>Explanatory variable</i>	β	p	95% CI	N
<i>H1</i>	Donation (0-5 GHS)	Eigenvector centrality	2.81	<0.01	(0.89, 4.73)	207
<i>H3b</i>	Same leader preference in dyad (0,1)	Proximity (SNN) of dyad	0.03	0.05	(0.00, 0.05)	4,257
		Proximity (MPD) of dyad	-0.03	0.05	(-0.06, 0.00)	4,257
<i>H3c</i>	Same project preference in dyad (0,1)	Proximity (SNN) of dyad	0.01	0.49	(-0.02, 0.04)	4,227
		Proximity (MPD) of dyad	0.02	0.25	(-0.01, 0.04)	4,227

All estimates calculated as in Table 4 above, subset to treated participants only.