

From Chatter to Action: How Social Networks Inform and Motivate in Rural Uganda*

Jennifer M. Larson[†], Janet I. Lewis[‡] and Pedro Rodríguez[§]

March 16, 2018

Abstract

From election campaigns to public service announcements, numerous political activities and policy interventions hinge on the spread of new information that motivates behavior. However, few studies directly examine the process by which information spreads via word-of-mouth, or compare that to the separate process by which those who learn the information act on it. Using a novel design that seeded valuable, actionable information in rural Uganda, we show that both processes depend on a group's social network, but in different ways. Information spread via a straightforward contagion process. Behavior did not spread so simply; those who acted were socially proximate to the earliest actors and saw endorsement by their most intimate ties. Puzzlingly, while those most central in the network were most likely to hear the information, those who ultimately acted were the *least* central among the informed. Connections to highly-connected peers may generate pressure to refrain from doing something new.

9,894 words

*We thank Erin Baggott, Pablo Barberá, Robert Bates, Andrew Bertoli, Brett Carter, Andrew Coe, Cali Ellis, Ben Graham, Franziska Keller, James Lo, Jonathan Markowitz, Bryn Rosenfeld, Amy Erica Smith, Dustin Tingley, the Yale University Leitner Seminar participants, and especially Pia Raffler for helpful comments on pieces of this project. We are also grateful to Vianney Mbonigaba and Innovations for Poverty Action's Uganda office for crucial advice about implementing the study; to Juliet Ajilong and an excellent team of research assistants for their efforts collecting the data in Uganda; to Batuhan Gorgulu for careful data entry assistance; and to the study participants in Abalang and Mugana for generously giving their time. We received support for this project from Harvard's Institute for Quantitative Social Science. Any views expressed are the authors' and do not necessarily reflect the official policy or position of the United States Naval Academy, Department of Defense, or the U.S. Government.

[†]**Corresponding Author.** Assistant Professor of Politics, New York University, jenn.larson@nyu.edu.

[‡]Assistant Professor of Political Science, U.S. Naval Academy.

[§]Ph .D. Candidate, Department of Politics, New York University, and Instituto de Estudios Superiores de Administración, Caracas, Venezuela.

1 Introduction

For electoral candidates, rebel leaders, and policymakers alike, changing people’s behavior by providing them with new information is a fundamental goal. Candidates for public office try to influence participation and vote choice with information about their policy platforms, character, and past performance. Rebel group leaders aim to attract new recruits by persuading local villagers about their cause and fighting prowess. Public health officials seek to contain new diseases by teaching people to alter their sanitation practices. Non-governmental organizations work to motivate citizens to hold officials accountable by providing information about public service quality and corruption.

In sum, outcomes ranging from voting behavior to violent conflict to public health crises depend not only on new information spreading widely, but also on people being motivated to act on that information. Despite how foundational these processes are to core theories in political science and to the success of myriad policy interventions, few studies directly examine information dissemination; even fewer consider the potential differences between hearing information and acting on it.¹

Yet surely these processes often operate differently. For example, while certain attributes may increase the chance that someone hears news about a politician’s involvement in a corruption scandal, different attributes may increase the chance that someone who heard the news turns out to vote for an opposing candidate. Likewise, it is one thing for a villager to learn that international health workers have arrived nearby to deal with a local outbreak of Ebola. It is quite another for him to stop tending to sick family members in the usual ways and instead take them to the quarantine tent. This gap between information provision and behavioral change may account for the limited effect of numerous, costly programs that

¹The one exception of which we are aware is Ferrali et al. (2018) which separately considers the spread of a new mobile platform and its use. They explore the village-level conditions under which behavior with positive externalities would spread through social networks.

seek to improve governance and development outcomes via information provision to citizens (Lieberman, Posner and Tsai, 2014).

Consequently, while it is well known that information can spread through social networks (Banerjee et al., 2013; Mobius, Phan and Szeidl, 2015; Larson and Lewis, 2017), that peers can influence a person’s behavior (Sinclair, McConnell and Michelson, 2013; Bond et al., 2017; Eubank et al., 2017), and that knowing something is a prerequisite to acting on it, little is known about exactly *how* a social network operates to convert that information into action. Studies rarely have all of the necessary features to unpack this process, which include trackable, novel information; a separate measure of hearing and acting on information; and a rich measure of social networks.

We implemented a study in Abalang, a rural village in the Teso region of Uganda, that offers an unusual view of how new information travels naturally through social networks, and how people receiving the information decide whether or not to act on it. We chose seven households at random and provided them with information in person that in three days, an event would be held just outside the village at which every adult in attendance would receive a valuable block of soap in exchange for taking our survey. After seeding this information, the research team left the village for the three interim days, and then returned to host the event during the next three days. Although a soap event may seem innocuous, it was so unusual that some villagers regarded the news with suspicion, envisaging risks to attending. Despite the care of our enumeration team, a villager who heard the news could not be certain that the event would be held, that attending would be valuable, that attending would not entail risks (including of a supernatural form, discussed below), or that attending would be viewed favorably by fellow villagers.

Of the roughly 1,400 residents of Abalang and its outskirts, 138 attended the event. Surveys at the event inquired about a person’s sources of information about the event, as well as demographic and social network information. Once the event concluded, surveys

were conducted throughout the village to complete our mapping of social networks and to measure the reach of information. Our study is unusual in that we have a record of those who both heard the information and changed behavior based on it—138 people heard and attended the event— as well as a measure of the reach of information that did not change behavior—an additional 130 people in our data heard about the event but did not attend.

We show that, as expected, social networks played a key role in the dissemination of information throughout the community, as well as the motivation to act on it. However, *how* social networks informed is importantly different from the way that social networks motivated action.

We begin by examining evidence for contagion. Information appears to have spread through the community in this straightforward way. Information spread exclusively through face-to-face interactions along social ties of various types; the more of a person’s social contacts who heard about the event—both in absolute number and as a proportion of peers—the more likely a person was to have heard about the event. Attendance, on the other hand, did not spread so simply. It is robustly *not* the case that the more of a person’s social contacts attended the event, the more likely the person was to have attended, even conditional on hearing about the event.

Instead, the features of social networks that predict attendance are those that would be useful for establishing the credibility of the information. The strongest predictor of attendance is social proximity to one of the eleven villagers who attended on the first day. Occupying a network position one step closer to an early attender—a friend-of-a-friend instead of a friend-of-a-friend-of-a-friend, say—is associated with a 16% increase in one’s likelihood of attending. This relationship is larger and stronger even than social proximity to a seed, a person who had first access to the information, who was personally visited by an enumerator, and who was given an official information sheet. Social contacts who test out new behavior first and are able to credibly convey its benefit or low risk appear crucial to its subsequent

popularity.²

Moreover, while general social endorsement of attending is unrelated to attendance, the endorsement by *certain kinds* of network ties is related. If more of the people with whom a person shares meals or personally visits attended, that person is more likely to have attended. While any tie may inform, only intimate ties motivate. Combined, these findings are consistent with other work showing that social ties can serve as latent sources of trustworthy connections to important resources or information (Barr, Ensminger and Johnson, 2010; Cruz, Labonne and Querubin, 2017).

Furthermore, we document a surprising relationship between network centrality and the outcomes of hearing and attending. Being in a more centrally-located network position (in the sense of being near lots of other people, or high closeness centrality), and having connections to peers who themselves have lots of connections (high eigenvector centrality) is, as expected, positively related to hearing about the event. People in these network positions have privileged access to information flowing through the network. However, among those who were informed, it is robustly the *less* central who acted on the information and attended. We discuss possible explanations, including that the most central face a higher opportunity cost to acting on new information, and that those with less well-connected peers are freer to take novel actions, though adjudicating among these mechanisms is left for future research.

As a whole, our findings suggest that social networks play a role in collective behavior via two processes: a simple transmission of information along social ties, and a more complex process that converts that information into action. In the present context of valuable, actionable information spread in Abalang, the latter appears to be a function of lending credibility to the information. Social ties, especially intimate ones, may operate to reassure, reinforce, vet, and pass judgment on the action suggested by the information. Consequently,

²This result hints that early attenders in our setting may serve a function similar to the “social referents” that are most influential in reducing bullying in the setting of Paluck, Shepherd and Aronow (2016).

social proximity to those who tested the action early is a powerful motivator, more powerful than proximity to those who heard first, and although the central network positions are indeed more likely to encounter new information, they are not necessarily more likely to be convinced to act on it.

These conclusions are drawn from the spread of a single type of information in a single setting. Of course much is left for future work, especially to establish the generality of these findings. However, features of this study give cause for optimism about generalizing beyond Abalang, Uganda. The fact that the information was brand new, seeded at random, and was regarded as uncertain and possibly risky means that we may be tapping into the way villagers would respond to a large range of new and potentially politically-relevant information, such as news that a rebel group has begun operating nearby, or news that a politician was involved in a scandal (we discuss this potential correspondence in the next section).³ To the extent that villagers feared attending and found safety in confirmatory information from early attenders, our intervention may be picking up behavior relevant to collective action. Additionally, news spread from person to person exclusively through in-person exchanges. This mode of sharing information is thought to be particularly important in rural, developing country contexts (Banerjee et al., 2013), and for this reason we see the setting of rural Africa as a virtue; however, such exchanges serve as an important means of informing and deliberating in developed country settings as well (e.g. Gerber and Green, 2000).

Beyond the insights into how social networks function to motivate behavior, our results also highlight important policy implications. When interventions include information campaigns, for instance about quality or corruption of local governance, best health practices, fertilizer use, and so on, an important design consideration is whom to choose as the injection points. Our findings suggest potentially different answers depending on whether the goal is

³Testing sensitive political information directly was ruled out due to ethical considerations.

exclusively to spread information or to change behavior based on the information. While seeding information with people who have high network centrality may be optimal for the wide reach of information (Banerjee et al., 2013), a different set of injection points may be optimal for encouraging the most people to act on it. In Abalang (and we suspect elsewhere), hearing information is not sufficient for someone to act on it. Rather than target the most central, our findings suggest that targeting a tight-knit pocket within the network, central or not, and encouraging them to act early may generate the largest cascade of behavior. Others close to them, especially their most intimate social contacts, are likely to follow suit.

2 Relationship to Existing Literature

In countless theories of political behavior, the outcome of interest depends on how well people are informed about something new. Retrospective voting theories suggest that elected officials will be held accountable at the polls for their performance if voters receive information about politician quality (Przeworski, Stokes and Manin, 1999). Theories of ethnic politics explain outcomes such as voting behavior and public goods provision in terms of how easily information flows within but not between ethnic groups (Chandra, 2004; Miguel and Gugerty, 2005; Habyarimana et al., 2009). Nascent rebel groups' likelihood of survival depends on their ability to seed favorable information about themselves among the local civilian population (Lewis, 2017). Rumors that spread too widely can spark conflict (Varshney, 2003) while wide-reaching gossip about misbehavior can stave off conflict (Fearon and Laitin, 1996; Larson, 2017). Protests depend on what people know about an upcoming protest and expectations as to other peoples' thoughts about it (Chwe, 2000; Siegel, 2009).

The set of relationships that comprise a social network are important sources of new information in any context, especially in the developing world where informal ties often serve as a vehicle for information and services that elsewhere could be provided through other chan-

nels (Banerjee et al., 2013). Theory that seeks to explain the spread of something through a network, “diffusion,” borrows heavily from epidemiology. According to these models of contagion, individuals “infected” with ideas or behavior are contagious and can infect network neighbors (Newman, 2000; Jackson and Rogers, 2007). Hence, collective behavioral changes can emerge when enough prospective participants are infected with the desire to participate by enough other participants, resulting in cascades of large-scale collective action (Marwell, Oliver and Prahl, 1988; Gould, 1993). Variants acknowledge that behavior may spread according to relatively more complicated rules than a disease, requiring repeated exposure, or exposure to enough others who are infected (Valente, 1996; Chwe, 1999; Dodds and Watts, 2004; Chiang, 2007; Centola and Macy, 2007; Centola, 2013). According to these theories, if a person is sufficiently exposed to enough others who will undertake a certain action, the person will take that action.

Theories of political behavior that emphasize information dissemination, and theories of diffusion through networks in general, offer a number of important insights, but two gaps in the literature are particularly troublesome. First, despite widespread agreement on its importance, direct investigation of information diffusion through social networks and its consequences is rare in rural, developing country contexts. Second, few empirical studies in any context examine the outcomes of learning information and acting on it *separately*.⁴

As we show, classical models of diffusion explain the spread of information through Abalang well, but perform poorly as an explanation for the spread of acting on it. While this could mean that networks are irrelevant for behavior, we show that instead networks may matter for behavior in important ways that differ from classical diffusion. We further show that without separate measures, a study is unable to disentangle “information effects”

⁴This issue also affects the social science theories enumerated above for which information diffusion is an important mechanism. For example, in theories for which information diffusion leads to improved collective action among co-ethnics, processes of information diffusion are rarely considered distinctly from the processes of behavior.

from other effects that can have important, distinct policy implications. Better understanding of these distinctions may help advance knowledge about why information interventions do not always have their intended effect on political and other actions. For example, field experiments examining interventions to provide information about corruption to voters in Mexico (Chong et al., 2015), MP performance to voters in Uganda (Humphreys and Weinstein, 2012), and how parents can facilitate their child’s learning in Kenya (Lieberman, Posner and Tsai, 2014), showed respectively that this did not produce the expected increase in opposition vote share in Mexico or Uganda, nor did it increase parental involvement in education in Kenya.

Our study measures information spread and its consequences in a rural, developing country setting by eliciting social networks with surveys (see Conley and Udry, 2010; Barr, Ensminger and Johnson, 2010; Alatas et al., 2016; Apicella et al., 2012; Banerjee et al., 2013; Cai, De Janvry and Sadoulet, 2015; Larson and Lewis, 2017; Eubank et al., 2017; Ferrali et al., 2018). Existing empirical network studies document that, in general, behavior responds to the behavior of network neighbors (Munshi, 2004; Bandiera and Rasul, 2006; Conley and Udry, 2010). However, these studies typically find that, conditional on network neighbors sharing information, there is no further effect of networks on behavior. For instance, in the study most similar to our own, Banerjee et al. (2013) seed information about a new microfinance program with a small number of villagers and then track ultimate participation in 43 villages in India. While they find strong evidence that the social network is responsible for spreading the information about the program, they find no effect of the network on participation beyond that due to spreading information.⁵ Likewise, Cai, De Janvry and Sadoulet (2015) seed information about a new weather insurance product, track ultimate uptake, and find no evidence of a social network effect beyond the information effect. Ac-

⁵While Banerjee et al. (2013) shares our interest in separating the spread of information from the spread of behavior, no data were collected on the spread of information: these effects are estimated from a structural model using participation and social network data.

ording to these studies, information spreads through networks, but conditional on hearing the information, the behavior of social network neighbors has no effect on whether people act on the information or not.

We implement a similar study, seeding information and measuring behavior in response to it. As an important departure from previous studies, we separately track the spread of the information, even among those who did not act in response to it.⁶ We show that, as expected, information spreads more easily than behavior. We also reproduce the result of Banerjee et al. (2013) and Cai, De Janvry and Sadoulet (2015): conditional on hearing the information, the behavior of a person’s full set of social ties is not significantly related to a person’s behavior. The general endorsement of one’s social ties appears not to matter for acting on new information. However, our unusually rich data allow us to dig deeper, revealing that nevertheless, networks still matter beyond informing.

We find that endorsement by ties *in general* is unrelated to changed behavior, but the endorsement by particular kinds of ties, especially those most intimate, is related. We also find those who act are more likely to have connections to certain others, especially the small number of people willing to act very early. Combined, these patterns are consistent with individuals seeking out certain trusted ties to vet information and judge its credibility. The network may well be responsible for behavior above and beyond providing the information that allows it, but studies that fail to track the information through the network would be unable to detect these effects even when they are present. Relatedly, without separate measures of hearing and acting, we would mask the surprising result that while people in central network positions had easier access to information, conditional on having heard, those

⁶An exception is Mobius, Phan and Szeidl (2015), which tracks the spread of information leading to participants’ guesses in a game the authors designed and implemented in a US university. Participants were all given some information, and were told that a majority of the information was correct. In such a setting, seeking out strangers was just as good as seeking out social contacts. In our setting, in which the information was far outside the realm of experience of the villagers and in which we did not establish common knowledge about the information beforehand, if social contacts play a role verifying and reinforcing information, our design should reveal this function.

in the least central positions were most likely to act on it.

3 Social Networks in Rural Uganda

We implemented our study in Abalang, a village in the Teso region of Uganda.⁷ The village has approximately 1,400 residents, is comprised predominantly of peasant farmers, and is largely ethnically homogeneous. Detailed demographic information can be found in Section 1 of the Supporting Information.

In our intervention, we seeded information with 7 households, one selected at random from each of seven equally-sized geographic subregions. For each selected “seed” household, a Ugandan enumerator who was not from the village personally visited and shared the information that starting in three days an event would be held at which all adults who take a survey would receive a large block of soap. The seeds were told that they were welcome to tell others, were asked a few basic questions about their household, and were given a sheet of paper containing the same information they were told about the event. Enumerators left the village and stayed away for the next three days.

On the fourth day, the survey event began. The event was hosted at a church just outside of Abalang. A total of 138 people from Abalang or its outskirts attended the event, which was held over three days. Attendance exhibited a pattern of “early adopters” leading a cascade: 11 attended the first day, 81 the second day, and 46 the third day.⁸

During the week following the event, enumerators conducted surveys door-to-door throughout Abalang. All adults in all households within view of the seed households were invited

⁷We implemented the same study simultaneously in a neighboring village of Mugana. News traveled poorly there and only one person from Mugana acted on it. Since the villages’ networks are distinct, we restrict analyses to Abalang.

⁸On the second day, more arrived at the event site than our team had time to survey; those who could not be surveyed were given a coupon that would allow them to attend on the third day. Only those with coupons from the second day were surveyed on the third day. For this reason, when we contrast early with late attenders, we pool the second and third day attenders.

to take the survey, and at least one adult in all other households in Abalang was invited to take the survey. In total, 328 individuals were surveyed.

Surveys administered at the event asked demographic information, general networks questions, and questions specific to learning and spreading word about the event, including from whom the respondent heard the information and whom they told. Surveys administered after the event asked demographic information, general networks questions, whether or not the respondent knew about the event and, if they claimed to know, who they heard from and whom they told. Additional details about the design can be found in Section 2 of the Supporting Information.

Our data are novel in two respects. First, in addition to a measure of who heard and acted on the information—a record of who attended the event—we also have a measure of who merely heard the information. Our surveys reveal that an additional 130 individuals throughout Abalang heard about the event but did not attend.⁹ By conditioning our analyses of attendance on hearing about the event, we can explore the role of networks in encouraging attendance beyond an information effect.

Second, because our event was held over multiple days, we can separate out those classically called “early adopters” in technology adoption studies. Here we refer to the 11 who attended on day one of the event as the “early attenders.” As we demonstrate below, social proximity to these individuals turns out to be the most robust determinant of attendance.

3.1 The Uncertainty of Seemingly Benign Information

One important feature of our design is the novelty of the information that we seeded and the action that it motivated. An event hosted by outsiders is unusual for residents of Abalang. Extensive daily, in-person briefings between one of the authors and the local enumeration

⁹Surveys asked those who said they heard what was given out at the event; all who claimed to have heard knew the correct answer, namely “soap”.

team confirmed that many villagers found the event to be far from their realm of experience.

Villagers found information about the event surprising. Uncertainty about brand new experiences contributes to two possible costs. First, travel to the event could be a substantial investment of time—respondents who did attend reported traveling between 2 and 180 minutes to reach the church, with a mean travel time of 50 minutes. Any doubt about whether the event would in fact be held or would in fact offer soap could render this travel time not worthwhile. Second, despite the team’s care to convey the information clearly, some villagers expressed concerns about the possible presence of witchcraft or devil worship at the event.

Although it would be easy to classify the news of “free soap” as benign, risk-free, and even boring, the reaction of the villagers suggests it was not regarded this way. If social networks play a special role in environments of uncertainty or perceived risk, we expect to detect this in Abalang.

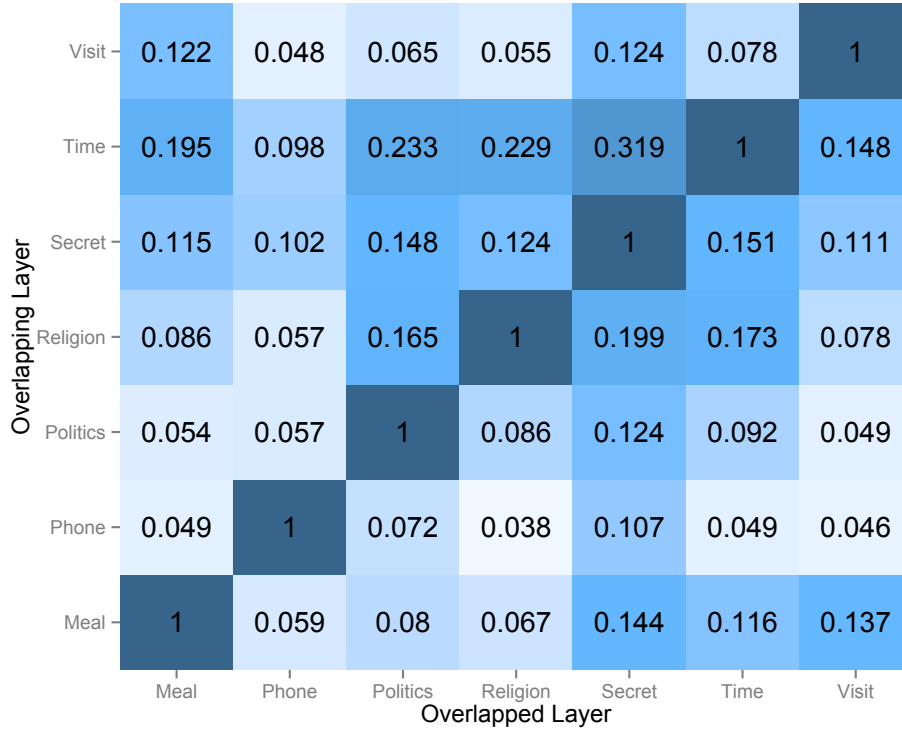
3.2 The Social Network of Abalang

Both the event and the post-event surveys collected information on seven dimensions of a social network. Respondents were asked to name up to five other individuals with whom they discuss politics, discuss religion, speak on the phone, share secrets, share meals, spend time, and whose homesteads they visit, for a possible maximum of 35 total names offered. The exact text of the network elicitation questions can be found in the Supporting Information, along with additional information about each network.

We selected these networks to maximize coverage of opportunities for word-of-mouth communication that are present in a rural village. Figure 1 shows the extent to which ties in one network are also present in each of the other networks. While there is considerable overlap, each network type contributes substantial information about a person’s social relationships.

Our main analyses use the union of these seven networks to capture an aggregate social

Figure 1: Heatmap of Overlap Between Networks



Note: we define overlap as the proportion of links of the overlapped layer (x-axis) contained in the overlapping layer (y-axis). Percentages are low in part because respondents could name anyone, including those outside of our sample. The same figure including only links that connect two respondents can be found in the Supporting Information.

network. In Section 4.3, we disaggregate the network to determine the role of each type of link.

3.3 Who Heard, Who Attended

Our data contain two outcomes: learning new information (hearing about the event), and acting on new information (attending the event). Our measure of who acted on the information is perfect: we observed attendance at the event. Our measure of who heard the information is perfect for those who attended (who could not have attended without hearing the information we seeded), and self-reported for those who heard but did not attend. In the

post-event survey, we asked those who claimed to have heard about the event to name the item given to attenders. All who claimed to have heard correctly named “soap.” While we cannot be sure that those who claimed they did not hear were being truthful, the great extent of people reporting that they did hear suggests that saying so was not generally perceived to be costly.

Table 1 reports raw comparisons of mean demographic and network information for those who neither heard nor attended, those who heard but did not attend, and those who attended. Relative to non-attenders, those who attended were significantly more likely to be female. No other demographic difference is statistically significant at the .05 level. Network measures, on the other hand, differ starkly between those who heard and those who did not, and between those who attended and those who heard but did not attend.

A few differences in network position stand out. First, those who heard but did not attend are much more eigenvector central in the network than those who did not hear, and are also significantly more eigenvector central than those who attended. Second, both those who heard and those who attended have a larger proportion of their network neighbors who heard. Finally, both those who heard and those who attended are significantly closer to early attenders than those who neither heard nor attended. Interestingly, distance to a seed— a person with first knowledge of the information and given an official information sheet— is less separating: those who heard are closer to a seed, but those who attended are not.

In both the event and post-event survey, in addition to inquiring about social network ties in general, we asked respondents to name people who told them and whom they told about the event. By comparing these initial sources of information about the event to the reported social network ties, we can identify the kinds of relationships that transmitted information about the event.

Figure 2 shows the proportion of ties of each type that were reported to have transmitted information about the event. All seven relationships were capable of transmitting informa-

Table 1: Base are those who neither heard nor attended. Hearers are those who heard but did not attend. Attenders are those who both heard and attended the event. P-values in Hearers and Attenders columns are from comparison to Base; in Hear v. Attend column from comparison of Hearers and Attenders.

	Base	Hearers	Attenders	Hear v. Attend
Prop. Female	0.59	0.73*	0.79**	
Age	39.43	39.88	37.57	
Prop. Married	0.79	0.91*	0.83	*
Prop. Catholic	0.63	0.63	0.68	
WallMat	0.10	0.13	0.14	
Educ	3.34	3.33	3.23	
Prop. Unemp	0.03	0.02	0.04	
Prop. PartTime	0.42	0.47	0.46	
Prop. FullTime	0.15	0.10	0.18	*
Prop. Retired	0.39	0.42	0.32	*
NumPeers	12.51	14.67*	15.80***	
AvgNumPeersPeers	9.90	11.24**	9.52	***
PropPeersHeard	0.73	0.88**	0.92***	**
PropPeersAttended	0.14	0.18*	0.18*	
DistSeed	2.14	1.85**	2.02	*
DistEarlyAttender	2.27	1.96**	1.75***	**
AvgDist	3.92	3.78**	3.80**	
Eigen. Centrality	0.05	0.09***	0.06	**
Number of People	59	130	138	

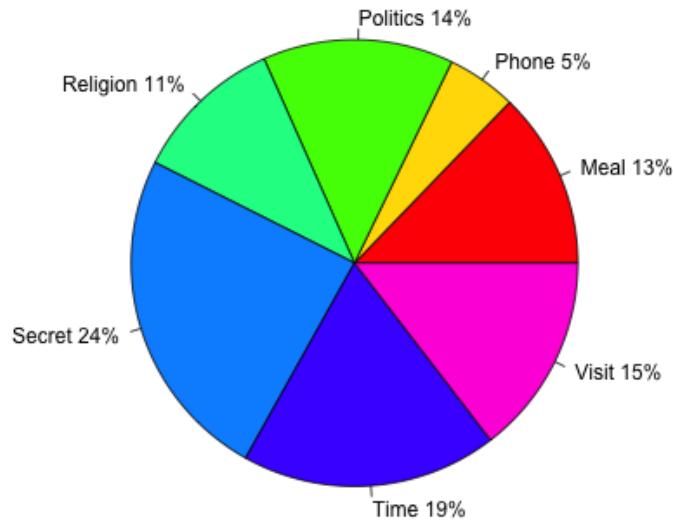
*p<0.1; **p<0.05; ***p<0.01

tion about the event. Ties that connote sharing secrets and spending time are particularly conducive—about a fourth and a fifth of secret and time ties in our data were reported to have also directly transmitted information about the event. The phone network is least conducive. While this may seem counterintuitive, respondents were also asked by what means they learned about the event. Exactly zero respondents learned about the event over the phone, and fewer than a third of respondents own a phone.¹⁰

Inquiring about a person’s source and target of information reveals only a small part of

¹⁰That five percent of ties in the phone network were also ties along which news of the event passed does not contradict the fact that no one learned over the phone. Even though five percent of phone ties match with informing ties, this does not mean the person was told *while talking on the phone*. This match simply means that the person who was the source or recipient is related socially by this function. As Figure 1 shows, some of these relationships appear in other networks as well.

Figure 2: Percentage of links of each type a respondent said were used to transmit information about the event.



the process by which people become informed and motivated. This question detects the most salient initial sources and targets of information, and confirms that social networks provide opportunities to pass information along. However, social networks also provide opportunities to discuss, verify, vet, establish social judgment about, and reinforce new information. Our direct questioning of initial sources and targets only scratches the surface of this process. In the next section, we use information about who heard, who attended, and respondents' social networks to dig deeper into the process by which social networks inform and motivate.

3.4 Overview of Methods

Our approach in the next section is to relate the network position of individuals to their propensity to hear, and to their propensity to attend conditional on hearing. The network position of one respondent in Abalang is, by definition, not independent of the network position of other respondents in Abalang. For this reason, we risk attributing too much precision to the results of regressions that rely on independence assumptions.¹¹ We take a number of measures to ensure that our results are statistically and substantively significant. When we use a parametric model, we employ a high threshold of statistical significance, and verify the result with a battery of alternate specifications. Some of our results rely on comparisons of the same model (other than the dependent variable) on the same network; comparisons of precision and magnitude hold the dependency structure constant. We further confirm results with a nonparametric test which takes the dependencies in the network as given. Because many network statistics are themselves correlated, we verify the importance of seemingly significant statistics with placebo tests which also allow us to compare results within the dependency structure of our particular data. These verifications and robustness checks can be found in the Supporting Information.

4 Results

We begin by assessing evidence for the contagion of information and of behavior through the social network. If straightforward contagion were at play as stipulated in the diffusion literature, the larger the proportion of a villager’s neighbors that exhibit the outcome, the more likely the villager is to exhibit the outcome as well.

¹¹Exponential Random Graph Models and Latent Space Models handle dependencies well when the dependent variable is a link, but our dependent variables of interest are at the level of the node.

4.1 Assessing Simple Contagion

We perform three tests for simple contagion. First, we use logistic regression to test the unconditional relationship between the proportion of one’s neighbors who exhibit the outcome (hear or attend) and whether one exhibits the outcome. If contagion were simply a matter of exposure to others in a social network, the effect of network neighbors should be apparent unconditionally— regardless of demographic or other network attributes. Second, we respecify the logistic regressions to condition on the number of network neighbors. Third, we use a non-parametric approach to examine the extent of clustering in the social network with respect to hearing and to attending. All three tests strongly support the simple contagion of hearing and cast strong doubt on the simple contagion of attending.

Table 2 reports the results of the first and second tests for the outcome of hearing about the event, displaying the marginal effects of each variable. Here we see strong evidence supporting simple contagion: the larger the proportion of a person’s network neighborhood that heard about the event (`PropPeersHear`), the more likely it is that the person heard about the event. A person whose full neighborhood heard is 68% more likely to have heard than a person for whom none of her neighbors heard. Controlling for the size of the neighborhood (`NumPeers`) changes the relationship little. This finding is consistent with a straightforward contagion process in which greater exposure to hearing about an event made a person more likely to “catch” the information.

Table 3 reports the results of the first and second tests for attendance for all who heard about the event. In contrast with the spread of information, the spread of behavior does not appear to exhibit simple contagion. It is *not* the case that the more of one’s network neighbors who attended the event, the more likely one is to have attended.¹²

¹²Because the tests for attending include only those who heard about the event and so were eligible to attend, the sample size is smaller than the tests for hearing. It is conceivable that the difference in hearing and attending is an artifact of a difference in power and not a difference in the process that spread the outcome. To assess this, we take random samples of size 268— the number that heard about the event— from the data used for the specifications in Table 2 and rerun the analyses on the smaller samples. Doing

Table 2: Relationship between proportion of peers who heard and hearing about the event

	P(Hear About the Event)	
	(1)	(2)
PropPeersHear	0.680*** (0.117)	0.647*** (0.120)
NumPeers		0.006** (0.003)
Adj. R-Squared	0.141	0.156
Observations	326	326

*p<0.1; **p<0.05; ***p<0.01

Note: Reported values are the marginal effects for the average observation. Network statistics calculated for undirected aggregate social network.

Table 3: Relationship between proportion of peers who attended and attending the event

	P(Attend the Event)	
	(1)	(2)
PropPeersAttend	0.083 (0.224)	0.086 (0.225)
NumPeers		0.005 (0.005)
Adj. R-Squared	0	0.005
Observations	268	268

*p<0.1; **p<0.05; ***p<0.01

Note: Reported values are the marginal effects for the average observation. Data include all respondents who heard about the event. Network statistics calculated for undirected aggregate social network.

The Supporting Information contains additional analyses that show that these results are not sensitive to including those who did not hear in the attendance regressions, to using the existence of a tie to one who heard or attended instead of the proportion of peers who heard or attended, or to including various sets of control variables.

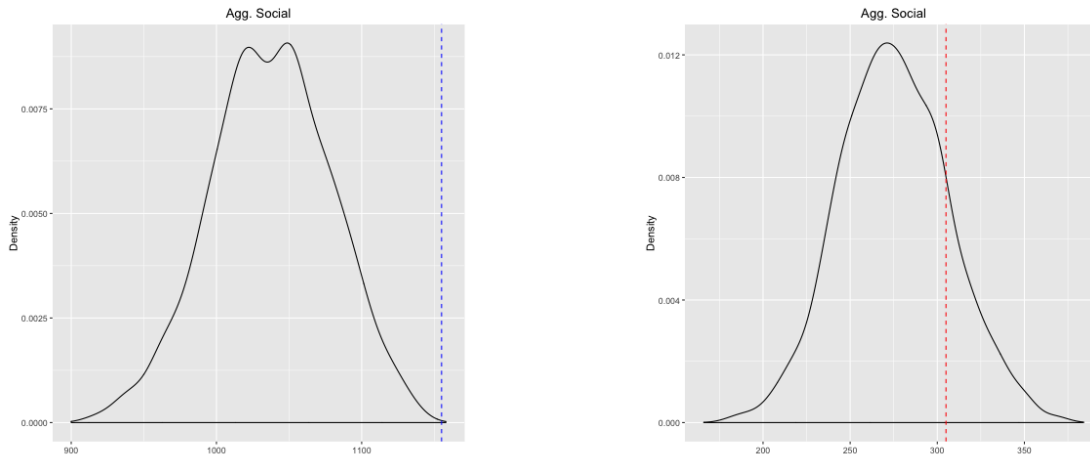
As a third test for contagion, we turn to a non-parametric approach that naturally accounts for the dependencies in network data.¹³ If an outcome spread through a network from node to node via a contagion process, then those nodes exhibiting the outcome should be more connected to one another in the network than a randomly selected set of nodes of the same size from the same network would be. By randomly sampling a set of nodes from the network of size 268– the number that heard about the event– and comparing the density of ties among that set to the density observed among the 268 whom we know heard, we find that a density as high as the observed value would only be expected by chance less than .1% of the time; a statistical significance at the .001 level. Performing the same on sets of nodes of size 138– the number of attenders– and comparing the density reveals that we would observe a density at least as high as the density among the actual attenders 18% of the time, falling short of conventional thresholds for statistical significance. Figure 3 shows the sampling distributions generated by this test. The non-parametric approach confirms that hearing about the event likely followed a contagion process while attending the event did not.

If the social network played a role in encouraging attendance, it was not via a simple contagion process. We now turn to a deeper investigation of the determinants of attendance.

so 10,000 times produces estimates significant at the .05 level in every case. The estimate in the simple regression ranges from .47 to .98, and the estimate in the conditional regression ranges from .44 to .95 in the shrunken samples. The smaller sample size is not masking a contagion effect for attendance– if the data used to estimate the models for hearing were as small, the results would retain their statistical significance.

¹³Of course, since the two outcomes are compared for the same network, the extent of underlying dependence among network neighborhoods is the same. Consequently, the relative statistical precision in the parametric models is still meaningful.

Figure 3: Sampling distribution of density among groups of people the same size as the group of all hearers (268, left), and the same size as the group of all attenders (138, right). The vertical line indicates the density of the observed group of hearers and attenders, respectively. Hearers cluster in the network consistent with a simple contagion process; attenders do not.




4.2 Social Networks as Sources of Credibility

Next we examine evidence that the network functions as a source of credibility. In this conceptualization, people can use their ties to vet or verify information. Consider two types of people who could be useful in establishing the credibility of this intervention’s information. One is the “seeds.” These seven individuals were personally visited by enumerators and given an information sheet about the event. Their information was official and they received it first. The second is the “early attenders.” These eleven individuals chose to attend the event on the first day.¹⁴ We calculate the network distance separating a person from their closest seed: if a person has a seed as a network neighbor, DistSeed is 1. If a person has no seed as a neighbor, but one of her neighbors has a seed as a neighbor, DistSeed is 2, and so on. We calculate the same for distance to an early attender, DistEarlyAttend . The Supporting Information Section 6 shows that a person directly connected to an early attender is at least 11.5% more likely to attend than a person whose closest connection to an early attender is

¹⁴Only five of the eleven early attenders were seeds. The correlation between people’s shortest path to a seed (DistSeed) and shortest path to an early attender (DistEarlyAttend) is .36. The maximum value for both is 4.

Figure 4: Proportion of those in the sample at different distances to an early attender who heard and attended. Fewer of those farther from an early attender attended.



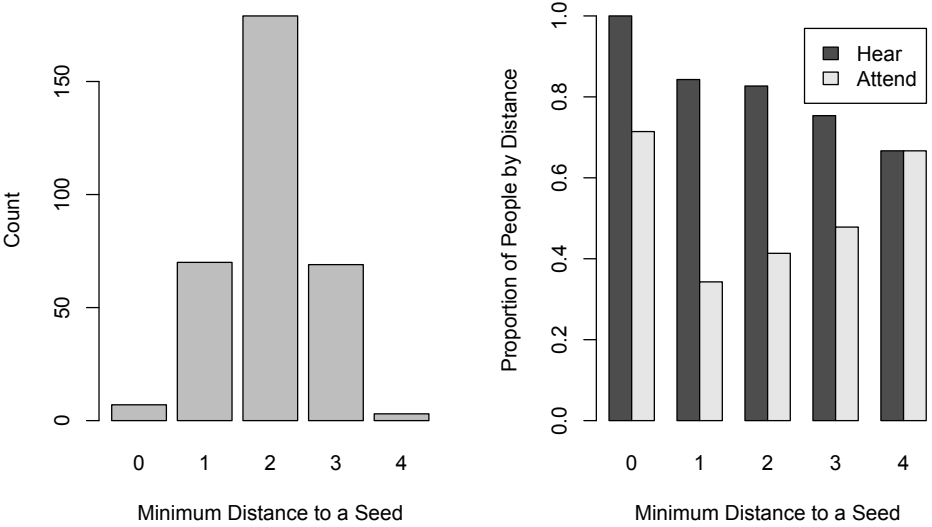
a friend-of-a-friend. A person as close as possible to an early attender is 34% more likely to attend than a person as far as possible in this network from an early attender.

Figure 4 shows the extent of hearing and of attendance at an increasing distance to an early attender. At network distances farther out from the early attenders, fewer and fewer people heard or attended. For contrast, Figure 5 shows the same by distance to a seed instead of to an early attender. As the analyses in SI Section 6 confirm, distance to a seed is not related to attendance in the same way. While fewer and fewer people heard as the distance from a seed increases, more and more people attended.¹⁵

On the one hand, the seeds were privy to the official version, and possessed a means of corroborating their message: showing the information sheet. On the other hand, given the novelty of an event like this one in an otherwise rural, remote village, despite the enumerators' assurances that soap would be plentiful and that the purpose was benign, there may have

¹⁵In the Supporting Information, we show that the importance of proximity to an early attender is robust to a battery of alternate specifications, and is supported by the results of a set of placebo tests using distance to 11 respondents selected at random using different sets of criteria.

Figure 5: Proportion of those in the sample at different distances to a seed who heard and attended. *More* of those farther from the seeds attended.



been doubt about whether soap would in fact be available, or even about the true purpose or ulterior motive of the hosts. In an environment of uncertainty or possible perceived risk, social proximity to one who tests out attending, reports the low risk, and credibly demonstrates the existence of soap would reasonably be motivating.¹⁶ SI Section 7 compares the demographic characteristics and network positions of the eleven early attenders to the 127 later attenders. The two groups do not differ significantly in demographic attributes, but they are distinguished by their network positions. The early attenders occupy a close-knit community within the network that is near the seeds.

¹⁶One enumerator reported that, upon arrival at the event, an attender commented: “oh, it’s just a white lady.” Early attenders willing to accept the risk could report back to their social contacts about the safety of the event and credibly demonstrate the existence of soap.

4.3 Disaggregating Network Type

Our data that measure the social network along seven dimensions can shed light onto which types of relationships are conducive to motivating attendance. Additional information about each dimension can be found in the Supporting Information.

Table 4 separates the social network into its seven constituent layers. Each of these layers is treated as a separate network. Now the proportion of peers that attended is calculated with respect to the number of peers in this network only (as opposed to the union of the seven). This effectively measures the extent of contagion along each link type.

Table 4: Relationship between attending the event and neighbors in each network type who attended

	P(Attend the Event)						
	Time	Phone	Politics	Religion	Meal	Visit	Secret
PropPeersAttend	-0.172 (0.157)	0.315* (0.167)	0.229 (0.142)	0.131 (0.137)	0.278** (0.126)	0.252** (0.117)	0.009 (0.111)
Adj. R-Squared	0.003	0.014	0.012	0.003	0.017	0.015	0
Observations	263	175	183	251	225	229	226

*p<0.1; **p<0.05; ***p<0.01

Note: Reported values represent the marginal effects for the average observation. Network statistics calculated for undirected network comprised of a single tie type.

Although there is no evidence for contagion of attendance in the aggregate social network, here we see that contagion may be present with respect to two link types. In both the sharing meals and the visiting homes networks, having more network neighbors who attended is associated with attending. That is, the more people with whom a person shares meals or visits that attended, the more likely a person is to attend. The magnitude is large as well; a person with all of her meal partners attending is 28% more likely to attend than a person with none of her meal partners attending, and a person with all of her visit partners attending is 25% more likely to attend than a person with none of the people she exchanges

visits attending.¹⁷

One interpretation of these results is that these quite intimate relationships serve to vet news of the day. Through discussions (perhaps at meal time or during a visit), information is assessed and right actions are established. These intimate meal-sharing ties are used to vet new actions.

While this interpretation cannot be fully confirmed with our data, we examine the extent to which individuals who attended cluster in the visits and the meals network, using the same non-parametric density simulations used to test for contagion in Section 4.1. By randomly drawing a set of nodes from each network the size of the set of attenders in that network and measuring the extent to which they are connected, we can compare how interconnected those who attended are in terms of that type of link relative to its sampling distribution. The Supporting Information contains the results. Only in two networks is the extent of interconnectedness among attenders higher than would be expected by chance: the meals and the visits network. Those who attended are significantly more interconnected in terms of sharing meals and visiting one another than would be expected by chance, further supporting the importance of these intimate ties.

4.4 Digging Deeper into Attendance

To fill out the story of who hears and who attends, we add a battery of demographic and network controls. Table 5 confirms the consistent story for hearing about the event: hearing is robustly contagious, so that the more network neighbors who heard, the more likely one is to have heard. Moreover, being eigenvector central, female, socially proximate to an early attender, and wealthier (estimated by the quality of wall material) are positively associated

¹⁷The marginal effects for the phone and politics networks are substantively large, though both are imprecisely estimated. Their size and precision change dramatically with different sets of control variables. The marginal effects in the meal and visits networks are more stable across specifications, and confirmed by the nonparametric tests reported below. See the Supporting Information, Section 6.

with hearing.

Table 6 shows the same specification for attending. Even conditional on own demographic, neighbors' demographic, and network characteristics, attending is still not contagious in the full social network— the proportion of one's peers who attend is unrelated to one's own attendance. The story continues to be one of social proximity to an early attender. For the average villager in our sample, having a path to an early attender that is one link shorter— moving from having a friend-of-a-friend to having a friend who attended on day 1, say— is associated with being 12% to 15% more likely to attend. Next we zoom in on the negative relationship between eigenvector centrality and attendance.

4.5 Role of Network Centrality

Network theory and existing empirical studies hold that greater network centrality should be associated with greater access to information spreading through a network (Kempe, Kleinberg and Tardos, 2003; Borgatti, 2005; Ballester, Calvó-Armengol and Zenou, 2006; Banerjee et al., 2013, 2014). Indeed, we find that greater network centrality is in general positively associated with hearing information about the event. The closer a person is to other people (low AvgDist), and the more a person is connected to other highly connected people (high Eigen), the more likely a person is to hear information spreading through her network.

It would be easy to conclude that the relationship between network centrality and acting on the information should be positive as well. However, among our sample of those who heard, centrality measures are consistently *negatively* related to acting on the information. Given that a person has heard the information, being more central is associated with a lower likelihood of attending.

Table 7 explores this relationship by stripping out other controls and assessing the connection between network position and attendance. While the size of one's neighborhood— one's “degree centrality”— is positively associated with attending (though with varying precision),

Table 5: Relationship between peers who heard and hearing, conditional on other network attributes and ego and peer demographic characteristics

	P(Heard about the Event)	
	(1)	(2)
PropPeersHear	0.502*** (0.116)	0.499*** (0.123)
NumPeers	0.002 (0.003)	0.002 (0.003)
Eigen	0.720*** (0.275)	0.741** (0.290)
DistEarlyAttend	-0.065** (0.030)	-0.061** (0.029)
Female	0.143** (0.061)	0.170** (0.068)
Age	0.0002 (0.001)	-0.001 (0.001)
Catholic	0.033 (0.039)	0.014 (0.041)
Educ	0.003 (0.013)	0.001 (0.013)
Married	0.052 (0.068)	0.043 (0.065)
WallMat	0.077** (0.036)	0.082** (0.034)
FemalePeers		-0.043 (0.066)
AgePeers		0.001 (0.002)
CatholicPeers		0.010 (0.064)
EducPeers		-0.033 (0.023)
MarriedPeers		0.069 (0.117)
WallMatPeers		-0.143 (0.095)
Adj. R-Squared	0.23	0.258
Observations	310	306

*p<0.1; **p<0.05; ***p<0.01

Note: Reported values represent the marginal effects for the average observation. Network statistics calculated for undirected aggregate social network.

Table 6: Relationship between peers who attended and attendance, conditional on other network attributes and ego and peer demographic characteristics

	P(Attend the Event)	
	(1)	(2)
PropPeersAttend	-0.259 (0.294)	-0.078 (0.322)
NumPeers	0.010* (0.006)	0.012** (0.006)
Eigen	-1.321** (0.608)	-1.525** (0.680)
DistEarlyAttend	-0.123** (0.053)	-0.154*** (0.056)
Female	-0.009 (0.089)	0.098 (0.105)
Age	-0.002 (0.002)	-0.003 (0.002)
Catholic	0.016 (0.071)	-0.038 (0.078)
Educ	-0.037 (0.025)	-0.038 (0.027)
Married	-0.159* (0.093)	-0.178* (0.098)
WallMat	0.055 (0.092)	0.071 (0.102)
FemalePeers		-0.401* (0.207)
AgePeers		0.001 (0.005)
CatholicPeers		0.237* (0.135)
EducPeers		-0.021 (0.064)
MarriedPeers		0.135 (0.246)
WallMatPeers		0.034 (0.232)
Adj. R-Squared	0.069	0.104
Observations	256	252

*p<0.1; **p<0.05; ***p<0.01

Note: Reported values represent the marginal effects for the average observation. Network statistics calculated for undirected aggregate social network. Data include all respondents who heard about the event.

the other measures all relate negatively. The more central a person is, measured in terms of the size of her neighbors' neighborhoods, her average distance to everyone else in the network, and her eigenvector centrality, the less likely she is to attend (though again with varying precision).¹⁸

Table 7: Relationship between network centrality and attending the event

	P(Attend the Event)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
NumPeers	0.005 (0.005)				0.015** (0.007)	0.013** (0.005)	0.005 (0.004)
NumPeersPeers		-0.033*** (0.009)					-0.033*** (0.009)
AvgDist			0.064 (0.097)		0.325** (0.148)		
Eigen				-0.857** (0.386)		-1.428** (0.557)	
Adj. R-Squared	0.004	0.04	0.001	0.016	0.018	0.036	0.044
Observations	268	268	268	268	268	268	268

*p<0.1; **p<0.05; ***p<0.01

Note: Reported values represent the marginal effects for the average observation. Network statistics calculated for undirected aggregate social network. Data include all respondents who heard about the event.

Table 8 presents the raw comparison between the twenty villagers in the sample with the highest eigenvector centrality and the twenty with the lowest. The most eigenvector central in the sample are less female and less catholic (lending credence to the specifications of the last section that added these demographic characteristics as controls). Of course, due to the definition of eigenvector centrality, they also mechanically have larger neighborhoods, have neighbors with larger neighborhoods, and are closer to everyone on average. Consequently they also have lower distances to both the seeds and the early attenders.

¹⁸SI Section 9 provides evidence supporting the robustness of the centrality result. The biggest concern is that the relationship between centrality and attendance is an artifact of sampling all who attended, boosting their centrality. We show that our results are robust to a number of specifications that account for sampling, and that in fact our method of sampling should render our results on centrality conservative.

Table 8: Comparison of the twenty villagers with the highest eigenvector centrality and the twenty with the lowest eigenvector centrality in the social network.

	Bottom Eigen Central	Top Eigen Central
Female	0.75	0.30**
Age	42.58	39.89
Catholic	0.95	0.32***
WallMat	0.11	0.10
Married	0.79	0.89
Educ	3.11	3.70
Unemp	0.00	0.10
PartTime	0.50	0.25
FullTime	0.11	0.10
Retired	0.39	0.55
NumPeers	7.20	22.70***
NumPeersPeers	6.71	13.05***
DistSeed	2.75	1.35***
DistEarlyAttend	2.75	1.80***
AvgDist	4.49	3.48***
PropPeersHear	0.90	0.90
PropPeersAttend	0.18	0.17
Eigen	0.004	0.349***
Hear	0.80	0.90
Attend	0.55	0.35
AttendGivenHeard	.69	.39**
Number	20	20

The notable comparisons are in terms of hearing and attending. While more of the top eigenvector-central villagers heard about the event, a significantly smaller proportion of those who heard attended.

This raw comparison, combined with the controlled comparisons above, imply the following role of centrality: central network positions are indeed privileged in terms of access to information. When information is flowing through a network, the most central are most likely to hear it. However, among those who receive the information, it is the *least* central who are most likely to act on it. This latter finding would be masked by studies that only measure acting on information; as a group, those who act on information are more central

than those who do not. However, this is due to their access to information. Conditional on receiving information, it is the least central who are most likely to act on it.

There are many reasons why greater network centrality may discourage those who receive information to act on it. It could be that the most central have less need for soap, or have a higher implied opportunity cost of attending the event than the less central. Table 8 shows that the twenty most and least central are similar in terms of the quality of their house’s wall material (a proxy for wealth) and employment status (which speaks to opportunity cost), which cast some doubt on these explanations. It could also be that the most central are interested in acquiring the soap, but are better able to send someone else to attend and acquire the soap on their behalf than the less central. This would suggest a positive relationship between centrality and the proportion of peers who attended. Table 8 shows this is not the case, if anything the proportion of peers who attended is slightly smaller for the more central. Alternatively, if network neighbors serve as channels of peer pressure and social judgment, those with more influential neighbors— e.g. neighbors with large neighborhoods or that have high eigenvector centrality— may face greater novelty-resisting pressure. In the presence of new opportunities, those less central may be less encumbered by expectations of their peers. While our data cannot fully adjudicate among explanations for the negative relationship between network centrality and attendance among those who heard, we encourage future researchers to explore this finding that has potentially large consequences for interventions aiming to motivate actions with novel information.

5 Discussion and Conclusion

By randomly seeding novel information with individuals in Abalang, Uganda; detecting who learned it and who acted on it; and measuring personal information sources and social networks, we offer a rare, direct examination of the process by which social networks help

inform and motivate in a rural, developing country context.

We show that information spreads through the network along a myriad of social ties, and that in general, the more of one’s social contacts who hear the information, the more likely one is to hear the information. The pattern, verified by a number of parametric and non-parametric specifications, is consistent with a process of simple contagion assumed by many theories of information diffusion.

On the other hand, the process by which people become motivated to act on new information is more complicated. Notably, behavior— in this case, attending an event— does not spread so easily through the network, and does not follow a pattern of simple contagion. It is *not* the case that the more of a person’s social contacts attend, the more likely a person is to attend.

Instead, we find that while the proportion of one’s peers who attend is unrelated to attending, social proximity to someone who was willing to attend the event before most others— the “early adopters” — is strongly related to attending. In fact, social proximity to these informal sources that “tested” the event is more important than social proximity to people who were the initial recipients of the information and the accompanying official information sheet.

Furthermore, attenders are densely connected to one another by strong, intimate ties— these people pay visits to one another’s homesteads and share meals together much more than other groups of the same size. The larger the proportion of people with whom one shares a meal or homestead visits attended the event, the more likely a person is to have attended. Intimate connections to others who attended may serve to encourage attending.

Network theory suggests that central positions within a network offer greater access to information. We find support for this proposition, as network centrality is associated with hearing the information. However, we find a surprising relationship between network centrality and *acting* on the new information. While the most central are more informed,

among those who are informed, it is the less central who act on it.

These findings are consistent with social networks serving an important role beyond merely informing. Network ties serve as channels of brand new information, but they appear to do more than this as well. Social contacts can also serve to reinforce, verify, and pass judgment on topics like whether new information is worth acting on. These ties give a person access to the information that early attenders bring back from their test of the event, and provide a forum for deliberating future actions. When one's most trusted contacts endorse an action, one is more likely to take the action as well. The potential for negative judgment may also deter new, risky, or unconventional behavior.

Our results suggest a number of important avenues for future research. The core group of early adopters appear to have been important to the eventual high attendance of the event. This pattern is consistent with critical mass theory and empirical studies of cascades in technology adoption, though little is known about how to encourage this initial group. This study hints that a tight-knit group may be a fruitful target, establishing a local pocket of common knowledge that makes attending less costly within the group, and encourages attendance by others outside the initial group later on. Also, this group of early adopters was self-selected; whether interventions which require early participation by a few are equally effective as those relying on voluntary early participation remains to be seen. Moreover, other work has found that information may flow differently within and across ethnic groups. While Abalang is ethnically homogeneous, understanding how these determinants of motivation interact with ethnic diversity will be an important next step.

Our results also highlight the importance of being sensitive to a local context before implementing interventions aimed at changing behavior. While the behavior we aimed to encourage – attending a survey event to receive some soap– seemed benign *ex ante*, to those in an isolated village, it was far enough from local experience to be interpreted as uncertain and even risky. Interventions that encourage seemingly low-cost, low-risk, objectively beneficial

actions may still be interpreted differently by the recipients. It may be tempting to suppose that all that will be required to change behavior is to spread true, sensible information. However, if the recipients of the information *regard* the action encouraged to be risky, even merely socially risky, then the spread of behavior will likely take a very different form from the spread of information.

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.
- Apicella, Coren L, Frank W Marlowe, James H Fowler and Nicholas A Christakis. 2012. “Social networks and cooperation in hunter-gatherers.” *Nature* 481(7382):497–501.
- Ballester, Coralia, Antoni Calvó-Armengol and Yves Zenou. 2006. “Who’s who in networks. wanted: the key player.” *Econometrica* 74(5):1403–1417.
- Bandiera, Oriana and Imran Rasul. 2006. “Social Networks and Technology Adoption in Northern Mozambique.” *The Economic Journal* 116(514):869–902.
- Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo and Matthew O Jackson. 2013. “The Diffusion of Microfinance.” *Science* 341(6144).
- Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo and Matthew O Jackson. 2014. Gossip: Identifying central individuals in a social network. Technical report National Bureau of Economic Research.
URL: <https://goo.gl/gZSBfY>
- Barr, Abigail, Jean Ensminger and Jeffrey C. Johnson. 2010. Social Networks and Trust in Cross-Cultural Experiments. In *Whom Can We Trust? How Groups, Networks and Institutions Make Trust Possible*. Russell Sage Foundation Press chapter 3, pp. 65–90.
- Bond, Robert M, Jaime E Settle, Christopher J Fariss, Jason J Jones and James H Fowler. 2017. “Social endorsement cues and political participation.” *Political Communication* 34(2):261–281.

- Borgatti, Stephen P. 2005. "Centrality and network flow." *Social networks* 27(1):55–71.
- Cai, Jing, Alain De Janvry and Elisabeth Sadoulet. 2015. "Social Networks and the Decision to Insure." *American Economic Journal: Applied Economics* 7(2):81–108.
- Centola, Damon M. 2013. "Homophily, Networks, and Critical Mass: Solving the Start-up Problem in Large Group Collective Action." *Rationality and society* 25(1):3–40.
- Centola, Damon and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American journal of Sociology* 113(3):702–734.
- Chandra, Kanchan. 2004. *Why ethnic parties succeed: patronage and ethnic head counts in India*. Cambridge, UK: Cambridge University Press.
- Chiang, Yen-Sheng. 2007. "Birds of Moderately Different Feathers: Bandwagon Dynamics and the Threshold Heterogeneity of Network Neighbors." *Journal of Mathematical Sociology* 31(1):47–69.
- Chong, Alberto, Ana De La O, Dean Karlan and Leonard Wantchekon. 2015. "Does Corruption Information Inspire the Fight or Quash the Hope? A Field Experiment in Mexico on Voter Turnout, Choice, and Party Identification." *Journal of Politics* 77(1):55–71.
- Chwe, Michael Suk-Young. 1999. "Structure and Strategy in Collective Action." *American journal of sociology* 105(1):128–156.
- Chwe, M.S.Y. 2000. "Communication and Coordination in Social Networks." *Review of Economic Studies* 67(1):1–16.
- Conley, T.G. and C.R. Udry. 2010. "Learning about a New Technology: Pineapple in Ghana." *The American Economic Review* 100(1):35–69.

- 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.
- Dodds, Peter Sheridan and Duncan J Watts. 2004. “Universal Behavior in a Generalized Model of Contagion.” *Physical review letters* 92(21):218701.
- Eubank, Nick, Guy Grossman, Melina Platas and Jonathan Rodden. 2017. “Social Networks and Turnout in Low Salience Elections: Evidence from Uganda.” Working Paper.
URL: <https://goo.gl/Ym3Npo>
- Fearon, J.D. and D.D. Laitin. 1996. “Explaining Interethnic Cooperation.” *The American Political Science Review* 90(4):715–735.
- Ferrali, Romain, Guy Grossman, Melina Platas and Jonathan Rodden. 2018. “Peer Effects and Externalities in Technology Adoption: Evidence from Community Reporting in Uganda.” SSRN.
URL: <https://goo.gl/NcGSvv>
- Gerber, Alan S. and Donald P. Green. 2000. “The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment.” *The American Political Science Review* 94(3):653–663.
- Gould, Roger V. 1993. “Collective Action and Network Structure.” *American Sociological Review* pp. 182–196.
- Habyarimana, James, Macartan Humphreys, Daniel N. Posner and Jeremy M. Weinstein. 2009. *Coethnicity: diversity and the dilemmas of collective action*. New York: Russell Sage Foundation.

- Humphreys, Macartan and Jeremy M. Weinstein. 2012. "Policing Politicians: Citizen Empowerment and Political Accountability in Uganda." Unpublished Manuscript.
URL: <https://goo.gl/kiJyLG>
- Jackson, Matthew O and Brian W Rogers. 2007. "Relating Network Structure to Diffusion Properties through Stochastic Dominance." *The BE Journal of Theoretical Economics* 7(1).
- Kempe, David, Jon Kleinberg and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM pp. 137–146.
- Larson, Jennifer M. 2017. "Networks and Interethnic Cooperation." *Journal of Politics* 79(2):546–559.
- Larson, Jennifer M and Janet I Lewis. 2017. "Ethnic Networks." *American Journal of Political Science* 61(2):350–364.
- Lewis, Janet I. 2017. "How Does Ethnic Rebellion Start?" *Comparative Political Studies* 50(10):1420–1450.
- Lieberman, Evan S., Daniel N. Posner and Lily L. Tsai. 2014. "Does Information Lead to More Active Citizenship? Evidence from an Information Intervention in Rural Kenya." *World Development* 60:69–83.
- Marwell, Gerald, Pamela E Oliver and Ralph Prahl. 1988. "Social Networks and Collective Action: A Theory of the Critical Mass III." *American Journal of Sociology* pp. 502–534.
- Miguel, E. and M.K. Gugerty. 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics* 89(11-12):2325–2368.

- Mobius, Markus, Tuan Phan and Adam Szeidl. 2015. "Treasure Hunt: Social Learning in the Field." NBER Working Paper.
URL: <https://goo.gl/9Lhbtm>
- Munshi, Kaivan. 2004. "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution." *Journal of development Economics* 73(1):185–213.
- Newman, Mark EJ. 2000. "Models of the Small World." *Journal of Statistical Physics* 101(3-4):819–841.
- Paluck, Elizabeth Levy, Hana Shepherd and Peter M Aronow. 2016. "Changing climates of conflict: A social network experiment in 56 schools." *Proceedings of the National Academy of Sciences* 113(3):566–571.
- Przeworski, Adam, Susan C. Stokes and Bernard Manin. 1999. *Democracy, accountability, and representation*. Cambridge, U.K.: Cambridge University Press.
- Siegel, David A. 2009. "Social networks and collective action." *American Journal of Political Science* 53(1):122–138.
- Sinclair, Betsy, Margaret McConnell and Melissa R Michelson. 2013. "Local canvassing: The efficacy of grassroots voter mobilization." *Political Communication* 30(1):42–57.
- Valente, Thomas W. 1996. "Social Network Thresholds in the Diffusion of Innovations." *Social networks* 18(1):69–89.
- Varshney, A. 2003. *Ethnic Conflict and Civic Life: Hindus and Muslims in India*. New Haven: Yale Univ Pr.