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

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

April 16, 2021

Abstract

From public health to political campaigns, numerous attempts to encourage behavior begin with the spread of information. Of course, seeding new information does not guarantee action, especially when it's difficult for receivers to verify this information. We use a novel design that introduced valuable, actionable information in rural Uganda and reveals the intermediate process that led many in the village to hear the information, but only some to act on it. We find that the seeded information spread easily through word-of-mouth via a simple contagion process. However, acting on the information spread less easily; this process relied instead on endogenously created social information that served to vet, verify and pass judgment. Our results highlight an important wedge between information that a policy intervention can best control and the behavior that ultimately results.

9,863 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 for generously giving their time. We received support for this project from Harvard's Institute for Quantitative Social Science.

†Associate Professor of Political Science, Vanderbilt University, jennifer.larson@vanderbilt.edu.

‡Assistant Professor of Political Science, George Washington University.

§Postdoctoral Fellow, Data Science Institute, Vanderbilt University, and Instituto de Estudios Superiores de Administración, Caracas, Venezuela.

1 Introduction

For NGOs, electoral candidates, and rebel leaders alike, changing people's behavior by providing them with new information is a fundamental goal. Doing so is especially challenging, however, in situations where the information is difficult for receivers to verify and the behavior's potential returns are uncertain. For example, public health officials seek to increase uptake in healthy practices in rural villages through information campaigns countering unfounded rumors on potential side-effects. Opposition candidates in authoritarian regimes try to influence vote choice with information about their character and policy platforms. Rebel group leaders aim to attract supporters by persuading villagers that they are likely to succeed. Non-governmental organizations encourage citizens to speak up when witnessing corruption on the part of public officials.

In sum, outcomes ranging from public health crises to violent conflict depend not only on new information spreading widely, but also on people opting to act on difficult-to-verify information and in the face of uncertain outcomes. This two-step process is foundational to core theories in political science[1] and to the success of myriad policy interventions, and the two steps can operate differently. For example, someone may hear about a politician's alleged corruption while running errands, but mere market chatter may not be enough to motivate her to turn out to vote for the opposition. Likewise, it is one thing for a villager to hear that international health workers have arrived 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

[1]For example, retrospective voting theories indicate 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 is thought to depend on their ability to seed favorable information about themselves among the local civilian population (Larson and Lewis, 2018; Lewis, 2020). 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).

take them to the quarantine tent. This gap between information provision and behavioral change has been identified in numerous recent studies (e.g. Lieberman, Posner and Tsai, 2014; Chong et al., 2015; Dunning et al., 2019), and raises important questions about when and how it can be closed.

Recent studies suggest that part of the answer depends on social networks. It is well known that social networks spread information (Banerjee et al., 2013; Mobius, Phan and Szeidl, 2015; Larson and Lewis, 2017) and that peers can influence behavior (Sinclair, McConnell and Michelson, 2013; Bond et al., 2017). Less well known is *how* exactly networks convert information into action. This article joins a small, new literature in political science that uses original network data collected in the field to disentangle the processes by which social networks disseminate information versus how they motivate action (Ferrali et al., 2019; Eubank et al., Forthcoming; Nathan and Atwell, 2020).[2]

Specifically, we designed a study that zooms in on the social processes by which people come to hear new, hard-to-verify information and decide whether to act on it. We implemented our study in a rural village in eastern Uganda. We provided seven randomly chosen households 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. Of the roughly 1,400 residents of the village we studied 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,

[2]More broadly, our study joins a burgeoning literature that 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). Although social networks are thought to be particularly valuable in such contexts– and for this reason we see our study's setting of rural Africa as a virtue– they are also important for informing and deliberating in developed country settings (e.g. Gerber and Green, 2000).

surveys were conducted throughout the village to complete our mapping of social networks and to measure the reach of information. We thus 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 find that the seeded information spread through the community exclusively through face-to-face interactions along social ties of various types. Straightforwardly, and consistent with contagion models of information diffusion through networks (Chwe, 1999; Newman, 2000; Jackson and Rogers, 2007; Centola and Macy, 2007; Centola, 2013), a person who had more social ties that heard about the event was more likely to hear about it herself. However, the social process that motivated some to act on this information was more complicated. Having more social contacts who attended the event is robustly not predictive of a person's attendance, even conditional on hearing about the event.[3] In the language of Young (2009), simple contagion does not explain behavior in this village, so we probe the possibility of social influence and social learning as well.

To interpret our results, it is important to consider how villagers understood the information we shared. The behavior we elicited was perceived as uncertain or potentially sensitive by our respondents. This may have been due to the difficult-to-verify nature of our information, the relatively unusual way we communicated it, or the unusual nature of the event itself. We designed the study with the aim of ethically introducing this kind of uncertainty, since we sought to approximate substantive contexts sketched above, for which the information spread - e.g. about the promises of a new health practice or a new rebellion - is inherently difficult for villagers to confirm.[4] While we obtained all required approvals for

[3]This finding is consistent with Banerjee et al. (2013); Cai, De Janvry and Sadoulet (2015).

[4]We ruled out introducing new political or security information due to ethical considerations. We ruled out spreading health-related information due to practical considerations of ensuring that our information was useful but *new*.

our study, including a district-level approval, we intentionally did not coordinate with village officials prior to the event, which precluded villagers from confirming with local leaders the existence of the event or the reputations of its organizers.[5] Our intervention therefore relates most directly to cases of new information spreading that is difficult to verify; this could be villagers initially hearing novel news from a resident of another village, or from a not-fully-trusted outside entity like a new NGO or politician, or from a trusted local official but about a new behavior with uncertain returns. Further, although a soap event may seem mundane, it was unusual in this context – there had not been a similar event in this area in recent years, according to local officials.[6] Consequently, 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, or that attending would be viewed favorably by fellow villagers. In such a setting, we might expect people to turn to their social contacts before deciding what to do.

Indeed, we find evidence consistent with social learning (Young, 2009); a greater propensity for villagers to act on the information after observing early attenders of the event who could confirm its validity and utility. The eleven villagers who attended on the first day saw the event with their own eyes, experienced a low risk event, and received the promised block of soap. Social proximity to one of these eleven early attenders is the strongest predictor of subsequent attendance. 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 received the information via a personal visit

[5]Outside entities operating in Ugandan villages would typically share information first with the village Chairperson and elicit their help in spreading it, however this is not always the case due to variation in village official availability and quality. There are hundreds of villages in each Ugandan district; as of 2020, there were 146 districts and 70,626 villages in Uganda, according to Uganda's Electoral Commission. Accessed at https://www.ec.or.ug/electoral-commission-statistics

[6]We confirmed this with district level officials prior to our study, and with village officials immediately after our event.

by an enumerator. 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[7] and, notably, may be different individuals than those selected as seeds.

Our results further highlight the subtle role that network centrality plays. Villagers who are more central in the network (high closeness centrality) and who have connections to peers who themselves have lots of connections (high eigenvector centrality) are, as expected, more likely to have heard about the event. These network positions afford privileged access to information flowing through the network. However, among those who were informed, the *less* central were more likely to act on the information and attend. We discuss possible explanations, such as that those with less well-connected peers are freer to take novel, potentially risky actions, which is consistent with a social influence model in which conformity motives can act as accelerators or inhibitors of innovation (Young, 2009). It could also be that the most central face a higher economic –rather than social– opportunity cost to acting on new information. Adjudicating among these mechanisms is left for future research.

As a whole, our findings contribute new knowledge about how social networks play a role in the diffusion and adoption of new, difficult-to-verify, actionable information, first via a simple transmission of the seeded information, followed by a more complex process that creates secondary, social information that can motivate action. In the present context of valuable, actionable information spread in a village, the latter appears to be a function of lending credibility to the information. In sensitive situations, social ties may operate to reassure, reinforce, vet, and pass judgment on the action suggested by the seeded information. Consequently, 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 the seeded information, they are not necessarily

[7]This 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).

more likely to be convinced to act on it.

Beyond the insights into how social networks function to motivate (or inhibit) 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, security incidents, 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 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 if the information is sufficiently hard to verify and is encouraging behavior that is sufficiently unusual. When this is the case, 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.

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. After describing the design and ethics of the study in greater detail, and then presenting the results, in the conclusion we return to important considerations for the design of related future studies as well as promising directions.

2 Social Networks in Rural Uganda

We implemented our study in a village in the Teso region of Uganda.[8] The study village has approximately 1,400 adult 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 seven households, one selected at random from each of seven equally-sized areas within the village. 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 then 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 the village. A total of 138 people from the village 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.[9]

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

[8]We implemented the same study simultaneously in a neighboring village. News traveled poorly there, perhaps due to ethnic heterogeneity and mistrust [see: cite removed for anonymity] and only one person from that village acted on it. Since the villages' networks are distinct, we restrict analyses here to the village where action was more common. However, given other recent findings that different network structures correlate with ethnic demography (Eubank, 2019) and families' influence over village politics (Cruz, Labonne and Querubn, 2017), conducting similar studies in more villages that vary along these attributes will be an important next step for future research.

[9]On 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.

invited to take the survey, and at least one adult in all other households in the village 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 the village heard about the event but did not attend. 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.

2.1 The Uncertainty of the Seeded Information

One important feature of our design is the novelty of the information that we seeded and the action that it motivated. 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. Extensive daily, in-person discussions between one of the authors who was in Teso throughout the study and the local enumeration team indicate that many villagers found the event to be unusual. For example, despite the team's care to convey the information clearly, a few villagers expressed concerns about the possible presence of witchcraft at the event.

This uncertainty is further compounded by the opportunity cost of attending. 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.

If social networks play a special role in environments of uncertainty or sensitivity, we expect our design to capture this.

2.2 Ethics

When designing the study, we carefully considered the ethical issues posed by seeding new information in a community. Importantly, we wanted to provide information that, while novel and not straightforward to verify, would pose minimal, if any, risk of harm. With this goal in mind, we opted to share accurate information about a future event of our own design.

Participation was voluntary; we carefully trained enumerators to request informed consent to receive information about the event, and to respond to our event and post-event survey. We chose soap as the compensation for participating in our survey in consultation with several Ugandan colleagues who had experience as both researchers and with NGOs; our aim was to provide an item that would be genuinely useful to any household in the village.

We selected Teso because it was a region where one of the authors had conducted extensive qualitative fieldwork on a prior project and knew the local political and social terrain well. That author remained in Teso for the duration of the study to engage with local officials and directly oversee data collection. We chose the study village because of its size, demography, and rural character; it appeared to be roughly representative of typical villages in Teso.

We conducted the study with prior approvals from the authors' university Institutional Review Boards, from Uganda's National Council on Science and Technology, and from the relevant district-level officials. As discussed above, we waited until after our information was seeded and our event had ended before debriefing with the relevant village officials; we wanted to learn about how villagers would respond to new information that could not be easily verified with their local official.

When we learned on the last day of our event that some villagers had conveyed concern about the possibility of witchcraft at our event, we immediately arranged meetings with local village officials and local religious officials to understand the nature and extent of these concerns in the village. We learned from these meetings that any concerns that had existed were held by just a small number of villagers, and had reflected general skepticism about the existence of and intention behind the event, rather than a perception of physical danger.[10] We trained the research team that conducted surveys in the village after the event to offer and to share transparently any details of our study with villagers.

2.3 The Social Network of a Village in Eastern Uganda

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 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

[10] Connotations of witchcraft in African society do not necessarily include danger. For example, one scholarly volume explains, "Across the [African] continent, people see witchcraft less as extraordinary than as everyday and ordinary," and invoking witchcraft has been interpreted as communicating humor, irony, or skepticism (Moore and Sanders, 2003, 4, 5).

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.

Figure 1: Heatmap of Overlap Between Networks

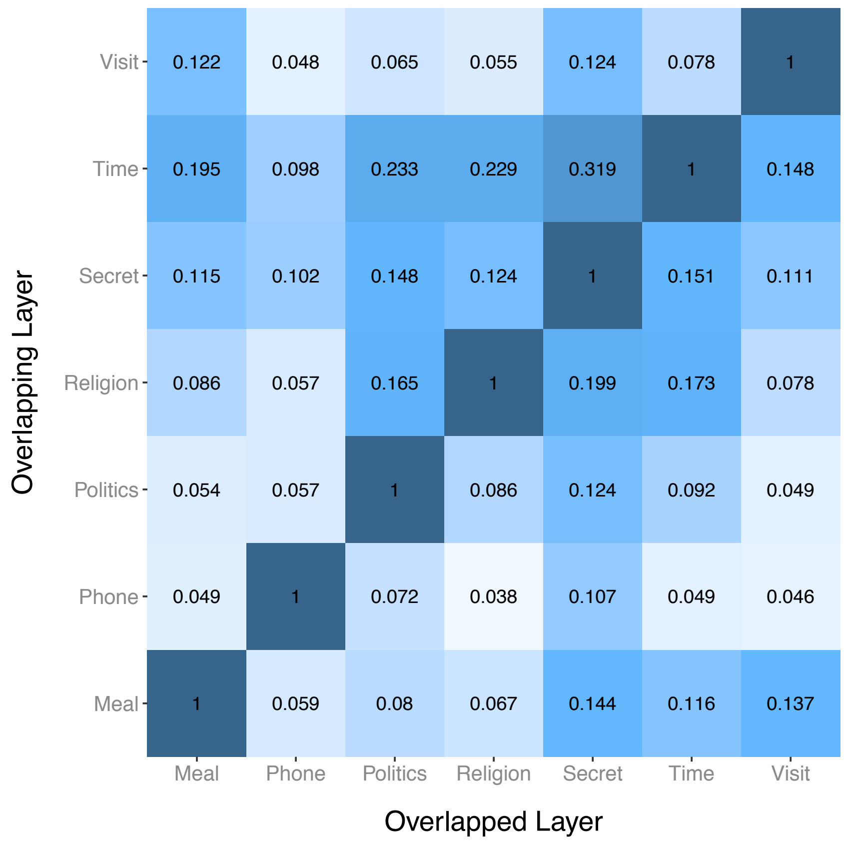

Note: we define overlap as the proportion of links of the overlapped layer (x-axis) contained in the overlapping layer (y-axis). Percentages are low in part because respondents could name anyone, including those outside of our sample. The same figure including only links that connect two respondents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

We begin our analyses using the union of these seven networks to capture an aggregate social network. In Section 3.3, we disaggregate the network to explore the role of each type of link.

2.4 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 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 (see Figure 2). A few differences in network position stand out. First, we look at eigenvector centrality, a standard measure of a node's influence in a network. A node has a high score (is influential) if it is itself connected to other highly connected nodes. We observe that 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 (Figure 2d). Second, both those who heard and those who attended have a larger proportion of their network neighbors who heard (Figure 2a). Finally, both those who heard and those who attended are significantly closer –are separated by a fewer number of edges–

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	
Housing material (1 = cement, 0 = mud)	0.10	0.13	0.14	
Educ. (3 = some primary, 4 = finished primary)	3.34	3.33	3.23	
Prop. unemployed	0.03	0.02	0.04	
Prop. working part time	0.42	0.47	0.46	
Prop. working full time	0.15	0.10	0.18	*
Prop. retired	0.39	0.42	0.32	*
Number of People	59	130	138	

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

to early attenders than those who neither heard nor attended (Figure 2b).[11] Interestingly, distance to a seed– a person with first knowledge of the information and given an information sheet from the information source– 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 3 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 information 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

[11]Distance between two nodes refers to the shortest number of edges connecting them. If two nodes are neighbors, their distance is equal to 1, the edge that connects them.

Figure 2: 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. Same legend applies to all figures.

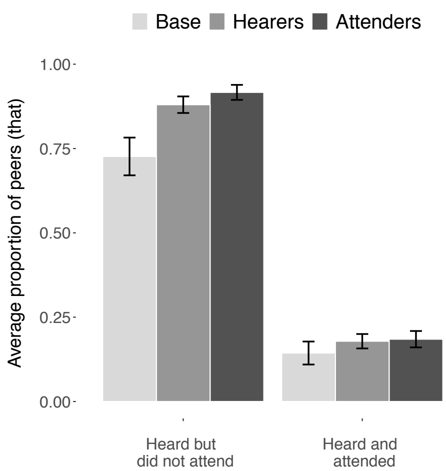

(a) Average proportion of peers that heard but did not attend and proportion of peers that heard and attended.

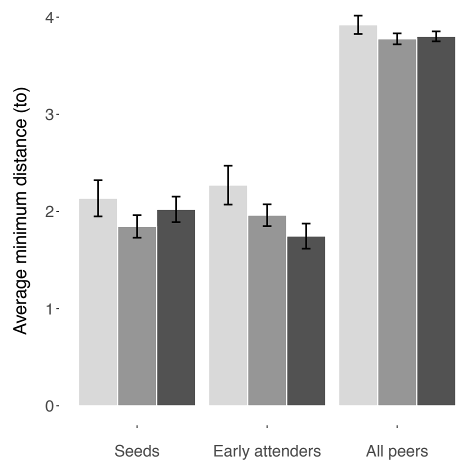

(b) Average distance to seeds, early attenders and all peers.

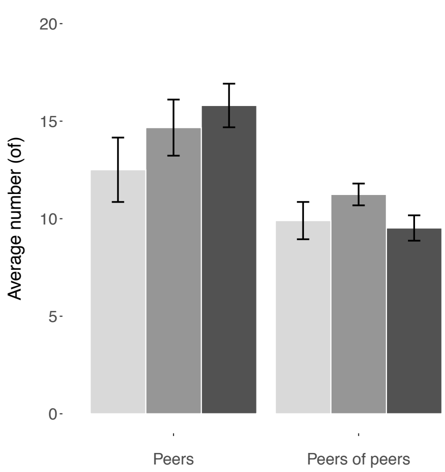

(c) Average number of peers and peers of peers.

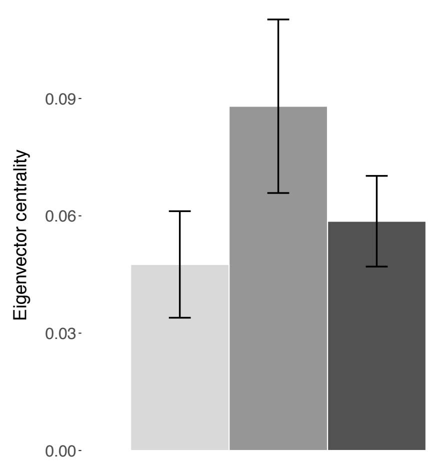

(d) Eigenvector centrality.

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.[12]

Figure 3: Percentage of links of each type a respondent said were used to transmit information about the event.

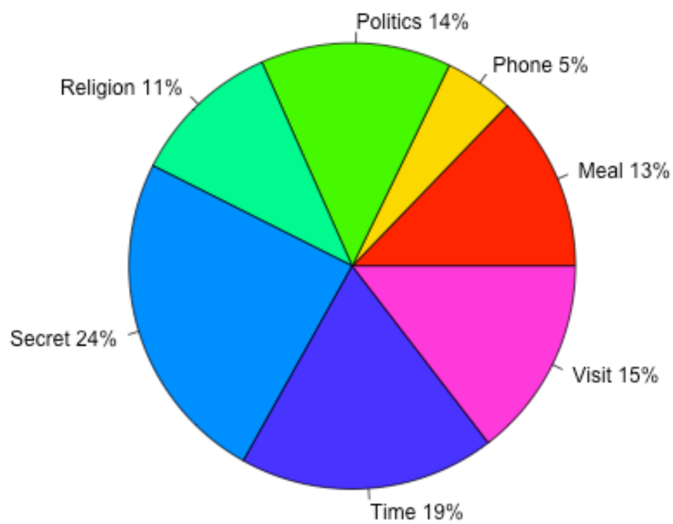

Inquiring about a person's source and target of information reveals only a small part of the process by which people become informed and motivated. This question detects the most salient initial sources and targets of information, and confirms that social networks provide

[12]That 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.

opportunities to pass information along. However, social networks also provide opportunities to generate social information – 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.

2.5 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 a village is, by definition, not independent of the network position of other respondents in that village. For this reason, we risk attributing too much precision to the results of regressions that rely on independence assumptions.[13] 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.

[13] Exponential 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.

3 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.

3.1 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 (Prop. of peers that heard), 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 (Number of peers) 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 2: Relationship between proportion of peers who heard and hearing about the event

	P(Hear About the Event)	
	(1)	(2)
Prop. of peers that heard	0.680***	0.647***
	(0.117)	(0.120)
Number of peers		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 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.[14]

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.[15] If an outcome spread through a network

[14]Because 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 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.

[15]Of 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.

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

	P(Attend the Event)	
	(1)	(2)
Prop. of peers that attended	0.083	0.086
	(0.224)	(0.225)
Number of peers		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.

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. Following this logic, we randomly sample, from the network, sets of 268 nodes – the number that heard about the event – and compare the number of ties among these sets to the number of ties observed among the 268 whom we know heard. We find that a number of ties 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 – reveals that we would observe a number of ties at least as high as the number among the actual attenders 18% of the time, falling short of conventional thresholds for statistical significance. Figure 4 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.

Figure 4: Sampling distribution of number of ties 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 number of ties of the observed group of hearers and attenders, respectively. Hearers cluster in the network consistent with a simple contagion process; attenders do not.

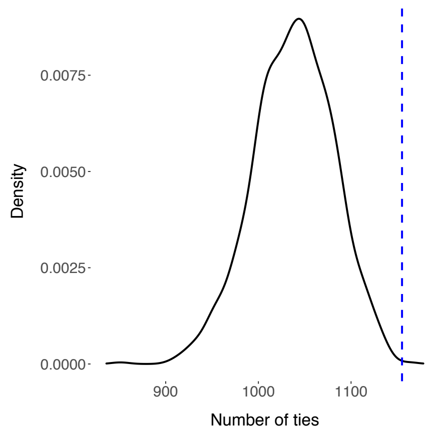

(a) Samples of the same size as the group of hearers.

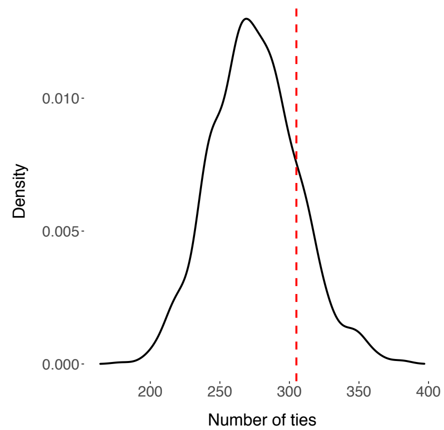

(b) Samples of the same size as the group of attenders.

3.2 Social Networks as Sources of Credibility

Next we examine evidence that the network functions as a source of credibility, in line with social learning models of diffusion (Young, 2009). 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; they received the information first. The second is the "early attenders." These eleven individuals chose to attend the event on the first day.[16] We calculate the network distance separating a person from their closest seed: if a person has a seed as a network neighbor, the shortest distance to a seed (Min. distance to seeds) is 1. If a person has no seed as a neighbor, but one of her neighbors has a seed as a neighbor, the shortest distance to a seed is

[16]Only two of the eleven early attenders were seeds. The correlation between people's shortest path to a seed (Min. distance to seeds) and shortest path to an early attender (Min. distance to early attenders) is .36. The maximum value for both is 4.

Figure 5: 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.

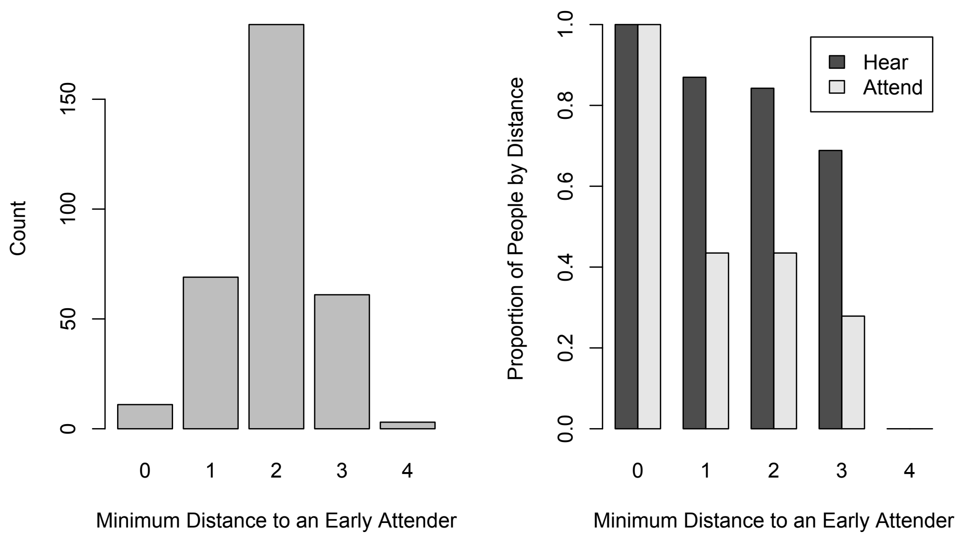

2, and so on. We calculate the same for distance to an early attender (Min. distance to early attenders). The Supporting Information Section 6 shows that a person directly connected to an early attender is at least 11.5% more likely to attend than a person whose closest connection to an early attender is 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 5 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 6 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.[17]

[17]In 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 6: Proportion of those in the sample at different distances to a seed who heard and attended. *More* of those farther from the seeds attended.

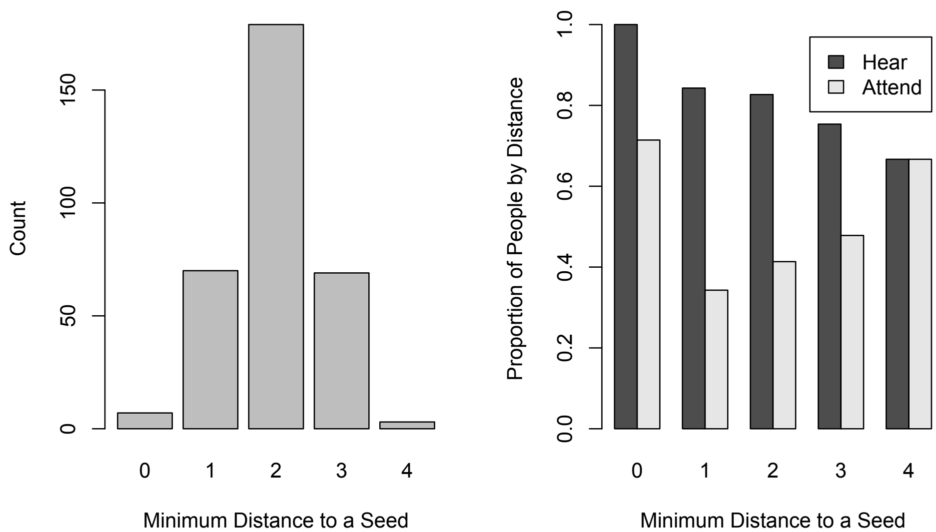

On the one hand, the seeds were privy to the original version, and possessed a means of corroborating their message: showing the information sheet. On the other hand, given the novelty of an event like this one in an otherwise rural, remote village, there may have been doubt about whether soap would 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.[18] 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.

[18]One enumerator reported that, upon arrival at the event and seeing one of the authors, 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.

3.3 Disaggregating Network Type

Our data allow us to consider qualitative differences in social relationship type. As Larson and Rodríguez (2020) point out, if people prefer to pass information along selectively, systematically making use of some ties and ignoring others, then aggregating network ties may miss this specialized diffusion. Table 4 separates the social network into its seven constituent layers, each 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).

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
Prop. of peers that attended	−0.172	0.315*	0.229	0.131	0.278**	0.252**	0.009
	(0.157)	(0.167)	(0.142)	(0.137)	(0.126)	(0.117)	(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.

Indeed, although we find 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.[19] The Supporting Information contains additional specifications and nonparametric tests that corroborate these rough magnitudes, as well as other information about each link type.

Motivation to attend appears to have spread more readily through the visits and meals links. Our data do not let us say precisely why relationships of these types were more relevant than others. Perhaps conversations about upcoming plans and their value happen most often in exchanges over meals and during visits, or perhaps those whom one visits and with whom one shares meals are those whose opinions matter most about right actions.

3.4 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 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

[19]The 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 also the Supporting Information, Section 10.

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)
Prop. of peers that heard	0.502***	0.499***
	(0.116)	(0.123)
Number of peers	0.002	0.002
	(0.003)	(0.003)
Eigenvector centrality	0.720***	0.741**
	(0.275)	(0.290)
Min. distance to early attender	−0.065**	−0.061**
	(0.030)	(0.029)
Female	0.143**	0.170**
	(0.061)	(0.068)
Age	0.0002	−0.001
	(0.001)	(0.001)
Catholic	0.033	0.014
	(0.039)	(0.041)
Educ. (3 = some primary, 4 = finished primary)	0.003	0.001
	(0.013)	(0.013)
Married	0.052	0.043
	(0.068)	(0.065)
Housing material (1 = cement, 0 = mud)	0.077**	0.082**
	(0.036)	(0.034)
Prop. of peers that are female		−0.043
		(0.066)
Age of peers		0.001
		(0.002)
Prop. of peers that are catholic		0.010
		(0.064)
Educ. of peers		−0.033
		(0.023)
Prop. of peers that are married		0.069
		(0.117)
Housing material of peers		−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.

negative relationship between eigenvector centrality and attendance.

3.5 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 Avg. distance to all), and the more a person is connected to other highly connected people (high Eigenvector centrality), 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), 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).[20]

[20]SI Section 11 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 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)
Prop. of peers that attended	−0.259	−0.078
	(0.294)	(0.322)
Number of peers	0.010*	0.012**
	(0.006)	(0.006)
Eigenvector centrality	−1.321**	−1.525**
	(0.608)	(0.680)
Min. distance to early attenders	−0.123**	−0.154***
	(0.053)	(0.056)
Female	−0.009	0.098
	(0.089)	(0.105)
Age	−0.002	−0.003
	(0.002)	(0.002)
Catholic	0.016	−0.038
	(0.071)	(0.078)
Educ. (3 = some primary, 4 = finished primary)	−0.037	−0.038
	(0.025)	(0.027)
Married	−0.159*	−0.178*
	(0.093)	(0.098)
Housing material (1 = cement, 0 = mud)	0.055	0.071
	(0.092)	(0.102)
Prop. of peers that are female		−0.401*
		(0.207)
Age of peers		0.001
		(0.005)
Prop. of peers that are catholic		0.237*
		(0.135)
Educ. of peers		−0.021
		(0.064)
Prop. of peers that are married		0.135
		(0.246)
Housing material of peers		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.

Table 7: Relationship between network centrality and attending the event

	P(Attend the Event)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Number of peers	0.005				0.015**	0.013**	0.005
	(0.005)				(0.007)	(0.005)	(0.004)
Number of peers of peers		−0.033***					−0.033***
		(0.009)					(0.009)
Avg. distance to all			0.064		0.325**		
			(0.097)		(0.148)		
Eigenvector centrality				−0.857**		−1.428**	
				(0.386)		(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.

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

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***
Housing material (1 = cement, 0 = mud)	0.11	0.10
Married	0.79	0.89
Educ. (3 = some primary, 4 = finished primary)	3.11	3.70
Unemployment	0.00	0.10
Employed part time	0.50	0.25
Employed full time	0.11	0.10
Retired	0.39	0.55
Number of peers	7.20	22.70***
Number of peers of peers	6.71	13.05***
Min. distance to seeds	2.75	1.35***
Min. distance to early attenders	2.75	1.80***
Avg. distance to all	4.49	3.48***
Prop. of peers that heard	0.90	0.90
Prop. of peers that attended	0.18	0.17
Eigenvector centrality	0.004	0.349***
Proportion that heard	0.80	0.90
Proportion that attended	0.55	0.35
Proportion that attended given they heard	.69	.39**
Number	20	20

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. This last explanation is in line with a social influence model according to which conformity motives can act as accelerators or inhibitors of innovation (Young, 2009). 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.

4 Conclusion

By randomly seeding novel information with individuals in a village in the Teso region of eastern 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 several parametric and non-parametric specifications, is consistent with a process of simple contagion assumed by many theories of information diffusion. Also, consistent with the proposition in network theory that central positions within a network offer greater access to information, we find that network centrality is associated with hearing the seeded information.

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 sources of secondary, "social information" is more important than proximity to people who were the initial recipients of the original information. We also 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 spreading the seeded information. 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 social information that early attenders bring back from their attendance 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, sensitive, or unconventional behavior. In the language of (Young, 2009), while the seeded information seems to have spread via simple contagion, social learning and social conformity both appear to have played a role in motivating behavior.

Our results suggest a number of important lessons and avenues for future research. First, the relationship we document between network centrality and behavior flags the importance for both theory and measurement to separately consider hearing information and acting on it. Without including in our study those who heard the information but did not act on it, we could not have observed the negative relationship between centrality and action among the informed, which would have generated misleading policy implications.

Second, these results highlight the importance of being sensitive to a local context before implementing interventions aimed at changing behavior. 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 uncertain, or even merely socially sensitive, then the spread of behavior will likely take a very different form than the spread of information.

Third, the fact that a core group of early adopters appear to have been important to subsequent behavior is consistent with critical mass theory, empirical studies of cascades in technology adoption, and the spread of political communication technologies (e.g. Centola, 2013; Foster and Rosenzweig, 1995; Ferrali et al., 2019) – but little is known about how to encourage this initial group. Our study hints that a tight-knit group may be a fruitful target for the initial group, establishing a local pocket of common knowledge that makes attending less costly within the group, and encouraging action by others outside the initial group later on. Note that in our study 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 via future research.

Lastly, our results corroborate methodological work about the importance of disaggregating social tie types when diffusion is link-specific (Larson and Rodríguez, 2020). We find that secondary, social information, such as about the veracity, value, and risk associated with a future event, is an example of link-specific diffusion – villagers selectively shared with social ties with whom they share meals and homestead visits over other types of ties. Measuring the relevant ties and disaggregating tie data are essential to uncovering link-specific spread.

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, Coralio, Antoni Calvó-Armengol and Yves Zenou. 2006. "Who's who in networks. wanted: the key player." *Econometrica* 74(5):1403–1417.

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.

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 Querubn. 2017. "Politician Family Networks and Electoral Outcomes: Evidence from the Philippines." *American Economic Review* 107(10):3006–37.

Dunning, Thad, Guy Grossman, Macartan Humphreys, Susan D. Hyde, Craig McIntosh, Gareth Nellis, Claire L. Adida, Eric Arias, Clara Bicalho, Taylor C. Boas, Mark T. Buntaine, Simon Chauchard, Anirvan Chowdhury, Jessica Gottlieb, F. Daniel Hidalgo, Marcus Holmlund, Ryan Jablonski, Eric Kramon, Horacio Larreguy, Malte Lierl, John Marshall, Gwyneth McClendon, Marcus A. Melo, Daniel L. Nielson, Paula M. Pickering, Melina R. Platas, Pablo Querubn, Pia Raffler and Neelanjan Sircar. 2019. "Voter information campaigns and political accountability: Cumulative findings from a preregistered meta-analysis of coordinated trials." *Science Advances* 5(7).

Eubank, Nicholas. 2019. "Social Networks and the Political Salience of Ethnicity." *Quarterly Journal of Political Science* (14):1–39.

Eubank, Nicholas, Guy Grossman, Melina Platas and Jonathan Rodden. Forthcoming. "Viral Voting: Social Networks and Political Participation." *Quarterly Journal of Political Science* .

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. 2019. "It Takes a Village: Peer Effects and Externalities in Technology Adoption." *American Journal of Political Science* 64(3):536–553.

Foster, Andrew D. and Mark R. Rosenzweig. 1995. "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture." *Journal of Political Economy* 103(6):1176–1209.

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.

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.

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.

Larson, Jennifer M and Janet I Lewis. 2018. "Rumors, Kinship Networks, and Rebel Group Formation." *International Organization* 72(Fall):871–903.

Larson, Jennifer M and Pedro Rodríguez. 2020. "The Attenuating Effect of Aggregating Networks." Working Paper.

Lewis, Janet I. 2020. *How Insurgency Begins: Rebel Group Formation in Uganda and Beyond*. Cambridge, U.K.: Cambridge University Press.

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.

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*

Moore, Henrietta L. and Todd Sanders. 2003. *Magical Interpretations, Material Realities: Modernity, Witchcraft and the Occult in Postcolonial Africa*. Routledge.

Nathan, Noah L. and Paul Atwell. 2020. "Channels for Influence or Maps of Behavior? A Field Experiment on Social Networks and Cooperation." *American Journal of Political Science* .

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.

Varshney, A. 2003. *Ethnic Conflict and Civic Life: Hindus and Muslims in India*. New Haven: Yale Univ Pr.

Young, H Peyton. 2009. "Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning." *American economic review* 99(5):1899–1924.