

Ethnic Networks  

Jennifer M. Larson New York University
Janet I. Lewis U.S. Naval Academy

Abstract: Active research on a wide range of political contexts centers on ethnicity's role in collective action. Many theories posit that information flows more easily in ethnically homogeneous areas, facilitating collective action, because social networks among coethnics are denser. Although this characterization is ubiquitous, little empirical work assesses it. Through a novel field experiment in a matched pair of villages in rural Uganda, this article directly examines word-of-mouth information spread and its relationship to ethnic diversity and networks. As expected, information spread more widely in the homogeneous village. However, unexpectedly, the more diverse village's network is significantly denser. Using unusually detailed network data, we offer an explanation for why network density may hamper information dissemination in heterogeneous areas, showing why even slight hesitation to share information with people from other groups can have large aggregate effects.

Replication Materials: The data, code, and any additional materials required to replicate all analyses in this article are available on the *American Journal of Political Science* Dataverse within the Harvard Dataverse Network, at <http://dx.doi.org/10.7910/DVN/W1TDGZ>.

Ethnic identity plays a fundamental role in collective action, especially in developing countries. Scholars have found a strong association between geographic concentrations of coethnics and higher levels of economic development and public goods provision (Habyarimana et al. 2009; La Ferrara 2002; Miguel and Gugerty 2005). A growing body of research also finds a link between subnational ethnic group concentration and organized, political violence (Cederman, Weidmann, and Gleditsch 2011; Kasara 2015; Larson and Lewis 2016; Toft 2003; Weidmann 2009). Furthermore, ethnic identity shared between candidates and voters in developing countries can bring about patronage-based ethnic block voting rather than policy- or programmatic-based voting (Bates 1983; Chandra 2004; Ejdemyr, Kramon, and Robinson 2015).

In several theories guiding such findings, ethnicity's importance is attributed to the way it allows people to *share information*. Foundational assumptions hinge on the premise that information flows more freely among coethnics than among members of different groups. While explanations for this difference vary, a common assumption is that information spreads better in ethnically homogeneous areas because ethnically homogeneous areas have denser social networks. These pivotal assumptions, however, remain largely implicit—and even when explicit, untested.

This article offers, to our knowledge, the first systematic, individual-level measurement of ethnic networks in a rural area of a developing country and direct investigation of their role in spreading information naturally (not in a lab).¹ We employed a novel experimental design

Jennifer M. Larson is Assistant Professor of Politics, New York University, Department of Politics, 19 W. 4th St., New York, NY 10012 (jenn.larson@nyu.edu). Janet I. Lewis is Assistant Professor of Political Science, U.S. Naval Academy, Political Science Department, 589 McNair Road, Annapolis, MD 21402 (janetilewis@gmail.com).

We thank Robert Bates, Andrew Coe, Nick Eubank, Guy Grossman, Lorian Jasny, Michael Joseph, Zane Kelly, Armando Razo, Pedro Rodriguez, Dustin Tingley, Stefan Wojcik, Josef Woldense, 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, the Department of Defense, or the U.S. government.

¹As of 2014, 63% of sub-Saharan Africans, 67% of South Asians, and 22% of Latin Americans and Caribbeans lived in rural settings (World Bank 2014).

American Journal of Political Science, Vol. 61, No. 2, April 2017, Pp. 350–364

©2016, Midwest Political Science Association

DOI: 10.1111/ajps.12282

that seeded identical information in the same way at the same time in two villages in rural Uganda that are similar in all measured attributes except ethnic composition: One is ethnically homogeneous, and the other is more diverse. The seeded information was that in three days an event would be held at which all adults in attendance would receive a valuable block of soap in exchange for taking a survey. After the event, at least one individual in all households in both villages was surveyed to complete the measurement of social networks and the reach of information.

The seeded information spread: 145 villagers attended the event. Every person who learned about the event did so by word-of-mouth, in-person interaction—no one learned by phone. As anticipated by existing theory, the information spread much less widely in the diverse village; more than eight times more individuals from the homogeneous village heard about the event. However, an investigation of the social networks in both villages calls into question the mechanism of network density. We find that the social network in the more diverse village is significantly denser, and this relationship holds when considering different conceptions of density, when comparing only the networks among coethnics in each village, and when controlling for other network differences. In other words, news spread much more widely in the homogeneous village *despite having a less dense social network*.

To better understand why a sparser network could spread information more widely, and to consider the generality of this result, we simulate an information-spreading process on the observed networks in the two villages. We show that in the presence of ethnic heterogeneity, if there is greater hesitation to share information between non-coethnics than there is between coethnics—even if this difference is small—greater network density can impede the spread of information, especially when that greater density is due to more cross-ethnic ties. The intuition is as follows: Given a limited number of opportunities to share news with a social contact on a given day, greater density increases the likelihood that some of those opportunities are with less trusted social contacts. Additional ties to those one trusts less crowd out the chance to interact with more trusted people to whom information flows more freely. Because of the exponential nature of information spread, when this dynamic recurs throughout a network, it can greatly impede the process of information spread overall.

We also show that under quite general conditions, heterogeneous areas would be expected to generate networks that are denser, and that the greater density would

be largely due to additional cross-group ties. The more the different ethnic groups specialize socially and economically in a heterogeneous area, the more the network would be dense and feature cross-group ties, and hence the more likely it would be to impede the spread of information.

The article makes several contributions to knowledge on how information spreads in rural areas and the role that ethnicity plays in the process. For example, while information communication technologies (ICTs) can affect information flows to and from marginalized communities (Grossman, Humphreys, and Sacramone-Lutz 2014), potentially improving political accountability, this article shows that face-to-face interaction still presents a crucial mode of transmitting sensitive information for rural villagers, at least in eastern Uganda. It also provides a rare look at how this works. Our experimental design, which combines both controlled comparison and detailed measurement of networks, allows us to provide direct support for a key assumption of theories of ethnic collective action: that information spreads more widely in communities that are ethnically homogeneous than in heterogeneous ones.

However, our results suggest that a usual rationale behind this assumption, that networks in homogeneous areas are more dense, may not be accurate in practice. We show why density in fact can be greater in heterogeneous areas and yet severely hinder information dissemination there. While our evidence base is modest, these stark results suggest that the focus on the ability to spread information within ethnic groups is well founded, but the connection to network density needs to be better understood. In particular, this suggests the need for more research on how the nature of ties affects information dissemination; contrary to Granovetter's (1973) groundbreaking argument, weak ties may not be strong for the purposes of sharing brand-new information. Further, we show why there is good reason to expect greater network density in diverse areas, given minimal assumptions about group specialization.

Finally, we present a logic of information dissemination and network formation that establishes the conditions under which the empirical findings here would generalize. These findings have implications for research about the relationship between ethnicity and various forms of collective action, as well as policy interventions that disseminate information to citizens—from those aimed at building a more informed electorate to those aimed at countering false rumors about, for example, how diseases like Ebola spread—in rural regions of developing countries.

Relationship to Existing Research

Coethnics' Informational Advantage

While numerous core political science theories posit that shared ethnic identity among individuals enhances their ability to act collectively, scholars invoke several mechanisms to explain *why* this may be the case. The most prominent mechanism of ethnic collective action emphasizes a coethnic advantage in socially sanctioning other group members (Fearon and Laitin 1996; Miguel and Gugerty 2005). Implied here is an assumption of coethnic informational advantage since coethnics are more likely to learn when a group member has defected from the group, and are more likely to learn how to locate the defector in order to punish her. A second, related set of explanations for ethnic collective action centers explicitly on the role of improved information dissemination among coethnics, which can facilitate spread of incendiary rumors (Varshney 2003), enhance the ability to form viable rebel groups (Larson and Lewis 2016), or serve as an obstacle to social mobility (Deutsch 1953). In a third family of explanation, ethnicity's primary function is that it serves as a cue of credible willingness for voters to supply votes and for patrons to provide favoritism or benefits (Bates 1983; Chandra 2004; Conroy-Krutz 2013; Ferree 2010; Posner 2005). A related variant of mechanisms posits that ethnicity confers advantages because group members share *preferences* over outcomes such as policy goals, types and locations of public goods, or even the wisdom of rebelling against the state. Even for such accounts of ethnic collective action that are not explicitly related to information flows, a coethnic advantage in information sharing is arguably implied, or at least helpful; the processes that lead to expectations of intra-ethnic patronage relations or shared preferences are often highly social ones, requiring communication and coordination.

Central to all of these theories of ethnic collective action, then, is an assumption that communication travels among coethnics more easily than it does among individuals of different ethnicities. However, few existing works present evidence about how information flows among and between groups.² To fill this gap, we seeded information in randomly selected households in two villages of eastern Uganda—one village that was quite homogeneous and one that was not—and then surveyed village

²An exception is Habyarimana and colleagues' (2009) groundbreaking book, which finds evidence for improved "reachability" among coethnics compared to those from different ethnic groups in the slums of Kampala. However, this research design does not allow for studying information transmission directly, nor in rural settings, where over 80% of Ugandans and the majority of sub-Saharan Africans reside.

members to understand their social network and the process of spreading the information we seeded. While this research design of closely examining two villages is, naturally, constrained in its ability to support general claims, it allows us to collect data rich enough within each village to peer inside the black box of ethnic group information networks. Doing so is essential to answering fundamental questions not only about whether, but especially about *why* ethnic homogeneity confers informational advantages.

Network Attributes That Facilitate Information Dissemination

Our approach also allows us to reconsider what attributes of networks facilitate information dissemination. In particular, many seminal contributions in political science specifically attribute coethnicity's felicitous effects on collective action to ethnic groups' "dense" networks (e.g. Chandra 2004, 71–72; Fearon and Laitin 1996, 719; Gubler and Selway 2012, 210; Miguel and Gugerty 2005, 2330). For example, in his classic work, *Bowling Alone*, Robert Putnam writes of the "dense social ties" that characterize "close-knit ethnic enclaves" (2000, 21). In another influential piece, Miguel and Gugerty (2005, 2330) state that they "assume that social sanctions and coordination are possible within groups due to the dense networks of information and mutual reciprocity that exist in groups but are not possible across groups." Yet the comparative density of networks among groups of coethnics and of non-coethnics has not been measured systematically to date in the developing country contexts that such works (except for Putnam 2000) typically aim to address. Moreover, why density may matter is rarely theorized; rather, the general logic offered is that denser networks have more links of interaction in them, and the more people in a network interact, the more likely they are to share information.

However, network density may not always work this way for two reasons. First, it can be that the links in a dense network are arranged in such a way that the reach of information is limited despite the number of links being large (Alatas et al. 2015; Jackson and Rogers 2007). Second, realistically, nodes in *social* networks do not always transmit information along every possible link. Many models of transmission or contagion draw on analogies to deterministic systems like a physical computer network in which a signal sent from one computer is automatically broadcast to all connected computers. However, in the less tangible network of social ties underlying village life, a person may not encounter every social contact at every moment, may encounter different social contacts at

different moments, and may differentially choose whether to share information during an encounter based on the nature of the tie. We argue that people in dense social networks may interact frequently, but some may lack sufficient trust or motivation to share information, especially somewhat sensitive news, in the same way with every social contact.

Our argument thus departs from Granovetter's classic argument about "the strength of weak ties" (Brown and Konrad 2001; Calvó-Armengol and Jackson 2004; Granovetter 1973), since we posit that certain types of "weaker," cross-ethnic ties inhibit information spread. This difference stems largely from the process of information spread we respectively envision; while the information process envisioned by Granovetter can be thought of as a demand-driven one through which potential receivers seek out information from their network—an unemployed person looking for information about employment opportunities, say—we instead focus on a supply-driven process. That is, our findings and theory apply to cases in which a person possesses brand-new information that others do not know exists, and the onus is on that person to choose with whom he shares it. Weak ties can in fact be weak conduits for new information in supply-driven processes like the one studied here.

Measuring Networks in the Field

This article also contributes to a growing literature in the social sciences that measures social networks in the field. Studies aiming to draw inferences at the level of the network face a sharp trade-off between the number of distinct networks measured and the resolution at which each network is measured.³ Indeed, existing studies of a large number of networks tend to sample a small proportion of each network's nodes (Alatas et al. 2015; Baldassarri 2015; Banerjee et al. 2013; Jackson, Rodriguez-Barraquer, and Tan 2012).

While data sets containing a large number of networks permit comparisons of group-level outcomes with high statistical power, comparing network features becomes problematic due to the small sample of nodes. Because a network's nodes are systematically related to one another, small random samples of nodes do not necessarily reveal properties of the larger network in the same way that random samples of independent observations

³For some research questions, coarse estimates of network properties or sparsely measured networks can suffice. This tends to be the case when the goal is to establish the existence (rather than the size or source) of peer effects or learning. When the research question has to do with network *structure*, more precise measures are needed.

reveal features of the broader population from which they are drawn. Sparse samples of nodes, even when randomly drawn, introduce imprecision and often bias the measurement of network properties (Lee, Kim, and Jeong 2006).

We contribute a new approach, which allows for an unusual comparison of two detailed networks. Our design measures a small number of networks with a large sample of nodes: We sample someone in all households in two villages. Our approach selects the small number of networks in a way that maximizes the difference along one dimension—ethnic diversity—and holds constant as much as possible about other demographics and the intervention. This permits an unusually detailed paired comparison of two networks, allowing us to offer direct support for claims that information spreads better in ethnically homogeneous areas. Furthermore, by accompanying this study with theory—that identifies a mechanism by which the outcome would hold in general and of which we show the plausibility via simulations on this rich network data—this article identifies a promising agenda for future empirical research.

Measuring Information Flows in Rural Uganda

Location

We introduced novel information and measured its spread through social networks in two villages. The villages are located in the Teso region of Uganda, approximately 10 kilometers north of Soroti town: Abalang in Arapai Parish, and Mugana in Aloet Parish.⁴ The two villages are approximately two kilometers apart at their closest points. They are demographically similar but differ in one key respect: Abalang is ethnically and linguistically homogeneous and Mugana is heterogeneous.

Both villages have approximately 1,400 residents, both are rural, and both are located in the same geographic region of Uganda. Mugana is a little wealthier, measured in terms of the quality of wall material in houses, and has more Catholic residents.⁵ The starkest difference between the two villages is in ethnic composition:

⁴The Teso region is policy-relevant as a post-conflict site: The region was home to a rebel group, the Uganda People's Army, which launched a rebellion against the Ugandan government from 1987 to 1992 in a village near Abalang. Fieldwork conducted by one of the authors reveals a relationship between the networks among civilians in the surrounding area and the rebels' success in becoming a viable threat to the government.

⁵Abalang has more respondents with middle-quality wall materials and fewer with high-quality wall material, and it also has a smaller total proportion of Catholics. Because those who attended

TABLE 1 DV: Respondent Heard about the Event

	(1)	(2)	(3)	(4)
Gender	1.044*** (0.350)	1.175*** (0.363)	0.900** (0.364)	0.995*** (0.375)
Age	0.001 (0.010)	0.003 (0.010)	0.004 (0.011)	0.004 (0.011)
Educ	0.159 (0.115)	0.122 (0.118)	0.171 (0.119)	0.145 (0.122)
WallMat	0.218 (0.203)	0.238 (0.202)	0.245 (0.212)	0.266 (0.214)
Married	0.766 (0.471)	0.821 (0.477)	0.712 (0.491)	0.742 (0.494)
PartTime	-0.386 (0.345)	-0.387 (0.348)	-0.454 (0.365)	-0.449 (0.367)
FullTime	-0.918 (0.526)	-0.738 (0.538)	-0.983 (0.540)	-0.841 (0.553)
Anglican	-0.313 (0.428)	-0.306 (0.431)	-0.295 (0.437)	-0.275 (0.441)
Pentacostal	0.364 (0.405)	0.378 (0.408)	0.305 (0.427)	0.318 (0.429)
Degree		0.039 (0.024)		0.029 (0.025)
PHoodHear			1.599*** (0.513)	1.572*** (0.516)
Constant	-1.953 (1.078)	-2.661** (1.173)	-2.987** (1.218)	-3.515*** (1.312)
Observations	198	198	194	194
Log Likelihood	-123.927	-122.495	-115.579	-114.846
Akaike Inf. Crit.	267.854	266.989	253.158	253.693

Note: Data are from the Abalang post-event survey.

p < .05, *p < .01.

Measured as ethnolinguistic fractionalization, Mugana, at .48, is much more ethnically diverse than Abalang, at .08. Measured another way, the size of the second largest ethnic group in Mugana is 36% of the population, and in Abalang it is only 4% of the population. Full demographic statistics for both villages can be found in Table 1 of the online supporting information.

Our village-level comparisons are bolstered by three details. First, Mugana and Abalang are similar in geography, size, age distribution, income distribution, employment status, technology adoption, clan structure, and religious attendance but differ starkly in ethnic composition. Second, the predominant ethnic group in Abalang is also

the event from Abalang were not the poorest in Abalang, and because the proportion of Catholic attenders is indistinguishable from the proportion of Abalang that is Catholic, we feel confident these small differences do not explain the large differences in the reach of information or attendance.

the largest ethnic group in Mugana. Third, we performed the intervention at the same time in the two villages, implemented it identically, and measured networks in the same way in both.

Intervention

The intervention randomly seeded novel information in both villages. Specifically, we seeded information with 7 randomly chosen households in Abalang and 10 in Mugana.⁶ For each selected “seed” household, a Ugandan enumerator who was not from the village personally visited and shared the information that starting in three

⁶We selected a larger number of seeds in Mugana due to its larger size. Information about our randomization design and the balance across each village’s seeds can be found in Tables 2 and 3 in the supporting information.

days, an event will be held at which all adults who attend and take a survey will receive a large block of soap. A limitation of this intervention is that it cannot directly speak to the spread of highly sensitive information, as in a conflict setting; transmitting such information would risk causing harm to villagers. However, the novelty of this information, and the difficulty of verifying it, gave it a mysterious and somewhat sensitive quality. Based on pretesting and one of the author's extensive, daily, in-person debriefing with the enumeration team, villagers found the news surprising and unusual.⁷ The information was seeded on the same day in both villages; seeding the 17 households took a full day. To mitigate against enumerator differences, the two enumerators each began in a different village, randomly determined, and switched villages halfway through the day. 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 was held at a church between and equidistant from Abalang and Mugana.

A total of 145 individuals attended the event. Strikingly, only 2 of the 145 were from Mugana or elsewhere in Arapai Parish.⁸ In total, 135 were from Abalang or Aloet Parish, and 8 were from farther villages.⁹ Respondents in attendance reported traveling between 2 and 180 minutes to reach the church, with a mean travel time of 51 minutes. The survey at the event collected information about demographics, networks, and learning about the event.

After the event, enumerators conducted surveys in both villages to collect information about demographics, networks, and knowledge of the event. Respondents who

⁷Furthermore, pretesting indicated that the value of soap is roughly equal across Ugandan villages. Post-event surveys reveal that a bar of soap lasts a household on average 6.8 days in Abalang and 6.2 days in Mugana, and the difference is statistically insignificant. Soap is a desired item throughout rural Uganda, where most people are subsistence farmers and use soap to wash their clothes and dishes. One large bar of soap costs about 70 cents.

⁸The difference in attendance is striking; however, in the supporting information, we show that the difference is not attributable to differences in access to the event, villagers' valuation of soap, conflicting events, interest in community events, clarity of information, or enumerator characteristics. We also show that the difference cannot be attributed to systematic differences in the selected seeds with respect to demographic characteristics, network position, or formal leadership roles.

⁹Attendees from Aloet Parish were mostly from Abalang, but also from the villages Central Aloet and Akum. In Arapai Parish, one attendee was from Mugana and the other from Amoru. The farther villages included Ogoloi, Akaikai, and Gweri, which are on the far side of Abalang from Mugana.

said they knew about the event but did not attend were asked what attendees received in order to verify that they in fact knew. The sampling procedure entailed first visiting all households within view of the seed households and inviting all adults to take the survey, and then visiting all other households in random order and inviting at least one adult to take the survey. In total, 226 individuals in Abalang and 237 in Mugana were surveyed in the post-event wave.

We leave analyzing event attendance for future work. Here, we focus on the spread of information about the event. The post-event survey reveals that news spread much more widely in Abalang, where it reached 62% of respondents, compared to Mugana, where it reached only 9% of respondents. In short, the same news seeded the same way about the same event at the same time spread much better in the homogeneous village.

The Importance of Word-of-Mouth Networks


Individuals in Abalang and Mugana were asked a battery of questions about their social lives in order to detect the networks along which information could spread in the two villages. Individuals were asked to name up to five people (for each category) whom they visit and with whom they discuss politics, discuss religion, spend time, share meals, share secrets, and speak on the phone. The exact text of the networks questions can be found in the supporting information. The "social network" is a directed network constructed from a union of these seven networks.¹⁰ This design measures informal relationships of the form sometimes called "quodidian interactions" (see, e.g., Parkinson 2013), the kind of meaningful relationships that may generate opportunities to spread information in person via word of mouth.

The alternative communications technology in the two villages is cell phones. However, cell phone penetration is low. Only 24% of respondents in Abalang and 27% in Mugana have phones. Moreover, those who do have and use telephones use them largely to reach people outside the village.¹¹ Phones, when used, are primarily for long-distance communication and not for relaying

¹⁰The analyses presented here use the union of the seven networks constructed as the closed network among respondents. The supporting information demonstrates robustness to the inclusion of ties to nonrespondents as well as the inclusion and exclusion of certain types of ties, and it considers the potentially different impact of each network.

¹¹When phone owners were asked to name up to five people whom they call, 72% of the names offered by respondents in Abalang and 65% in Mugana are of people outside the village.

FIGURE 1 Abalang (left) and Mugana (right) Social Networks, Visualized with Force-Directed Layout



Note: Those who heard are shaded black; node size is proportional to degree.

day-to-day local news. Most tellingly, of the 269 individuals who said they were told about the event or who reported telling someone else about the event, exactly 0 reported hearing or passing the news along over the phone.¹² Every respondent who learned the seeded information learned via word of mouth in a face-to-face interaction.

Our data affirm the role of networks in spreading news. Figure 1 combines the networks survey information with the outcome of the intervention to generate a visualization of the Abalang and Mugana social networks that highlights the nodes who heard about the event in each village. Quick inspection reveals two comparisons: First, there are many more links in Mugana's network, a comparison we substantiate below, and second, those who heard tend to be near each other in the network. The latter is the first hint that this network captures something relevant to the information-spreading process.

A logistic regression of hearing about the event on demographic and network features corroborates this. Table 1 reports the results of four specifications, with additional specifications that disaggregate tie type in the supporting information.¹³

Specification (1) regresses hearing about the event on only demographic characteristics. The only significant

predictor of hearing is being female, which is associated with being more likely to hear. Adding network features to the specification suggests that having network neighbors who heard is part of the explanation. Having a larger proportion of ties in the social network who have heard about the event (PHoodHear) is consistently and significantly associated with a greater likelihood of hearing about the event. While this is not a causal specification, it supports the search for a networks-based explanation and shows that the network is related to information spread, even after accounting for the demographic attributes of the respondents.

Social Networks in Abalang and Mugana

The network measured in Mugana includes 234 respondents, whereas the one in Abalang includes 216; comparisons of network features must account for the different sizes. Table 2 presents an overview of the comparisons. Although both networks exhibit similar levels of reciprocated ties¹⁴—about 20% of the ties in both—the network in Mugana is significantly denser, even after statistically controlling for the number of nodes (the procedure is described in Section 2.4 of the supporting information). We now turn to a deeper investigation of this difference.

¹²Phone penetration among event attenders corroborates this—24% of attenders have phones, exactly the same proportion as Abalang in general (where most attenders were from). If phones were used to spread the information and encourage attendance, we would expect a higher proportion of phone ownership among attenders than in the village as a whole.

¹³As shown in the supporting information, network neighbors in the visiting others, sharing meals, and sharing secrets networks who have heard about the event are most predictive of hearing about the event, which is consistent with the trust explanation offered below.

¹⁴The extent of reciprocated ties is often used as some indication of measurement error. The similar rate in both villages suggests that there is no systematic difference in measurement error between the two networks.

TABLE 2 Network Comparisons between Abalang and Mugana

Networks within Villages		
	Abalang	Mugana
Total Nodes	216	234
Total Links	660	965
Reciprocal	.20	.21
Average Degree	6.1	8.2
Density	.014	.018

Note: Comparisons are made using the post-event surveys and closed networks such that only links among survey respondents are included (see Banerjee et al. 2013, and the discussion in the supporting information). Density in Mugana is statistically significantly larger, determined using a simulated comparison in which nodes were randomly removed from Mugana until it reached the size of Abalang; the supporting information describes the procedure in greater detail. $p =$ probability of observing density as high in simulated Abalang as it is in Mugana. $p < .001$.

Density

Given conventional wisdom about dense networks within ethnic groups and in light of the outcome that news spread poorly in the ethnically heterogeneous village, the most surprising difference between the two villages' networks is that Mugana's network is denser than Abalang's. In fact, Mugana's network is denser in every sensible meaning of *dense*.

In the strict network sense, a network's density is the ratio of links to the total number of links that could possibly be present among a set of nodes. By this measure, Mugana is denser: .018 compared to .014, a difference that is robust to the difference in size between the two networks and is substantively large.¹⁵ It follows that the total number of links present in Mugana's social network is greater and that the average number of links per person is greater. The difference even holds when we break apart the aggregate: Individuals in Mugana listed a greater total number of people in response to the seven social networks questions (17 names vs. 14 in Abalang), with a similar unique proportion (.76 vs. .74 in Abalang). By the strict

¹⁵The total number of possible directed links among a set of n nodes is $n(n - 1)$: There are 46,440 links possible in Abalang's network and 54,522 links possible in Mugana's network. Given these denominators, a difference in density of .004 is large. The magnitude is easier to interpret in terms of average degree. This difference in density translates to each individual having on average 38% more social ties in Mugana than in Abalang. Since the number of nodes is larger in Mugana, attaining a high density is even more demanding. Section 4 of the supporting information presents further evidence of the robustness of the density comparison as well as differences in density by network type.

network definition or a more colloquial one, Mugana's social network is denser than Abalang's.

This relationship holds even at a lower level of resolution: The networks among only coethnics—the “ethnic networks”—within the heterogeneous village are denser than the ethnic network in the homogeneous village, even when controlling for size. Table 3 breaks down the network statistics by ethnicity in both villages and helps unpack the source of greater density. The network density among only the Iteso in Abalang is .015, whereas the network density among only the Iteso and among only the Kumam in Mugana is .021 and .029, respectively. Thus, both ethnic groups in Mugana are more densely connected among coethnics than the ethnic group in Abalang is connected among coethnics. This is true even when controlling for the size of the groups: If Abalang's ethnic network had as few nodes as Mugana's ethnic networks, it would still not be as dense.

In addition to the greater intra-ethnic density, Mugana also exhibits an abundance of interethnic ties. Over a third of social connections in Mugana are between non-coethnics. For comparison, if individuals formed social ties strictly at random without any coethnic preference in Mugana, the expected proportion of cross-ethnic ties would be .48. The observed proportion of cross-ethnic ties, .39, is inconsistent with a purely random formation of ties blind to ethnicity, but it is also inconsistent with a strict preference for coethnic ties; in fact, the value is closer to random formation than perfect coethnic separation.¹⁶

Figure 2 shows a visualization of the Mugana network, this time arranged so that all of the Iteso nodes are displayed on the left and all of the Kumam nodes are on the right. The horizontal lines show the extent of interethnic social relationships. This figure starkly shows just how prevalent cross-group ties are in Mugana's social network.

Summary of Observed Network Differences

The villages of Abalang and Mugana are similar in all measured demographic characteristics except for ethnic composition: Abalang is ethnically homogeneous and Mugana is more ethnically diverse. Despite the greater diversity, the social network in Mugana is substantially denser than the social network in Abalang. This is true in terms of


¹⁶Section 3 of the supporting information shows that the observed proportion of interethnic ties, .39, would fall within 95% of the sampling distribution of tie formation processes in which the probability of a coethnic tie falls within the range of [.58, .64]. Any higher probability of coethnic ties would make an observed proportion of non-coethnic ties of .39 statistically improbable.

TABLE 3 Comparison of Whole Village Networks and the Networks among Coethnics within Villages

		Ethnic Networks		
		Abalang	Mugana	
Village Networks	Node Count	216	234	
	Link Count	660	965	
	Density	.014	.018	
		Iteso	Iteso	Kumam
Ethnic Networks	Node Count	203	144	76
	Link Count	615	423	164
	Density	.015	.021	.029
	Cross-Group Tie Count			374
	Cross-Group Tie Prop.			.39
	Average Coethnic Degree	6	5	
	Avg. Non-coethnic Degree	.1	3.2	
Normalized Comps	Abalang Iteso Density < Mugana Iteso Density, $p < .001$.			
	Abalang Iteso Density < Mugana Kumam Density, $p = .01$.			

Note: Normalized comparisons are the result of simulations in which nodes are removed from Abalang at random to match the size of the smaller ethnic network to which it is compared; p is the proportion of simulations in which the density in the simulated shrunken Abalang is larger than the comparison network’s density. The Kumam comparison is relatively less significant despite the larger absolute difference because Kumam is a much smaller network; if Abalang were that small, by chance its density would be that high 1 in 100 times.

FIGURE 2 Mugana’s Network with the Iteso Tribe (left) and Kumam (right)



Note: In all, 39% of social ties are between non-coethnics.

network density, total number of links, total number of links per capita, total number of other individuals listed in answer to network questions, and total number of distinct individuals listed in answer to network questions. Additionally, the two ethnic subnetworks within Mugana (among the Iteso and among the Kumam) are themselves denser than the ethnic network within Abalang (among the Iteso). Cross-ethnic ties are abundant in Mugana.

The outcome of the intervention is consistent with existing theory—news spread less well in the heterogeneous area. The mechanism, however, is not. Contrary to

the expected story that news would spread well in the homogeneous village *because the social networks are denser*, news in fact spread well in the homogeneous village *despite the social network being less dense*.

To help sort out idiosyncrasy from regularity, and to offer predictions for future investigations with a larger number of networks, we turn in the next section to a theory of information spread and a theory of network formation. Through the first, we show that the greater density in heterogeneous societies is not necessarily favorable to the spread of information. Through the second, we show there is a wide range of plausible conditions under which

heterogeneous societies would in fact be expected to have denser social networks.

Barriers to Information Spread

In this section, we outline the logic for why a network may be denser but be substantially worse at spreading information. The logic relies on variance in willingness to spread news along certain social ties. A social network determines the set of people with whom one may have the opportunity to share information. However, the social network neither guarantees that every opportunity is realized—some days one may encounter few of one’s social contacts—nor does it guarantee that a person is equally willing to share information with every social tie, even given an opportunity.

In light of this characterization, we consider a model of information spread in which individuals only have a limited number of opportunities to share information at any moment, and which allows an individual’s willingness to share information to vary by tie. Here we conceptualize this willingness as trust, given the growing body of evidence suggesting higher levels of trust at the local level among coethnics in sub-Saharan Africa than that between non-coethnics (see especially Robinson 2015, 2016). Section 7 of the supporting information presents evidence from the Afrobarometer survey corroborating the existence of these trust differentials. In the model, when an individual happens to encounter a coethnic, she is more likely to share information than when she happens to encounter a non-coethnic.

We simulate this process on the observed networks in Abalang and Mugana using observed ethnicity of the nodes to show how differences in trust can inhibit the spread of information. Before discussing the model and simulation results, we first present suggestive evidence that a process like the one described may have been at play in the two villages.

First, we unpack the role of the proportion of one’s network neighborhood that hears about the event (PHoodHear) in the specifications of Table 1 by looking at the subset of ties to others who have heard that are the same as the respondent in terms of tribe, clan, the union of one’s and one’s spouse’s clan, and religion. Table 4 shows that the proportion of one’s network neighbors who have heard and are the same tribe (PHoodTribeHear) is a stronger predictor of hearing than any other characteristic of one’s neighbors who have heard. If a large proportion of a respondent’s social ties is the same tribe and has heard about the event, the respondent is more likely to have heard about the event.

Furthermore, we can consider the subset of respondents who heard the news and ask what makes them more likely to pass that news along. Table 5 shows that having a higher level of education consistently predicts a greater likelihood of spreading news, and even when controlling for the extent to which one spends a typical week with members of the same clan and the same religion, having a higher proportion of one’s social contacts as members of the same tribe (PHoodTribe) is a strong predictor of passing the news along to someone.

In short, the data suggest that those who had more opportunities to encounter a coethnic were more likely to spread the information conditional on hearing it.

Simulated Information Spread

We consider a discrete-time opportunity model of information spread in which an agent who knows information may pass it along when given the opportunity. Ex ante, agents have a type $\tau \in \{A, B\}$ and are arranged in a social network g . In $t = 0$, one randomly chosen individual, the “seed,” is endowed with some information. In $t = 1$, the seed encounters a randomly selected subset of size x of his neighbors in the social network. The seed passes the information in each of the x opportunities probabilistically: For each neighbor he encounters, he passes the information along with probability p_{same} if the neighbor is his same type, p_{dif} if a different type, with $0 \leq p_{same} \leq 1$ and $0 \leq p_{dif} \leq 1$. In $t = 2$, the seed’s network neighbors who received the information in $t = 1$ each encounter x of their neighbors at random and pass the information in the same way, and so on in subsequent periods.

This type of information-spreading process captures a scenario in which individuals encounter some, but not necessarily all, of their social contacts on any given day. The reason for the encounters may have nothing to do with passing along information: In a day, someone may need to speak with a teacher, visit the market, and have a meal with her friend. These encounters present opportunities to talk about news of the day. A person may tell everyone she encounters, or she may be hesitant to share in some encounters.

In the simulations, we consider the consequences of a gap between p_{same} and p_{dif} for the extent and rate of information spread. For simplicity, we set $p_{same} = 1$ so that anytime someone has the opportunity to tell someone else of the same type, he does so. The p_{dif} term will vary between 1, the instance in which information transmits perfectly to individuals of a different type who are encountered, and 0, the instance in which no information spreads between people of different types even when

TABLE 4 DV: Respondent Heard about the Event

	(1)	(2)	(3)	(4)	(5)
Gender	1.175*** (0.363)	1.044*** (0.374)	1.155*** (0.367)	1.166*** (0.369)	1.113*** (0.371)
Age	0.003 (0.010)	0.005 (0.011)	0.002 (0.011)	0.002 (0.011)	0.005 (0.011)
Educ	0.122 (0.118)	0.179 (0.124)	0.111 (0.119)	0.126 (0.120)	0.150 (0.122)
WallMat	0.238 (0.202)	0.255 (0.211)	0.238 (0.202)	0.260 (0.203)	0.257 (0.207)
Married	0.821 (0.477)	0.765 (0.495)	0.825 (0.481)	0.817 (0.484)	0.828 (0.492)
PartTime	-0.387 (0.348)	-0.342 (0.366)	-0.362 (0.358)	-0.336 (0.359)	-0.417 (0.360)
FullTime	-0.738 (0.538)	-0.733 (0.552)	-0.735 (0.541)	-0.770 (0.541)	-0.872 (0.548)
Anglican	-0.306 (0.431)	-0.164 (0.446)	-0.288 (0.437)	-0.216 (0.443)	-0.037 (0.456)
Pentacostal	0.378 (0.408)	0.278 (0.430)	0.310 (0.414)	0.381 (0.418)	0.515 (0.427)
Degree	0.039 (0.024)	0.027 (0.025)	0.033 (0.024)	0.033 (0.025)	0.033 (0.024)
PHoodTribeHear		1.497*** (0.486)			
PHoodClanHear			0.795 (0.624)		
PHoodUnionClanHear				0.932 (0.551)	
PHoodReligHear					1.215** (0.533)
Constant	-2.661** (1.173)	-3.584*** (1.313)	-2.608** (1.194)	-2.838** (1.220)	-3.266** (1.273)
Observations	198	194	194	194	194
Log Likelihood	-122.495	-114.746	-118.897	-118.249	-117.030
Akaike Inf. Crit.	266.989	253.491	261.795	260.498	258.059

Note: Data are from the Abalang post-event survey.

p < .05, *p < .01.

presented with an opportunity. The figures label the level of distrust, where $distrust := p_{same} - p_{dif}$.¹⁷

Figure 3 shows the result of the above information process on the observed Abalang and Mugana social

¹⁷Although a natural interpretation is trust, the simulations are agnostic about the interpretation of p_{dif} . Instead, $p_{dif} < 1$ could represent understanding, that perhaps an attempt is always made to share information across types when given the opportunity, but sometimes due to differences in language or experience, the message becomes garbled and the intended recipient does not learn the information. Additionally, p_{dif} could indicate credibility on the receiving end—a potential recipient of information may find it less credible when its source is a different type.

networks with types represented as ethnicities, Iteso and Kumam. Each plot shows the outcome of simulated information spread for a different value of distrust with x set to 2.¹⁸ The vertical axis shows the proportion of the village that has heard the news by each time step on the horizontal axis. The curves display the average over

¹⁸In Abalang, the average ratio of the number of people a respondent said she told to the total number of unique social contacts she listed is about 1:3. Consequently, we set the value of x to be 2, which is one-third of respondents' average number of social contacts in the networks on which the simulations are performed. See Section 6.2 of the supporting information for additional detail.

TABLE 5 DV: Passing News of the Event Along to Another

	(1)	(2)	(3)	(4)	(5)
Demographic Controls	<i>insig</i>	<i>insig</i>	<i>insig</i>	<i>insig</i>	<i>insig</i>
Educ	0.542*** (0.210)	0.637*** (0.230)	0.514** (0.216)	0.511** (0.215)	0.482** (0.215)
VisitOth	0.438 (0.322)	0.192 (0.342)	0.254 (0.339)	0.292 (0.342)	0.318 (0.338)
ClanWeek	0.202 (0.253)	0.126 (0.263)	0.105 (0.273)	0.193 (0.270)	0.203 (0.258)
ReligWeek	0.102 (0.294)	0.056 (0.301)	0.018 (0.309)	0.108 (0.297)	0.142 (0.297)
Degree		0.049 (0.034)	0.048 (0.032)	0.047 (0.032)	0.045 (0.032)
PHoodTribe		2.868** (1.396)			
PHoodClan			0.924 (0.748)		
PHoodUnionClan				0.092 (0.615)	
PHoodRelig					-0.717 (0.682)
Constant	-3.288 (2.065)	-6.822*** (2.615)	-3.591 (2.164)	-3.863 (2.127)	-3.244 (2.202)
Observations	115	115	115	115	115
Log Likelihood	-66.451	-62.846	-64.581	-65.353	-64.807
Akaike Inf. Crit.	158.903	155.691	159.161	160.707	159.613

Note: Data are from the Abalang post-event survey, using the subset of respondents who heard about the event.

p < .05, *p < .01.

the simulations, with the rug plot at each time step displayed vertically for each village. On average, news spreads less rapidly in Mugana when p_{dif} decreases by even a small amount relative to p_{same} (i.e., when distrust increases), and the difference in the reach of information between the two villages is increasing in distrust of non-coethnics.

The direction of this result is straightforward: In Abalang, since nearly everyone is the same type, even large reductions in p_{dif} do little to impede the spread of information. Most opportunities to spread information are encounters with people to whom information flows without hesitation. In heterogeneous Mugana, on the other hand, the more hesitation to share with a different type (the smaller p_{dif}), the slower information reaches many others in the community.

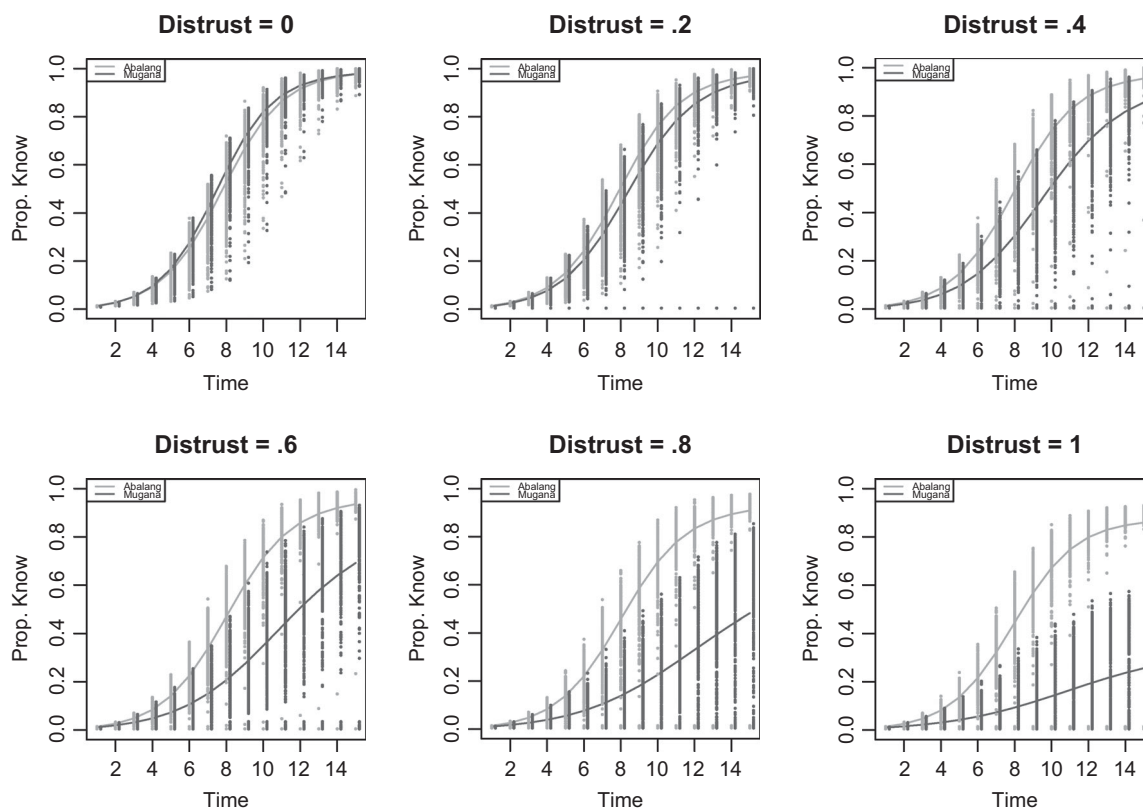
Perhaps more surprising is the magnitude of the result. At the extreme, when $distrust = 1$, after 14 time steps, only about 20% of Mugana would know on average (even though nearly everyone would have known if $distrust = 0$). When $distrust = .6$, an empirically

plausible value,¹⁹ the average difference in information reach is large: In expectation, when over 80% of Abalang has heard, well under half of Mugana has heard. Because information spread has an exponential character—each person can tell x others, who can each tell x others in the next time period, and so on—an unwillingness to pass news to the other type has substantial consequences for the process of informing the village. Section 4 of the supporting information shows that the difference between the two villages persists even when certain tie types are excluded.

Finally, the range of the rug plots in Figure 3 highlights the importance of the choice of seed in heterogeneous villages. As distrust increases, a much larger set of individuals will be seeds from whom information fails to spread widely.

¹⁹While it is difficult to measure p_{dif} accurately, 2005 Afrobarometer survey responses about trust of coethnics versus those from different ethnic groups enable us to back out a speculative empirical estimate for rural Uganda of $p_{dif} = .44$ and $distrust = .56$; see Section 7 of the supporting information for detail.

FIGURE 3 Proportion of Individuals in Abalang and Mugana Who Have Received News



Note: The figure displays the proportion of individuals who have received the news after a certain number of time steps in the simulated information spread with x set to 2 and $distrust$ set to $\{0, .2, .4, .6, .8, 1\}$. The rug plot for each set of simulations is displayed vertically.

In heterogeneous groups, even small barriers to sharing information across ethnic lines compound into large differences in the reach of information. Counterintuitively, dense ties can impede information spread—if those ties are among non-coethnics. If opportunities to spread news to a non-coethnic actually result in the news being shared less often than opportunities to spread news to a coethnic, then the more social ties one has to non-coethnics, all else equal, the greater the chance that an opportunity will fail to spread news. Hence, density in this environment can be counterproductive to information spread since ties to the less trusted crowd out other ties along which information passes more freely.

Forming Dense Networks in Heterogeneous Areas

The last section presented the logic for why denser networks that contain a large proportion of cross-ethnic ties may be worse for information spread than sparser networks. Why, then, would heterogeneous areas have dense networks that feature cross-group ties? We present a

simple network formation model in full in Section 8 of the supporting information, which shows that in the presence of social or economic specialization across ethnic groups, heterogeneous areas would likely form such networks. The intuition is as follows.

Suppose villagers have a type (e.g., ethnicity) and that villagers form social ties in order to fill certain social or economic roles. Suppose also that any villager prefers to fill a role with a same-type tie if possible, but at a minimum, she needs to fill every role and she needs some number of social ties to be at least minimally satisfied. If all villagers are the same type, the optimal bundle of social ties for everyone will trivially be all same-type ties. If instead villagers are different types, the optimal bundle depends on who is available to serve in the needed social or economic roles. If ethnic groups are fully unspecialized, so that for any role there are possible social ties that could fill it in either ethnic group, then the optimal bundle of ties contains no cross-type ties. The more the ethnic groups are specialized, the more the optimal bundle contains cross-group ties, and the larger the optimal bundle of ties is. In short, individuals in ethnically

homogeneous areas form fewer social ties, and individuals in ethnically heterogeneous areas form more social ties so long as ethnic groups have some positive level of specialization; the greater the specialization of roles, the more of these are likely to be cross-group ties.

Scholars of developing countries in general (Horowitz 1985, 108–13) and Africa in particular (e.g., Bates 2000) have noted a common tendency of ethnic division of labor. Recent econometric work has also found evidence for this observation (Michalopoulos 2012). While our data cannot speak directly to specialization in Abalang and Mugana, the supporting information contains indirect evidence, including a high incidence of cross-ethnic ties in Mugana in each of the seven measured networks (Section 4), and some evidence of specialization by occupation (Section 8).

Conclusion

Through a field experiment in two villages of eastern Uganda, this article has sought to bring rich evidence to bear on foundational assumptions about ethnicity's role in information spread. By selecting similar villages and introducing information to randomly selected sources in each, we have documented the wider reach of information in the ethnically homogeneous village compared to the ethnically heterogeneous village. A virtue of the design is that although the information was new, its spread was natural—individuals were left to transmit information in whatever way they typically would through their natural environment. Similar news introduced in similar ways in similar villages had much greater reach in the homogeneous village.

Moreover, we have shown that on every measure, the network of in-person interactions, or “social network,” in the heterogeneous village is in fact denser than the social network in the homogeneous village. This suggests that in this pair of villages, news spread more widely in the homogeneous village *despite having the less dense network*. Greater density may not be the mechanism facilitating the spread of news in homogeneous areas, as previously believed. In fact, we show a mechanism by which denser networks may be worse at spreading information in heterogeneous areas. Even slight hesitancy to share information with non-coethnics can impede dissemination, especially in the presence of cross-group ties. In addition, we show why heterogeneous areas may be likely to develop networks of this form. Therefore, the article presents a new logic of how density relates to both ethnic demography and information dissemination, and it offers a new technique for further experimental inquiry on these topics.

These findings have several implications for related, existing research. While Granovetter's classic work and more recent studies have argued that even “weak ties” between network subgroups can facilitate cohesion and information flows between them (Brown and Konrad 2001; Calvó-Armengol and Jackson 2004; Granovetter 1973), our findings instead suggest the *weakness of weak ties*. If weak ties transmit information at a lower rate—a plausible assumption when the subgroups in question are distinct ethnic groups—then the larger the proportion of social ties occupied by them, the worse they are for information spread overall. This insight also finds support in Varshney (2003), which found that sensitive information was more likely to be shared across members of ethnic groups who were both members of civic associations, and thus had strong ties, than those who were bound by weaker, everyday ties.

Moreover, information spread can be impeded even when there is some trust between ethnic groups; the process is also slowed even if individuals *do* usually share information with the non-coethnics they encounter. In other words, the finding that a large number of cross-group ties can impede the spread of information can be consistent with the findings of contact theory (e.g., Allport 1954) or more recent work by Robinson (2015, 2016) that finds that local ethnic diversity improves trust. Unless interethnic interactions restore trust *to the level of coethnic trust*, then the presence of more cross-group ties may still be unhelpful. The more that trust improves across the ethnic groups, the more helpful the cross-group ties will be at spreading information widely.

References

- Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken. 2015. “Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia.” Technical report, National Bureau of Economic Research.
- Allport, Gordon W. 1954. *The Nature of Prejudice*. Reading, MA: Addison-Wesley.
- Baldassarri, Delia. 2015. “Cooperative Networks: Altruism, Group Solidarity, Reciprocity, and Sanctioning in Ugandan Producer Organizations.” *American Journal of Sociology* 121(2): 355–95.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. “The Diffusion of Microfinance.” *Science* 341(6144): doi:10.1126/science.1236498.
- Bates, Robert. 2000. “Ethnicity and Development in Africa: A Reappraisal.” *American Economic Review* 90(2): 131–34.
- Bates, Robert H. 1983. “Modernization, Ethnic Competition, and the Rationality of Politics in Contemporary Africa.” In *State versus Ethnic Claims: African Policy Dilemmas*, ed.

- Donald Rothchild and Victor A. Olunsorola. Boulder, CO: Westview Press, 152–171.
- Brown, Deborah Wright, and Alison M. Konrad. 2001. “Granovetter Was Right: The Importance of Weak Ties to a Contemporary Job Search.” *Group & Organization Management* 26(4): 434–62.
- Calvó-Armengol, Antoni, and Matthew O. Jackson. 2004. “The Effects of Social Networks on Employment and Inequality.” *American Economic Review* 94(3): 426–54.
- Cederman, Lars-Erik, Nils B. Weidmann, and Kristian Skrede Gleditsch. 2011. “Horizontal Inequalities and Ethnonationalist Civil War: A Global Comparison.” *American Political Science Review* 105(3): 478–95.
- Chandra, Kanchan. 2004. *Why Ethnic Parties Succeed: Patronage and Ethnic Head Counts in India*. Cambridge: Cambridge University Press.
- Conroy-Krutz, Jeffrey. 2013. “Information and Ethnic Politics in Africa.” *British Journal of Political Science* 43(2): 345–73.
- Deutsch, Karl W. 1953. “The Growth of Nations: Some Recurrent Patterns of Political and Social Integration.” *World Politics: A Quarterly Journal of International Relations* 5(2): 168–95.
- Ejdemyr, Simon, Eric Kramon, and Amanda Lea Robinson. 2015. “Segregation, Ethnic Favoritism, and the Strategic Targeting of Local Public Goods.” Presented at the American Political Science Association Annual Conference. <http://bit.ly/2aqMCy3>.
- Fearon, James D., and David D. Laitin. 1996. “Explaining Interethnic Cooperation.” *American Political Science Review* 90(4): 715–35.
- Frere, Karen E. 2010. “The Social Origins of Electoral Volatility in Africa.” *British Journal of Political Science* 40(4): 759–79.
- Granovetter, Mark S. 1973. “The Strength of Weak Ties.” *American Journal of Sociology* 78(6): 1360–80.
- Grossman, Guy, Macartan Humphreys, and Gabriella Sacramone-Lutz. 2014. “‘I wld like u WMP to extend electricity 2 our village’: On Information Technology and Interest Articulation.” *American Political Science Review* 108(3): 688–705.
- Gubler, Joshua R., and Joel Sawat Selway. 2012. “Horizontal Inequality, Crosscutting Cleavages, and Civil War.” *Journal of Conflict Resolution* 56(2): 206–32.
- 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.
- Horowitz, Donald L. 1985. *Ethnic Groups in Conflict*. Berkeley: University of California Press.
- Jackson, Matthew O., Tomas Rodriguez-Barraquer, and Xu Tan. 2012. “Social Capital and Social Quilts: Network Patterns of Favor Exchange.” *American Economic Review* 102(5): 1857–97.
- Jackson, Matthew O., and Brian W. Rogers. 2007. “Relating Network Structure to Diffusion Properties through Stochastic Dominance.” *BE Journal of Theoretical Economics* 7(1): 1–16.
- Kasara, Kimuli. 2015. “Does Local Ethnic Segregation Lead to Violence? Evidence from Kenya.” *Quarterly Journal of Political Science*. Forthcoming. <http://bit.ly/2aqxWVN>.
- La Ferrara, Eliana. 2002. “Self-Help Groups and Income Generation in the Informal Settlements of Nairobi.” *Journal of African Economics* 11(1): 61–89.
- Larson, Jennifer M., and Janet I. Lewis. 2016. “Rumors, Kinship Networks, and Rebel Group Formation.” Presented at Political Institutions and Economic Policy Conference, Harvard University. <http://bit.ly/2atLDkn>.
- Lee, Sang Hoon, Pan-Jun Kim, and Hawoong Jeong. 2006. “Statistical Properties of Sampled Networks.” *Physical Review E* 73(1): 016102.
- Michalopoulos, Stelios. 2012. “The Origins of Ethnolinguistic Diversity.” *American Economic Review* 102(4): 1508–39.
- Miguel, Edward, and Mary Kay Gugerty. 2005. “Ethnic Diversity, Social Sanctions, and Public Goods in Kenya.” *Journal of Public Economics* 89(11–12): 2325–68.
- Parkinson, Sarah Elizabeth. 2013. “Organizing Rebellion: Re-thinking High-Risk Mobilization and Social Networks in War.” *American Political Science Review* 107(3): 418–32.
- Posner, Daniel N. 2005. *Institutions and Ethnic Politics in Africa*. Cambridge: Cambridge University Press.
- Putnam, Robert D. 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster.
- Robinson, Amanda. 2015. “Ethnic Diversity, Segregation, and Ethnocentric Trust in Africa.” Unpublished manuscript, Ohio State University. <http://bit.ly/2aqxGX2>.
- Robinson, Amanda Lea. 2016. “Nationalism and Ethnic-Based Trust: Evidence from an African Border Region.” *Comparative Political Studies*. Advance online publication. doi:0010414016628269.
- Toft, Monica Duffy. 2003. *The Geography of Ethnic Violence: Identity, Interests, and the Indivisibility of Territory*. Princeton, NJ: Princeton University Press.
- Varshney, Ashutosh. 2003. *Ethnic Conflict and Civic Life: Hindus and Muslims in India*. New Haven, CT: Yale University Press.
- Weidmann, Nils. 2009. “Geography as Motivation and Opportunity: Group Concentration and Ethnic Conflict.” *Journal of Conflict Resolution* 53(4): 526–43.
- World Bank. 2014. “World Bank Development Indicators: Agriculture and Rural Development.” Washington, DC: World Bank. <http://data.worldbank.org/topic/agriculture-and-rural-development>.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher’s website:

- Intervention Design
- Networks Data
- Cross-Group Ties
- Different Types of Ties
- Accounting for Unsampled Ties
- Additional Village Information
- Coethnic Trust in Rural Uganda
- Network Formation Model