Reducing Prejudice Towards Refugees in Uganda: Evidence that Social Networks Influence Attitude Change*

Jennifer M. Larson† Janet I. Lewis‡

April 4, 2023

Word Count: 11,483

Abstract

Interventions aimed at reducing prejudice towards refugees have shown promise in industrialized countries. However, the vast majority of refugees are in developing countries. Moreover, while these interventions focus on individual attitude change, attitudes often do not shift in isolation; people are embedded in rich social networks. We conducted a field experiment in northwestern Uganda (host to over a million refugees) and find that perspective-taking warmed individual attitudes there in the short-term. We also find that the treatment effect spills over from treated households to control ones along social ties, that spillovers can be positive or negative depending on the source, and that peoples' attitudes change based on informal conversations with others in the network after the treatment. The findings show the importance of understanding the social process that can reinforce or unravel individual-level attitude change towards refugees; it appears essential to designing interventions with a lasting effect on attitudes.

*We are grateful to Anthony Kamwesigye, Kyla Longman, Emmanuel John Osuta, and Innovations for Poverty Action (IPA)'s Uganda office for their critical role in carrying out data collection, and to the study participants for generously giving their time. We received helpful feedback from Scovia Aweko and participants of the Overcoming Prejudice Against Immigrant Minorities conference, hosted by the Identity & Conflict Lab at the University of Pennsylvania, as well as from John Marshall and participants of the Empirical Models of Political Economy conference at Caltech. We gratefully acknowledge funding for this project from an IPA Peace and Recovery Program grant and George Washington University.

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

‡Associate Professor of Political Science, George Washington University, janetilewis@gwu.edu.

1 Introduction

Each year, millions of people are forcibly displaced to countries outside their birth country. How refugees fare in a new location depends in part on the attitudes of the people already living there. Existing research has explored ways to induce greater warmth in attitudes towards groups of others, with some success. For instance, respondents who participate in a conversation in which they are induced to take the perspective of refugees and other outgroups tend to feel more positively towards them (Broockman and Kalla, 2016; Adida, Lo and Platas, 2018; Simonovits, Kezdi and Kardos, 2018; Kalla and Broockman, 2020; Williamson et al., 2021).

However, these studies have occurred mostly in industrialized countries, whereas the vast majority (85%) of the world's refugee population is in developing countries. Furthermore, while scholars and organizations typically administer and measure effects of interventions like these exclusively at the individual level, individuals' beliefs and attitudes are not developed, or changed, in isolation. Individuals are embedded in rich social networks. Even if a person's mind was changed during an intervention, what happens once she returns to her usual social life? Will friends and family support the change, push against it, or will they themselves be persuaded by it? The durability of an intervention's effect may well depend on this "social processing" that occurs afterwards.

We conducted a field experiment that addresses both concerns. The experiment assesses the effectiveness of an intervention aimed at shifting a host population's attitudes towards refugees in four villages in the West Nile region of Uganda. Uganda is an important developing country setting for studying host-refugee relations since it hosts the world's third largest refugee population (UNHCR, 2022). This region borders South Sudan and the Democratic Republic of the Congo, the origin of over 80% of Uganda's refugees. Furthermore, as a departure from previous studies, our design not only measures the effectiveness of the perspective taking intervention immediately and in the longer term, but it also directly measures the village social networks and the social processing that occurs within them after the intervention.

Specifically, our research team conducted a baseline survey of all village households that measured attitudes towards refugees, household characteristics, and the interactions that comprise the village social networks.[1] In a randomly chosen half of all households, a perspective-taking treatment was also administered. Three weeks later, the team followed up with an endline survey of all households in each village which measured attitudes again and also probed experiences with social processing.

We find that in all four villages, perspective taking did indeed change individuals' short-term attitudes to be warmer towards refugees on average. As expected, a treatment that invites a respondent to meaningfully consider the experiences of refugees and to discuss their views non-judgmentally can lead to greater warmth in a developing country setting like Uganda. We also find that in this setting, as is typical for anti-bias interventions, some of the warming of attitudes erodes over the course of the three week interim on average. However, we also show that this average erosion belies a wide variety of responses and a rich process that took place in the interim.

We present evidence that social processing is indeed present (and prevalent). After our intervention, respondents spoke with one another, especially their peers in the village social networks. They were speaking about refugees and doing so more often than usual. In fact, treating some in the village led to both the treated households *and the control* talking about refugees more often than usual. By this mechanism, the intervention spilled over onto control households and further shaped the reactions of the treated.

Our unique design allows us to connect this social processing to changes in attitudes over time. Our findings suggest that the effect of the intervention evolved in response to these conversations with peers afterwards. This did result in an on-average erosion of the gains in warmth for the treated, but the individual changes can be better explained

[1] In developing countries, particularly in rural Sub-Saharan Africa, the relevant networks are highly local; a great deal of trusted news travels via informal, local word-of-mouth networks (Banerjee et al., 2013; Larson and Lewis, 2017; Larson, Lewis and Rodríguez, 2022).

as movement towards the attitudes of their social ties in the village network. Intriguingly, control households were also shaped by this social processing that followed the intervention. Control households' attitudes moved in response to the treated households receiving the intervention and then talking about it with their social ties. Control attitudes also warmed on average, and also moved towards the attitudes of their social network ties. Perhaps most intriguingly, the control households' average increase in warmth and movement towards social ties is even larger than that of the treated households.

We further show that spillovers from treated respondents do not occur uniformly; some treated respondents generate positive spillovers while some generate negative ones. Ultimately it is not merely receiving treatment, but how one reacts to the treatment received, that has important consequences for the attitudes of those near a treated person in the network. We find evidence that those who were especially persuaded by the treatment generated positive spillovers whereas those who reacted most negatively to the treatment (fortunately there were not many such people) generated negative spillovers through the social network.

These findings strongly suggest that to design interventions that can lead to enduring improvements in attitudes towards refugees in rural, developing country contexts, we need a better understanding of the social processes that can reinforce or unravel individual-level attitude change. This study thus serves as a proof of concept that this topic is both important and feasible to study, even in rural, low-income contexts where word-of-mouth (rather than online) networks serve as the primary means of communication and social vetting.

This paper just scratches the surface of what would be valuable to learn about social processing of individual-level interventions, raising exciting open questions about how this process works to determine whether the effect will be durable. For example, what are the individual attributes that drive some people to be better at moving network neighbors towards more pro-refugee attitudes after being treated? Are these different from attributes that make people more influential at shifting people towards more anti-refugee attitudes? Are some village network structures more amenable to spread of pro-refugee attitudes than others? We discuss promising ways to continue to advance this agenda in the conclusion.

2 Prejudice Reduction through Social Networks?

This paper seeks to contribute to two dynamic, distinct areas of social science research: One on prejudice reduction towards refugees, and another about information flows and social processing through networks. The former rigorously examines anti-refugee prejudice and other social barriers to refugee integration in the U.S. and Europe (e.g. Bansak et al., 2018; Hainmueller and Hopkins, 2015; Adida, Lo and Platas, 2018; Choi, Poertner and Sambanis, 2019; Hopkins, Sides and Citrin, 2019*a*; Williamson et al., 2021), but less systematic work has done so in developing countries (see Audette, Horowitz and Michelitch, 2020), where refugee populations are much larger and more sizeable relative to host populations (Blair et al., 2021). The more constrained resource environment in such contexts may exacerbate tensions; for example, in Sub-Saharan Africa, living near refugees drives lower levels of interpersonal trust and less support for inclusionary citizenship rules (Zhou, 2019).[2] This paper builds on a smaller body of work on prejudice and prejudice reduction in developing country settings (e.g. Paluck, 2010; Burns, Corno and Ferrara, 2018; Rosenzweig and Zhou, 2021) by studying the effectiveness of an intervention aimed at measuring and improving Ugandans' attitudes and behavior towards South Sudanese refugees.

Because individuals do not usually process attitude changes about outgroups in a vacuum – but instead engage in social processes of exchanging information and ideas among trusted sources – we also seek to directly study the social processing that individual-level interventions spark. Here we draw on insights from a broad range of work: In countless theories of political and social behavior related to intergroup conflict and mobilization, the outcome of interest depends on what people know and believe – which is often informed

[2]Note however that Zhou and Grossman (2022) finds, using evidence from Uganda, that higher public goods provision near refugee settlements can mitigate backlash against pre-refugee policies.

by what others they trust know and believe. For example, rumors about outgroups that spread too widely can cause intergroup tension and ultimately spark conflict (Varshney, 2003), while wide-reaching gossip about outgroup misbehavior can stave off conflict (Larson, 2017; Fearon and Laitin, 1996). Protests depend not only on what people know about an upcoming protest, but also their assessment of other peoples' expectations about it (Chwe, 2000; Siegel, 2009). Trusted social ties also appear to play an essential role in vetting narratives as people decide what to do amidst the uncertainty of conflict, according to in-depth analyses of contexts ranging from Uganda (Lewis, 2020) to Abkhazia (Shesterinina, 2021), Syria (Schon, 2021), and the Philippines and Thailand (Greenhill and Oppenheim, 2017).

The relationships that comprise a social network are important sources of trusted communication in any context, especially in the developing world where informal ties often serve as a vehicle for information and services that elsewhere could be provided through other channels (Banerjee et al., 2013). Theory that seeks to explain the spread of something through a network, "diffusion," borrows from epidemiology and views individuals "infected" with ideas or behavior as contagious to network neighbors (Newman, 2000; Jackson, Rogers et al., 2007). Hence, collective behavioral changes can emerge when enough prospective participants are infected with the desire to change by enough other participants, resulting in cascades of large-scale collective action (Marwell, Oliver and Prahl, 1988; Gould, 1993). Variants acknowledge that behavior may spread according to relatively more complicated rules than a disease, requiring repeated exposure, or exposure to enough others who are infected (Valente, 1996; Chwe, 1999; Dodds and Watts, 2004; Centola and Macy, 2007; Chiang et al., 2007; Centola, 2013). According to these theories, if a person is sufficiently exposed to enough others who will undertake a certain action, the person will take that action too.

We argue that reversal of prejudicial attitudes may operate similarly; changing one person's mind may be difficult if all of that person's peers still hold prejudicial beliefs. Likewise, as the contagion literature cited above suggests, if many of one's peers have begun to reconsider their own prejudices, then one may be encouraged to do so as well. Therefore, understanding how to reduce prejudice toward refugees at the level of a group requires understanding how people in the group do or do not talk about their views, and do or do not account for their peers' views when thinking about refugees.

Despite the likelihood that altering individual beliefs and attitudes might depend crucially on perspectives flowing through a person's social networks, systematic empirical research into this process is rare, especially in developing countries. Although the number of studies that carefully measure social networks with enough detail to potentially observe social processing has grown (e.g. Ferrali et al., 2020; Arias et al., 2019; Eubank et al., 2019; Atwell and Nathan, 2021), factors such as the onerous demands of network elicitation, combined with the difficulty of measuring information flows, has resulted in studies of individual-level treatments to change attitudes and beliefs that are divorced from studies of social processing in networks. This paper moves beyond prior studies of prejudice reduction by not only assessing the effectiveness of an intervention that has already proven successful in another context, but also examining the social process of prejudice reduction by measuring the social interactions that occur between the initial and eventual attitudes of the treated as well as "spillover" of the treatment onto control individuals. Doing so enables us to show the likely importance of social processing, especially for assessing longer term effects of individual-level interventions that aim to mitigate prejudice.

3 Study Site and Design

3.1 West Nile region of Uganda

We carried out this study in the West Nile region of northwestern Uganda, which borders South Sudan and the Democratic Republic of Congo (DRC). Uganda is an important context for understanding refugee-host country relations for several reasons. First, it hosts the largest refugee community in Africa; Uganda is home to about 900,000 refugees from South

Sudan, most of which are concentrated in West Nile (UNHCR Uganda Population Dashboard).[3] Uganda also has a strong national commitment to hosting refugees that is reflected in its progressive immigration policies, which include the right to education, employment, and plots of land for cultivation (Blair et al., 2021). Still, its population faces challenges absorbing these refugees that are common to host countries. Relations are often strained between the refugee population and Ugandans, some of who perceive refugees as unwelcome competition for local resources and services (World Vision, 2018; UNHCR, 2018; Search For Common Ground, 2021). Proximity to refugee settlements in Uganda is associated with higher levels of fear of crime, as well as higher electoral support for the incumbent President (Zhou and Grossman, 2022), whose party has been increasingly implicated in democratic erosion. In addition to the substantive importance of Uganda, several past studies have demonstrated the feasibility of collecting village network data there (e.g. Larson and Lewis, 2017; Ferrali et al., 2020; Eubank et al., 2021).

As described below, baseline data from our survey confirms that sizeable minorities of Ugandans in our West Nile study villages hold exclusionary attitudes towards refugees. And while refugee inflows do not typically lead to large-scale violence (Shaver and Zhou, 2021), concerning anecdotes indicate social tension and the potential for intergroup violence in West Nile. For example, when one of the authors recently asked an NGO leader working in West Nile about current Ugandan-South Sudanese relations in general, he responded that rumors are circulating in his village that South Sudanese people had beaten an ill Ugandan, leading to his death. These rumors, he said, are "fed by word of mouth" and made young people there "feel agitation" and "want revenge" against South Sudanese people. He also stated that some of the coexistence dialogue groups he leads between South Sudanese refugees and Ugandan nationals have recently broken out into physical, intergroup attacks.[4] The most severe recent case of intergroup violence was a 2020 attack on South Sudanese

[3]Refugee settlements also exist in western Uganda; most of the refugees in these settlements are from DRC. The vast majority (over 90%) of refugees in Uganda live separately from the host population, in refugee settlements.

[4]Author conversation via Skype with Pax Sakari, Director of Rural Initiative for Community Empowerment (RICE)-Uganda (January 2022).

refugees that left over 10 dead and 15 homes destroyed, resulting in police and military deployment to the area in order to prevent escalation.[5]

3.2 Study Design

We carried out our study in Uganda from February to August 2021. In each of the four study villages, a randomly selected set[6] of households received a perspective-taking treatment along with a survey to learn beliefs, attitudes, demographics, and social networks. The remaining control households were only surveyed. Treatment and control households were surveyed again approximately two weeks later.

Our intervention was a brief (roughly 10-15 minute) conversation in which the visitor non-judgmentally exchanges narratives about refugees with the individual, and encourages them to take the perspective of refugees. We modeled our intervention on Broockman and Kalla (2016)'s "perspective-taking" intervention because it has strong evidence of effectiveness, and because the intervention's simplicity and brevity make it easily scalable.[7] Further, evidence from Adida, Lo and Platas (2018)'s experiment shows the effectiveness of a similar perspective-taking exercise to decreasing prejudice towards refugees in the U.S.

Specifically, we shared a narrative about a single South Sudanese refugee's life and her perspective, and reminded the respondent that this refugee is part of a much larger group now residing in Uganda. While the structure of this intervention allows for natural conversation, it entails key components of the treatment including creating a non-judgmental

[5]Samuel Okiror, "Uganda calls in troops as violence flares between refugees and locals." *The Guardian* September 15, 2020.

[6]50% of households in villages 1 and 2, about 60% in 3 and 4.

[7]Kalla and Broockman (forthcoming) distinguishes among and tests the component pieces of the Broockman and Kalla (2016) intervention. Kalla and Broockman (forthcoming) finds that omitting the "analogic perspective-taking" and "vicarious perspective-giving" components does not diminish effects, and that interventions employing only "perspective-getting" narratives durably reduce exclusionary attitudes. For simplicity and in keeping with the rest of the literature, we term our intervention – which included all three components – perspective-taking. Kalla and Broockman (2020) tested the intervention in seven locations in the United States and found that it successfully reduced exclusionary attitudes towards transgender people and unauthorized immigrants for at least four months. Kalla and Broockman (2020) also show the effectiveness of several closely-related interventions that involve non-judgmental sharing of narratives about outgroups.

context for discussion, encouraging active processing, acknowledging contrary perspectives, and addressing concerns that the respondent surfaces about refugees. Additional detail on our intervention is in the appendix.

Overall, this study design seeks to capture contexts in which there is a stimulus in a community that prompts local discussions of refugees; this could be striking news about refugees that only a subset of the population learns, or an anti-prejudice program that only a portion of the villages' population receives.

3.3 Ethics

In carrying out the study, we took several steps to mitigate any potential harm to respondents and other community members. Since the study occurred during the global COVID-19 pandemic, we consulted extensively and regularly with local officials and public health information to ensure that in-person surveying only occurred when COVID transmission was low in the localities where we conducted the survey. All surveys were conducted either via phone or in-person outdoors, with the enumerator wearing a mask. Research team members offered masks to all respondents, and maintained social distance from them. Before requesting consent to participate in the survey, in addition to describing the study, the enumerators also provided information about COVID and best practices to prevent contracting it.

Additionally, we ensured that the information we presented about refugees was accurate, and portrayed refugees in a positive light, and that all survey data was kept confidential and encrypted. Participation in the survey and each component question was voluntary; we carefully trained enumerators to request informed consent. We conducted the study with prior approvals from the authors' university Institutional Review Boards, from Uganda's National Council on Science and Technology, from a local Ugandan IRB (Mildmay Uganda Research Center) and from the relevant district-level officials. In November 2022, we shared the study's preliminary findings with the leadership of our four study villages, and several villagers.

3.4 Study villages and Issue Salience

All four study villages are in northwestern Uganda.[8] They are similar in size (about 100 - 150 households each) and the average age of respondents, though vary considerably in other demographics such as levels of education, primary occupation, and religious affiliation. Table 1 reports average values of these features for each village. Villages 3 and 4 are religiously homogeneous communities with a strong majority of farmers with low levels of formal education. Villages 1 and 2 are relatively more religiously diverse, have more traders and other non-farming occupations, and higher levels of education.

	Vlg 1	Vlg 2	Vlg 3	Vlg 4	All
Age	35	38	39	40	38
Protestant	0.45	0.12	0.00	0.00	0.13
Catholic	0.38	0.84	0.00	0.92	0.51
Muslim	0.14	0.03	0.99	0.06	0.34
Farmer	0.24	0.49	0.83	0.76	0.60
Trader	0.22	0.24	0.04	0.13	0.15
No Educ	0.03	0.07	0.33	0.13	0.15
Primary Educ	0.29	0.59	0.53	0.68	0.53
Secondary Educ	0.27	0.17	0.12	0.12	0.17
College Educ	0.41	0.16	0.02	0.07	0.16
Lived > 5yrs	0.64	0.73	0.81	0.83	0.76
Baseline hhs	127	98	146	150	521
Endline hhs	116	85	142	145	488

Table 1: Average age; proportion of respondents who identify as Protestant, Catholic, or Muslim; proportion of respondents who report farmer or trader as their occupation; proportion of respondents who report receiving no education or at least some primary, secondary, or college education; and the proportion of respondents who have lived in the village for more than five years.

The proximity of these villages to borders with refugee-sending countries and, consequently, refugee settlements, makes it no surprise that refugees is a salient issue.[9] Table 2

[8] All are in West Nile region. Village 1 is in Arua district; village 2 is in Maracha district; villages 3 and 4 are in Yumbe district. Uganda has over 110 districts and a population of over 40 million.

[9] Refugees live in settlements within about 50 kilometers from the border and are provided land for subsistence farming. Our study villages were 30-60 minutes (by walking or public transit) to at least one

shows that many of our respondents were once refugees themselves.[10] Most of the respondents have personally met a refugee, with the highest frequency in Village 1 where 76% of respondents have done so.

	Vlg 1	Vlg 2	Vlg 3	Vlg 4	All
Has been refugee	0.24	0.30	0.34	0.40	0.32
Has met refugee	0.76	0.47	0.64	0.57	0.62
Num times came up last week	2.58	1.79	0.59	0.51	1.28
Heard from friend or family	0.69	0.55	0.24	0.14	0.38
Heard from radio	0.39	0.34	0.17	0.12	0.24
Heard from newspaper	0.06	0.03	0.00	0.01	0.02
Heard from TV	0.05	0.03	0.01	0.01	0.02
Heard from other	0.02	0.00	0.00	0.00	0.01

Table 2: Proportion of respondents in each village who were themselves a refugee at one time and who have met a refugee; the average number of times respondents reported that the issue of refugees came up for them in the previous week; and the proportion who reported that they had heard about refugees the past week from each source/ medium.

The topic of refugees also comes up regularly for many of our respondents, with varying frequency across the villages. In villages 3 and 4, just over one out of every two people said the issue came up in the past week; in villages 1 and 2, respondents reported the issue arising more often than weekly. Table 2 shows the breakdown of sources and media in which the issue arises. Across the board, interpersonal connections are the most prevalent source of refugee information, with radio taking second place. A context in which some information is learned from third party resources and much is learned from personal contacts is one with a lot of room for word-of-mouth sharing and processing.

settlement. The refugees are not integrated members of these villages, though interactions with refugees would be typical at water collection areas, markets, and sometimes in hospitals and schools, hence the high prevalence of having met a refugee among our respondents.

[10] These respondents fled Ugandan violence in the early 1980s, across the border into South Sudan and DRC (then Zaire), remaining for about a decade before returning.

4 Treatment Warms Individuals' Attitudes Towards Refugees on Average

Our primary dependent variable is an index of attitudes towards refugees that aggregates responses to six survey questions. The questions were designed to replicate survey instruments in Hopkins, Sides and Citrin (2019*b*) and Kalla and Broockman (2020), lightly modified to suit the Ugandan refugee context. Each asked the respondent to use a five-point scale to react to the statements:

- I would have no problem with refugees from foreign countries coming and living in my village.
- I believe that refugees just wouldn't fit socially in my community here in [name of village].
- I believe that refugees would be too large a burden on the resources of my community.
- I believe that refugees hold the same values as my community.
- Do you think the agricultural land set aside for use by refugees in Uganda to use for growing should be: [scale ranging from increased a lot to decreased a lot]?
- How likely is it that refugees will threaten the way of life in your community? [Scale ranging from very unlikely to very likely].

4.1 Baseline Attitudes

Figure 1 shows the baseline responses to each of the six questions for all respondents in the four villages, rescaled so the answer corresponding to the number 5 is always the most pro-refugee answer.[11] Baseline attitudes contain a fair bit of variation on all of the constituent questions, though over 60% say they would have no problem with refugees coming and living in their community.

Our pooled and village-level analyses use an index constructed from the sum of the rescaled responses to these six questions as the dependent variable, so the pro-refugee index ranges from 6 (the least pro-refugee answer to all six questions was selected) to 30

[11] The appendix shows these baseline attitudes also broken apart by village.

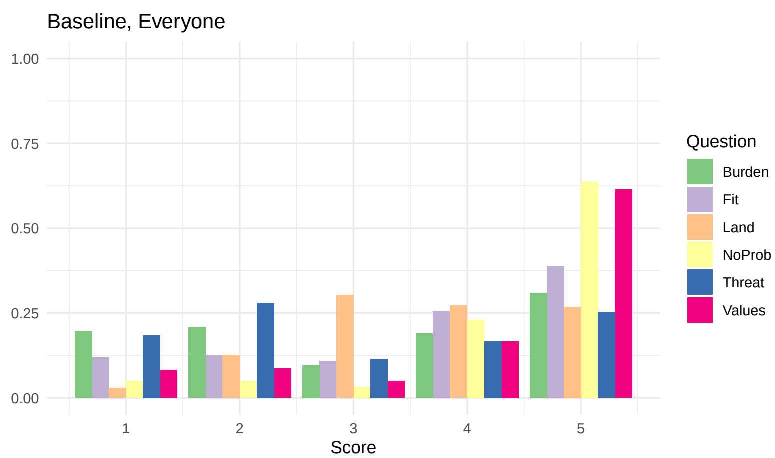

Figure 1: Baseline attitudes broken down by question for all villagers. Answers to each question range from least pro-refugee on the left to most pro-refugee on the right.

(the most pro-refugee answer to all six questions was selected). We refer to this index as the pro-refugee score, with higher values indicating warmer attitudes, and lower values colder ones.

4.2 Individual Short-Term Response to Treatment

On average, respondents' pro-refugee score increased 2.5 points in immediate response to treatment.

	V1	V2	V3	V4	All
Pro-ref score, bl	21.4	20.0	24.3	23.3	22.6
Pro-ref score, bl2	23.3	23.3	26.7	25.8	25.1
Short-term change	**1.9**	**3.3**	**2.3**	**2.5**	**2.5**
% s.t. change > 0	59%	70%	65%	70%	66%
% s.t. change < 0	17%	14%	14%	12%	14%
% s.t. change $= 0$	24%	16%	22%	18%	20%
n	59	50	88	92	289

Table 3: Average treatment effect on the treated separated by village and pooled.

After participating in a non-judgmental conversation in which respondents were invited to take the perspective of a South Sudanese refugee, respondents' answers to the six questions warmed by 2.5 points on the 6 to 30 point scale. This amount is over ten percent of the range of the scale, and is the equivalent of moving from the most negative to strictly positive in answer to one of the six questions. Table 3 further shows that this average is similar across villages, and that strong majorities of individuals moved warmer to comprise this average.[12] That individuals respond to a perspective-taking treatment by reporting substantially warmer attitudes towards refugees is an important confirmation that this style of treatment can also be effective in the short-term in a developing-country setting.

4.3 Individual Long-Term Response to Treatment

The effect of treatment on the pro-refugee score followed the same pattern in all four villages over time: those who received treatment immediately became more positive in their attitudes towards refugees on average. Then, after two to three weeks elapsed, they remained more positive towards refugees compared to their baseline attitudes, but the average increase was somewhat attenuated. Figure 2 shows this pattern separated out by village as well as pooled. It displays the mean pro-refugee score for the treated in the first baseline measure, the second, post-treatment baseline measure, and in the endline. The horizontal bars indicate the width of the standard error of each of the means. Stars label the baseline 2 and endline points to indicate the statistical precision of a difference in means t-test when compared with the baseline.

Although villages differ in how pro-refugee they start at baseline and in the magnitude of the gains, they all exhibit the same pattern: treatment causes the treated to hold warmer attitudes towards refugees, though some of the warmth appears to fade on average over time.

The apparent fading of warmth is not driven by any one constituent part of our attitude index that serves as the dependent variable. Figure 3 shows the change in pro-refugee score broken down into the index's six constituent questions. Each subfigure shows the proportion

[12] We note that some respondents did respond by moving more negative in their attitudes towards refugees. Only 14% of respondents did so. The possibility of backsliding is real, and has been the dominant effect in other settings.

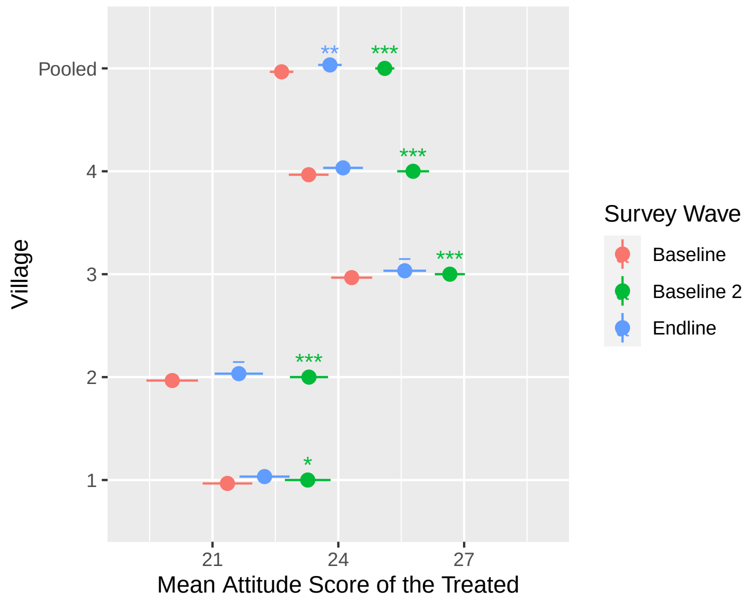

Figure 2: Mean attitude score of the treated and the standard error of the mean in each of the survey waves, pooled and separated by village. Results of two-sided difference in means t-test indicated by label, comparing baseline 2 to baseline 1 and endline to baseline 1. Symbol key $-p < .1$; $*p < .05$; $**p < .01$;$***p < .001$.

of treated respondents who chose each response for each of the six attitude questions, again rescaled so that larger numbers indicate warmer attitudes. The top left subfigure shows the baseline responses of the respondents who were assigned to treatment; the top right figure shows their responses at the end of the baseline survey after receiving the treatment; and the bottom shows their responses during the endline survey.

The figure shows that changes in treated respondents' pro-refugee scores in response to treatment were not concentrated in a single question. Some questions did see more movement than others. Almost 80% of treated respondents ended the baseline by choosing the most pro-refugee answer to "I would have no problem with refugees from foreign countries coming and living in my village." The same is true for long-term changes. By the endline, attenuation was also spread across the questions, though attitudes about land allocation and the possibility of refugees being a burden on resources saw the largest back-slide.

Language like "fades," "attenuates," and "wears off" is convenient to describe what hap-

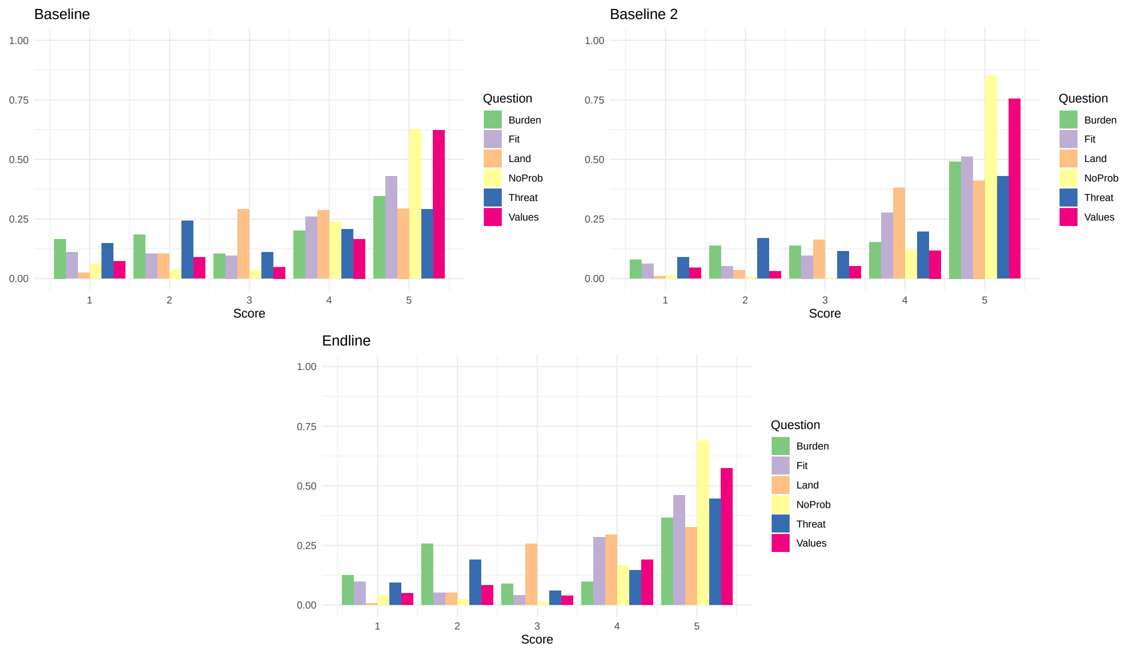

Figure 3: Baseline attitude score of the treated in the baseline before treatment (top left), after treatment (top right), and in the endline (bottom), pooled across villages and separated into the six constituent questions.

pens to the average effect over time. However, a closer look starts to reveal that this language might not fully capture the richness of the longer-term response. Figure 4 presents the same information as Figure 2, this time in terms of the *change* in score. The short-term change shown in Figure 4 is the difference between the treated respondents' pro-refugee score at the end of the baseline survey (after treatment) and their score measured at the beginning of that survey (before treatment).[13] Long-term change is the difference between the treated respondents' score in the endline survey (two to three weeks after treatment) and their initial baseline score. In this view, respondents in the center of the horizontal axis saw no change in their pro-refugee score. Respondents to the right saw warming in their score, and to the left saw cooling.

Figure 4 shows once again that the short-term and long-term response was greater warmth on average. However, it also makes clear that that average is comprised of substantial heterogeneity in individual responses. Most responded by increasing warmth in

[13]The same figure separated by village can be found in the appendix.

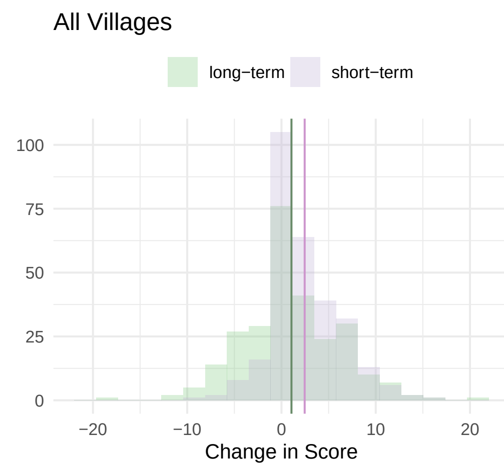

Figure 4: Change in attitude score of the treated in the short- and long-term, pooled across villages.

both the short- and long-term, by on average 2.5 and 1.1 point respectively, but the range in responses is wide. Some respondents were moved by more than ten points. Moreover, although the increase and attenuation is true on average, it is certainly not the case that all individuals became warmer in the short-term and then somewhat less warm in the long-term. In fact, all combinations of score changes are present in the data.

Figure 5 displays the responses of the treated another way, plotting the long-term change in attitudes against the short-term change. If classic individual attenuation were the primary explanation, we should observe the preponderance of points in the highlighted wedge. Points in this region represent respondents who responded positively to treatment (are on the right side of the plot) and remained positive (top half) but less so (beneath the 45-degree line). Indeed, many of our respondents are represented in this region. But many lie elsewhere on the plot. Some became warm and then got warmer in the long-term. Some became cooler immediately but then moved warmer in the long-term. Some got warm but then cooled substantially. Treated individuals responded as expected on average, but that average does not tidily describe even the bulk of treated respondents well. Understanding the variance in the response of the treated is one of two goals of Section 5. The other pertains to the control respondents.

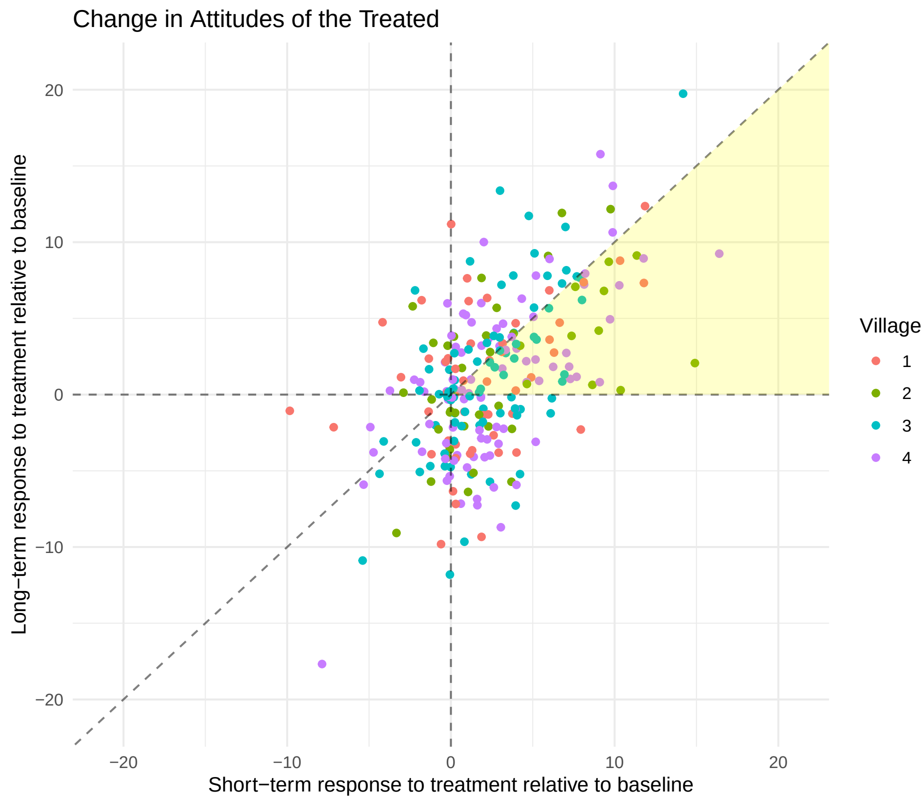

Figure 5: Change in attitude score of the treated; short-term change on the horizontal axis, long-term change on the vertical. If individual attitude changes were simply attenuating or wearing off, we should observe respondents primarily in the highlighted wedge, where short-term change is positive and long-term change is as well but with smaller magnitude.

4.4 Control Respondents and Spillover Effects

Treated respondents reacted to the treatment by reporting warmer attitudes towards refugees immediately, and their attitudes remained warmer than they started on average but less warm than the initial spike in response to treatment. Were treated individuals the only ones who were ultimately affected by this treatment? In the short-term the answer is yes, by design.[14] The treatment was administered privately; no other respondents were present. The second baseline measure of attitudes was collected immediately, without a chance to talk to anyone other than the enumerator.

[14] In other words, the stable unit treatment value assumption (SUTVA) holds in the short-term but might not in the long term, again by design.

In the long-term, the answer is less clear. Respondents had two to three weeks to live their lives between the baseline and the endline. In this window of time, respondents could have talked with others about their experience. Through talking, they may have exposed others, including individuals in the control condition, to a sort of secondary treatment.

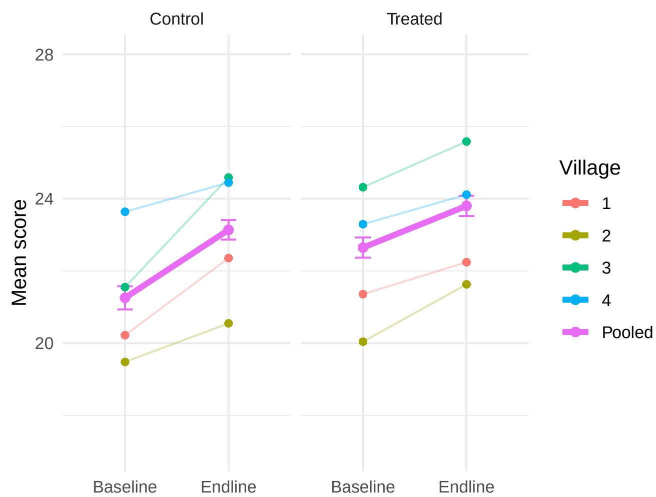

Figure 6: Change in mean attitude score of the treated compared to control in the long-term, pooled and separated by village. Bars show standard error of the pooled mean.

Figure 6 hints that the effect of treatment was not confined to the treated. The right panel contrasts the baseline and endline average pro-refugee score for the treated, displaying in another way the point made above that the treated experienced a warming of attitudes even in the long-term. The left panel shows, intriguingly, that the same pattern holds for the control group. In the two to three weeks between the baseline surveys in which some other individuals received a perspective taking treatment and the endline measure of attitudes, individuals in the *control* condition also became warmer towards refugees on average. This opens the possibility that the treated did not keep the treatment to themselves.

5 Social Processing

We consider the possibility of social processing of two types: (1) the treated, in forming their ultimate endline attitudes, talked to their social ties and reacted in some way, and (2) the control, in talking to the treated and others, also updated their attitudes in reaction to their social ties in some way.

Our design rules out the possibility of active social processing in the short-term. No social interactions occur between treatment and the measure of attitudes at the end of the baseline. In the two to three week interim between the baseline and endline surveys, though, respondents could continue a strictly individual response to treatment, or they could reach out to other people they trust to discuss their reaction, try out new positions, learn the impressions of others, and make a judgment about the socially correct response.

In order to evaluate whether respondents' reactions are consistent with this kind of social processing, we need to identify the set of other people that they might engage with to do so. To that end, we measure household social networks in each village.

5.1 Village Social Networks

In the baseline survey before measuring attitudes towards refugees, we elicited four types of social network ties among villagers. Each respondent was asked to name up to five adults in response to each of the following name generator prompts:

- the adult villagers whose homes you visit in a typical week who don't live in your household;
- the adult villagers who you share a meal with in a typical week who don't live in your household;
- the adult villagers who you go to if you need to borrow money who don't live in your household; and
- when you hear news or rumors that seem surprising or unusual, the adult villagers outside your household that you typically first turn to to chat about it.

These ties are intended to capture the kinds of interactions indicative of relationships that might be relevant for socially processing new information relevant to attitudes towards outgroups (Larson and Lewis, 2020). We use responses to these questions to construct a household network for each village. An undirected link is present between two households in a village's network if a member of one household listed a member of the other household in response to at least one of the four name generator questions. Figure 7 shows the resulting networks measured for each village.

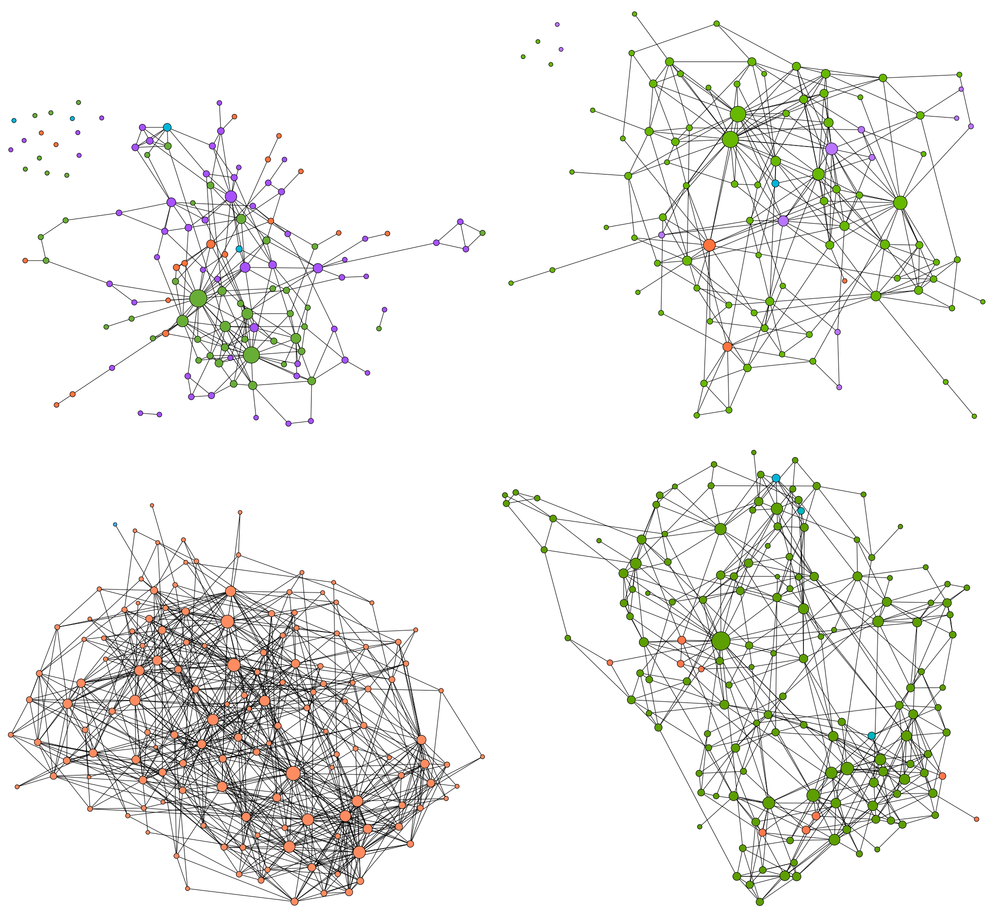

Figure 7: Village networks, 1 to 4 in order top left to bottom right. Nodes are households, sized proportional to degree in the network. Color indicates religion: green is Catholic, orange is Muslim, purple is Protestant, and blue is other.

5.2 Spillovers through Social Networks

Our approach assesses the case for social processing by triangulating from a variety of analyses. Taken together, they paint a picture strongly consistent with social processing taking place in each of the villages between the intervention and the endline survey. They also hint that the process may have been importantly different for treated and control respondents.

	V1	V2	V3	V4
Network Difference, Baseline	5.48	4.58	5.92	4.75
Network Difference, Endline	4.44	3.68	5.05	4.64
Network Dif for the Treated, Baseline	5.12	4.54	5.55	4.93
Network Dif for the Treated, Endline	5.00	4.07	5.06	4.68
Network Dif for the Control, Baseline	5.79	4.62	6.50	4.46
Network Dif for the Control, Endline	3.95	3.27	5.04	4.56

Table 4: Average absolute difference in network neighborhoods in the baseline compared to the endline. Calculated for the village networks overall, and separated out by treated and control nodes' neighborhoods.

We first show that respondents finished the study with views on refugees that were substantially more similar to their network neighbors' views than when they began. Table 4 computes a measure of network difference for every respondent in the network. This measure calculates the sum of the absolute differences between the respondent's prorefugee score and the score of each of her network neighbors and divides by the number of her network neighbors. A respondent with a baseline prorefugee score of 20, and two network neighbors who have baseline scores of 18 and 30, would have a network difference score of $(2 + 10)/2 = 6$ for the baseline. These are averaged over all respondents in the network to produce the network difference baseline score, and calculated in the same way using endline scores to produce the network difference endline score.

The first two rows show that network difference shrank from the baseline to the endline in all four villages. That is, people held endline refugee attitudes that were more similar to their network neighbors than their baseline attitudes were. The difference is largest in Village 1, where people became a whole point more similar to each of their neighbors. We might worry that network differences decreased mechanically due to the average increase in scores that are capped. If everyone reacted strictly individually to the treatment, and those reactions resulted in an average long-term increase that compressed more people's scores at the maximum value of 30, the network differences would shrink mechanically rather than due to any social processing. In the appendix, we show that the issue of hitting the cap can only explain a tiny portion of the decrease in network difference for these data.[15]

Table 4 also shows an intriguing difference between treated and control respondents. The bottom four rows decompose the change in network difference by treatment condition. In three of the four villages, the control respondents became much more similar to their network neighbors than the treated respondents did. Again zooming in on Village 1, while the treated respondents' endline scores become .12 points closer to their neighbors' scores, control respondents' scores become 1.84 points closer. This is the first of many pieces of evidence that suggest that social processing may have been especially prevalent among the control respondents.

We next investigate the role that the network may have played in ultimately determining the value of respondents' endline scores. We make use of individual-level measures of a respondent's social network position that might be relevant. The key reference set for any respondent is again their "network neighbors," the set of households to which they are directly linked through the relationships described above (sharing a meal, visiting, borrowing money, and chatting about rumors). Again, this is "neighbors" in the network sense– people to whom one is connected socially– and not in the geographic sense. We count how many of these network neighbors they have (`# Neighbs`), indicate whether any were treated (`Treated Neighbs`), compute the average baseline score of these neighbors (`Neighbs Bl Atts`), and account for the respondent's own baseline attitudes (`Baseline Atts`).[16]

[15] We also show that if the real changes to the baseline scores that we observe were shuffled at random in our observed networks, the network would not generate a decrease in the network differences of these sizes by chance; quite the contrary, network differences would go up, even accounting for the floor and ceiling of our prorefugee index.

[16] In the appendix, we show that the same results hold if we use a count of the number of treated neighbors

In Table 5, the relationships between these network features and a respondent's endline attitudes towards refugees are shown as coefficients in an OLS regression. Specification (1) regresses a respondent's endline score on the respondent's baseline attitudes, these three network variables, and their interaction with treatment to account for the possibility that spillovers work differently for treated and control (Vazquez-Bare, 2022).[17] This regression drops 14 observations from the 488 who remained in the endline because the network measures can only be calculated for respondents who have at least one network neighbor.

First, and unsurprisingly, a respondent's baseline attitudes are positively related to their endline attitudes, and the relationship is estimated with high precision. Individuals who started warmer towards refugees are likely to have warmer attitudes at endline. Still focusing on the first column of Table 5, treatment does not play a precisely estimated direct role (true for the marginal and interaction terms). What *is* consistently and precisely related to higher endline scores is having network neighbors with warmer baseline attitudes. In this model, the baseline scores of network neighbors are almost as related to a respondent's endline score as that respondent's own baseline score is (at least for the controls; an imprecisely estimated interaction term suggests the relationship might be attenuated for the treated). The relationship between network neighbors' views and endline views is evidence for social processing, and persists through a variety of specifications and added demographic and network controls.[18]

The next five columns in Table 5 add to our consideration a respondent's position in the village network relative to other potentially impactful reference households, chosen based on extreme baseline attitudes or extreme responses to treatment. The key variables

instead of an indicator for the existence of a treated neighbor.

[17]Our data and qualitative follow-up are consistent with a process that is more complicated than attitudes varying in response to different extents of exposure to treatment in the network (see Aronow and Samii, 2017). As we show below, the treatment received by some appears to be experienced differently than the treatment received by others, which has implications for how others exposed to their treatment are affected, and this process may be different for the treated and the control. Consequently, we start with the logic of Vazquez-Bare (2022) in that we account for the fact that the direct spillovers may affect the treated too, allow that to be different from the way they affect the control, and then directly examine the sources of spillover.

[18]The supporting information shows this, and also demonstrates that the same conclusions about the role of network neighbors' attitudes also hold in much simpler specifications than this flexible spillover model.

	DV: Endline Pro-Refugee Score					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	4.646	5.219	5.031	4.837	5.645*	5.706*
	(3.223)	(3.275)	(3.299)	(3.272)	(3.293)	(3.232)
Treated Neighbs	0.010	−0.394	−0.102	−0.317	0.354	−0.380
	(1.065)	(1.077)	(1.079)	(1.069)	(1.078)	(1.071)
# Neighbs	0.014	−0.014	0.001	−0.024	0.034	−0.022
	(0.040)	(0.041)	(0.042)	(0.042)	(0.042)	(0.043)
Baseline Atts	0.371***	0.349***	0.379***	0.383***	0.366***	0.341***
	(0.038)	(0.039)	(0.039)	(0.038)	(0.038)	(0.040)
Neigbs Bl Atts	0.315***	0.278**	0.335***	0.331***	0.312***	0.251**
	(0.108)	(0.110)	(0.111)	(0.109)	(0.109)	(0.111)
Dist to Warmest		−0.623***				−0.901***
		(0.238)				(0.260)
Dist to Coldest			−0.276			0.237
			(0.260)			(0.280)
Dist to Persuaded				−0.689***		−0.817***
				(0.237)		(0.244)
Dist to Backlashed					0.482**	0.785***
					(0.238)	(0.243)
Trt * Treated Neighbs	−1.157	−1.149	−1.377	−1.404	−1.818	−1.525
	(1.508)	(1.547)	(1.554)	(1.539)	(1.554)	(1.527)
Trt* # Neighbs	0.046	0.033	0.043	0.046	0.051	0.031
	(0.060)	(0.060)	(0.060)	(0.060)	(0.060)	(0.059)
Trt * Neighbs Bl Atts	−0.175	−0.200	−0.180	−0.179	−0.189	−0.209
	(0.146)	(0.147)	(0.148)	(0.147)	(0.147)	(0.145)
Constant	8.271***	11.343***	8.455***	9.859***	6.936***	12.353***
	(2.417)	(2.704)	(2.452)	(2.488)	(2.518)	(2.749)
Observations	474	470	470	470	470	470
R^2	0.206	0.216	0.206	0.219	0.211	0.248
Note:				*p<0.1; **p<0.05; ***p<0.01		

Table 5: Spillover analyses

calculate a household's network distance from the reference households. Network distance from one household to another counts the number of links in the shortest path that connects them in the network. If it takes a minimum of four hops along links to move from one of the households to the other, they are separated by network distance four. The supporting information provides the precise method by which these distance variables were constructed, as well as the selection of reference households in the "top set" of extreme scores and extreme reactions.[19]

These analyses use four new network variables. `Dist to Warmest` is the length of the shortest path between a respondent and the nearest household in the top set of warm baseline scores. A household that is directly linked to one of the warmest households has a `Dist to Warmest` value of 1. A household that is not directly linked to one of them, but is linked to a household that is linked to one of them, has `Dist to Warmest` value of 2, and so on. We do the same for the network distance to the respondents with the coldest baseline pro-refugee scores (`Dist to Coldest`), to the treated respondents whose attitudes warmed the most in response to treatment in the short-term (`Dist to Persuaded`), and to the treated respondents whose attitudes cooled the most in response to treatment in the short-term (`Dist to Backlashed`).

These analyses show that connections in the network to people who start very warm, to people who are most persuaded to become warm, and people who react most negatively to treatment are all related to endline attitudes in expectation. The farther a respondent is in the network from someone who started very warm, the colder their endline score is likely to be (and vice versa– the closer they are, the warmer their expected score). Likewise, the farther a respondent is from someone who was particularly persuaded by the treatment, the colder their endline score is expected to be. And, the farther a respondent is from someone who reacted negatively to the treatment, the *warmer* their attitudes end up. Being close to

[19] Villages 1 and 2 have multiple components in their networks. The extreme households are all in the giant component; that means the households outside the giant component are not connected to these households by any paths of finite length. Most of the households outside the giant component have degree 0, and so were not included in the network analyses anyway. Four households appear in components of size 2 in village 1, and so are also dropped from these analyses, reducing the sample size by four more.

people who start or become warm improves attitudes, as does being *far* from people who become colder. The final column (6) confirms that these relationships hold when combined in the same regression. These results are consistent with an interpretation that spillovers can be positive or negative, depending on the source.

The relationship between a respondent's network neighborhood's attitudes and the respondent's own endline attitudes continues to hold even when accounting for their social proximity to individuals with extreme starting attitudes and individuals with extreme responses to treatment. Notably, this is true *even controlling for the respondent's own baseline attitudes*.[20] Even when accounting for where a respondent's attitudes started and whether that respondent was treated, having network neighbors who are warmer towards refugees is related to a respondent becoming warmer towards refugees. The full regression in column (6) includes all four network distance variables at once. When all of these possible pathways of spillovers are accounted for, treatment may have a direct effect on baseline attitudes as well. Of course, the imprecisely estimated interactions complicates this interpretation, and suggest that the role of treatment warrants a closer look.

Table 6 does so, digging deeper into the different roles that the network may be playing for the treated and the control respondents. In these analyses, endline scores are regressed on all of the network features present in Table 6, now with the respondents separated into treated (left column) and control (right column) groups. We also add indicator variables for households that are serving as the top set to which the network distances are calculated.[21] Variables `Warmest` and `Coldest` are indicators for the respondents who have the warmest and coldest baseline scores (and to whom the distances in `Dist to Warmest` and `Dist to Coldest` are calculated). Likewise, variables `Most Persuaded` and `Most Backlash` are

[20]This is important because social networks tend to exhibit homophily on many dimensions; people have social ties with others who are like them. To the extent that a person's attitudes towards refugees are also something they tend to hold in common with network neighbors, a person's baseline score should account for this.

[21]For small sample sizes, we need to be careful that the households to which distances are calculated are not driving the distance results since their distances to themselves are 0. By adding an indicator variable for being the reference households, we can better distinguish whether it matters to be close to the most negative household or to *be* the most negative household.

	DV: Endline Pro-Refugee Score	
	Treated	Control
Treated Neighbs	−2.257*	−0.401
	(1.147)	(1.028)
# Neighbs	−0.020	−0.004
	(0.057)	(0.044)
Baseline Atts	0.428***	0.300***
	(0.072)	(0.061)
Neighbs Bl Atts	0.118	0.189*
	(0.119)	(0.111)
Warmest	0.277	−0.810
	(1.255)	(1.748)
Coldest	−3.679*	1.350
	(2.111)	(1.826)
Most Persuaded	2.565*	
	(1.519)	
Most Backlash	−4.727***	
	(1.343)	
Dist to Warmest	−0.214	−1.458***
	(0.446)	(0.410)
Dist to Coldest	−0.340	0.396
	(0.456)	(0.455)
Dist to Persuaded	−0.701*	−0.258
	(0.380)	(0.413)
Dist to Backlashed	−0.266	0.827**
	(0.386)	(0.411)
Constant	17.115***	13.816***
	(3.362)	(2.980)
Observations	261	209
R^2	0.308	0.261
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 6: Direct comparison of treated and control.

indicators for respondents who responded most warmly and most coldly to the treatment at the end of the baseline.

For both the treated and control, respondents' own baseline attitudes remain positively and significantly related to their endline attitudes. However strong social processing may turn out to be, individual processing remains a key part of the story for both treated and control. Only treated respondents have variation on `Most Persuaded` and `Most Backlash` by design since these are a function of responses to treatment. Those who were most persuaded by the treatment at the end of baseline remain about two and a half points warmer towards refugees at endline, and those who experienced the most backlash in response to treatment remain over four and a half points colder at endline. Extreme reactions appear to have persistence. Both the treated and control are eligible to have the most extreme start-of-baseline views. Interestingly, starting extreme is unrelated to endline views for the control. For the treated, on the other hand, those who started coldest still hold attitudes at the endline that are over 3 and a half points colder.

The role that the network plays in shaping endline attitudes appears to be different for treated and control respondents. The overall story for the treated is one about how others in the network received and reacted to treatment. For the control, the story is more about proximity to others with positive attitudes and distance from people who responded coldly to the treatment. Starting with the treated, the presence of treated network neighbors is negatively related to their endline attitudes. Of course this relationship is conditional on other network features that may relate to the treated neighbors, so we hesitate to conclude that more treated neighbors always moves a person negative, but exploring this possibility further is important for future research. A helpful starting point is the relationship between `Dist to Persuaded` and endline attitudes for the treated. Being far from those who were most warmed by the treatment is associated with lower endline scors while being close to them is associated with warmer scores. Combined, this suggests that for the treated, having social contact with not just those who recieved treatment but specifically those who received treatment and reacted favorably is important for warm endline scores (after all,

Figure 5 made clear that the average positive response to treatment is aggregating a wide variety of individual experiences).

For the control respondents, having treated neighbors does not seem to matter as much as the attitudes of those around them in the network, treated or not. If a control respondent's network neighbors' average pro-refugee score (`Neighbs Bl Atts`) were one point warmer, the respondent would be expected to have an endline score almost a fifth of a point warmer. Moreover, being closer to people who started especially warmly towards refugees (`Dist to Warmest`) and far from people who reacted most negatively to the treatment (`Dist to Backlashed`) are the most statistically significant predictors of warm endline pro-refugee scores for the control. Being one step closer to someone with one of the warmest baseline scores (a friend as opposed to a friend of a friend) is associated with an endline score almost one and a half points warmer, and being one step farther from someone who responded most negatively to the treatment is associated with over three quarters of a point warmer endline score.

Overall, these results are consistent with an interpretation that treatment kicked off reflection among the treated, which led to conversations with other villagers regardless of whether the conversation partners were treated themselves. This social processing helped shape ultimate attitudes based on the attitudes and reactions to treatment of social ties in the network.

A virtue of the current approach is an ability to peer into real social networks. The drawback is that the process that generated these networks might be correlated with factors that are relevant to the response to treatment. In other words, treatment was randomly assigned, but networks were not. One concern is that, although networks were measured pre-treatment, they might be correlated with unobserved factors that are themselves the true reason that attitudes landed where they did in the endline. If that were the case, then attitudes could appear to be related to network neighbors' endline attitudes without any active social processing.

To explore this possibility, we conduct a placebo test in which the new outcome is the measure of attitudes taken at the end of the baseline for the treated. This measure was taken after treatment was administered, but before the baseline survey ended. If the relationship between network characteristics and attitudes observed in Table 6 was truly due to active social processing (such as having discussions with network neighbors and determining their views), then we should not see the same relationships when the end of baseline attitudes are used as the dependent variable. This measure was taken before the respondent had a chance to leave and talk to anyone other than the enumerator. Any relationship with network features that appears in these specifications would be indicative of network characteristics potentially proxying for something other than active social processing.

Table 7 shows the results of this placebo test. Reassuringly, none of the network features' relationship to the end of baseline pro-refugee score are sizeable, nor are any estimated with precision. In most cases, the standard errors are much larger than the estimates. Also reassuringly, respondents' own baseline attitudes *do* still strongly predict their post-treatment attitudes; the warmer respondents started towards refugees, the warmer they were towards refugees after treatment at the end of the baseline survey. The same is true for our indicators for responding most warmly and most coldly to treatment at the end of baseline. Since this test uses the end of baseline measure of attitudes, the one used to construct these indicators, they should explain this measure of attitudes with high magnitude and precision, as they do. The important variables for this test are the number of treated neighbors, neighbors' attitudes, distances to those with extreme baseline views, and distances to those with extreme reactions to treatment. These do not explain variation in the end of baseline scores well. This means that these network characteristics only matter once a person has had a chance to turn to their networks and make use of them.[22] In short, the placebo test shows strong evidence of individual processing and no evidence of social processing, exactly what we would expect for the time period in which socializing with

[22]Social proximity to people who hold extreme views on refugees may be related to a respondent's own views due to past social processing, but their own baseline attitudes should account for this past network influence, as is borne out by this placebo test.

	Placebo DV: Baseline 2 Pro-Refugee Score
Treated Neighbs	0.432
	(0.651)
# Neighbs	0.020
	(0.033)
Baseline Atts	0.603***
	(0.041)
Neighbs Bl Atts	−0.011
	(0.068)
Warmest	−0.598
	(0.719)
Coldest	−0.385
	(1.229)
Most Persuaded	5.487***
	(0.856)
Most Backlash	−6.542***
	(0.778)
Dist to Warmest	−0.311
	(0.252)
Dist to Coldest	0.120
	(0.259)
Dist to Persuaded	−0.171
	(0.214)
Dist to Backlashed	−0.233
	(0.219)
Constant	12.337***
	(1.918)
Observations	278
R^2	0.647
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 7: Placebo test, re-running the analysis of table 6 using the second baseline attitude score for the treated as the outcome. Since respondents had no chance to engage in social processing before that attitude measurement, we should not see a relationship between network features and this outcome, as is indeed the case.

network neighbors was impossible.

6 Qualitative Validation of Social Processing

The logic of spillovers and social processing requires that people have a chance to learn what their network neighbors are thinking about refugees. Although it may be reasonable to assume that people might talk about these things, we directly investigate whether this mechanism could plausibly have been at play in our study.

A first direct measure is a question in the endline survey that asked respondents if they recalled having had at least one conversation with other villagers about refugees since our team first spoke with them. 53% of respondents said they had a specific memory of doing so, with substantial variation across the villages ranging from 37% in village 3 to 71% in village 2.

Additionally, and importantly for our spillover story, although a larger proportion of the treated respondents recalled and reported having a conversation about refugees, many control respondents did too. Table 8 shows the breakdown of respondents who reported having had at least one conversation with another villager about refugees since our team visited in the baseline, separated out by village and treatment condition. The first row indicates the proportion of each subset of respondents who said yes, they recalled having had a conversation. The next three rows show the proportion of these respondents who said refugees came up more often than usual, and whether they classified the information they heard in these conversations as mostly positive and mostly supportive of the idea of refugees coming to Uganda.

Villagers were talking about refugees, in many cases more than was typical before the study, and were hearing a mix of views on refugees in these conversations. This process involved not just treated respondents, but control as well.

A second piece of evidence also comes from a followup question in the endline asked

	V1		*V2*		*V3*		*V4*		*All*	
	T	C	T	C	T	C	T	C	T	C
Had Ref Convo	0.80	0.58	0.72	0.69	0.44	0.28	0.46	0.46	0.56	0.49
More Often	0.40	0.36	0.61	0.31	0.54	0.69	0.61	0.42	0.53	0.41
Mostly Positive	0.37	0.56	0.23	0.28	0.57	0.88	0.71	0.54	0.48	0.52
Mostly Supportive	0.51	0.56	0.32	0.21	0.57	0.81	0.66	0.65	0.53	0.52

Table 8: Respondents reporting in the endline that they have had a conversation with other villagers about refugees since our team first spoke with them, and the characteristics of those conversations, separated by treatment condition.

of respondents who recalled having had at least one conversation about refugees. We asked these respondents to name the villagers with whom they had these conversations. Effectively, this provides a spoke-about-refugees network. We can repeat the same network difference exercise as above, this time using as our network this record of who spoke to whom. We measure these links in the endline. These links are about interactions that occurred between the baseline and endline. For people listed who are in the village, we have a record of their (or someone in their household's) baseline scores. Putting these pieces together, we can observe whether people who conversed about refugees in the interim moved closer to one another in refugee attitudes between the baseline and the endline.

	V1	V2	V3	V4
Refugee Convo Difference, Baseline	5.29	4.70	5.99	4.51
Refugee Convo Difference, Endline	4.21	2.95	4.72	4.22

Table 9: Average absolute difference in network neighborhoods where the network is who reported in the endline having had a conversation about refugees since the baseline with whom. Compares the network difference in conversation partners' baseline scores with the endline scores. Conversation partners became much more similar after their conversations.

Table 9 shows the results. In all four villages, the people who conversed about refugees became more similar to one another in their attitudes. It is also informative to use the social network differences for the villages overall as a benchmark. In village 1 and 4, people conversed with people who were somewhat more similar to themselves in baseline views than their social network neighbors overall; in villages 2 and 3 people conversed with people whose baseline views were somewhat more different from their own than their social network neighbors overall. However, in all four villages, the conversation partners became much more similar to one another, even more so than their overall network neighbors did.

Finally, we collected a qualitative follow-up to our study about a year after it concluded. This follow-up entailed focus groups and one-on-one interviews with the local official (LC1) and a few villagers in each of the four villages. It was led by a researcher who was not a member of the original study's research team. Participants were asked what they remembered about the study and what their experiences with it were like. Many remembered the key details– a good sign since so much time had elapsed– and also reported experiences that we would label as social processing. Some mentioned seeking out others to see what they thought was going on. Some mentioned villagers seeking them out to do the same. Some mentioned attempts that resemble campaigning, explicitly aiming to change the views of others, especially on the issue of refugees coming to Uganda. These interactions led to conversations about refugees in which a variety of viewpoints were expressed. The qualitative follow-up points to a rich social process that contributed to the ultimate views of the villagers.

7 Conclusion and Questions for Future Work

Are perspective-taking exercises shown to reduce prejudice towards outgroups in the United States also effective in a highly resource constrained, developing country context? How do social networks reinforce or unravel the short-term improvements caused by interventions such as these? Can a light-touch intervention that accounts for social processing in networks better improve attitudes towards refugees at the community level? Answering these questions is pressing since, according to the UNHCR, the total number of refugees across the globe currently exceeds 80 million and is growing. Climate crises will likely increase the quantity of refugees globally in coming decades. A better understanding of how to improve their reception among existing populations is essential to the wellbeing of a substantial fraction of the world's population.

This paper moves a step towards answering these questions by showing that a perspective-taking intervention with proven effectiveness at reducing prejudice in industrialized countries can also reduce prejudice among Ugandan individuals towards South Sudanese refugees. It has also demonstrated that the intervention sparked a social process – an increased rate of conversations about refugees in the two weeks after our intervention – and coincided with improved average attitudes towards refugees *not only in treatment but also among control households* in the four villages where we carried out our study. That is, the intervention appears to have reduced prejudice on average for both the treated and control, likely through the indirect channel of discussions in the village that followed our intervention. These results highlight the importance of tracking and understanding spillovers in individual-level interventions such as these.

It also appears that individuals' experiences with the treatment matter for how the spillovers work. Our results are consistent with an interpretation that those who were most persuaded by the treatment during the baseline survey created positive spillovers whereas those who were most negatively influenced by the treatment created negative spillovers. Being close to the most persuaded but far from the largest backsliders led to the greatest warming in endline scores. Control individuals were caught up in the social process as well, but in a more diffuse way; the attitudes of those in their networks and their social proximity to others with extreme views whether or not they were treated seemed to matter the most for their ultimate endline scores.

This research raises many more questions than it answers, which opens a broad, pressing research agenda. Does social processing work differently across villages? Does the social processing look different in villages that view refugees as economic competitors compared to those who see them as economic partners? Although we detect something social happening across the board, these four villages are quite different in composition of occupation, level of education, religious affiliation, and, shown starkly in Figure 7, in social networks. Does the social network structure of villages– the density, the extent of isolated nodes, the length of paths between villagers– affect the character or the result of social processing?

Much is left to explore within villages and their networks as well. Average short-term reactions to the treatment were positive. This average includes most who responded positively and a few who responded negatively. Can we identify who the backsliders are likely to be in advance and find a way to harness the network to dampen the negative spillovers that appear to originate with them? Average long-term reactions were also positive, but this aggregate is also comprised of some positive and some negative reactions. Can we explain who ultimately becomes more positive and who moves negative in response to a treatment, and can we identify how the social network functions in this process?

In short, much remains for future work. We are hopeful that this study will help to lay the groundwork for future studies across a variety of contexts so that they may jointly fill out our understanding of social processing of changing attitudes towards outgroups.

References

Adida, Claire L, Adeline Lo and Melina R Platas. 2018. "Perspective taking can promote short-term inclusionary behavior toward Syrian refugees." *Proceedings of the National Academy of Sciences* 115(38):9521–9526.

Arias, Eric, Pablo Balán, Horacio Larreguy, John Marshall and Pablo Querubín. 2019. "Information provision, voter coordination, and electoral accountability: Evidence from Mexican social networks." *American Political Science Review* 113(2):475–498.

Aronow, Peter M. and Cyrus Samii. 2017. "Estimating Average Causal Effects Under General Interference." *Annals of Applied Statistics* 11(4):1912–1947.

Atwell, Paul and Noah L. Nathan. 2021. "Channels for influence or maps of behavior? A field experiment on social networks and cooperation." *American Journal of Political Science* .

Audette, Nicole, Jeremy Horowitz and Kristin Michelitch. 2020. "Personal Narratives Reduce Negative Attitudes toward Refugees and Immigrant Outgroups: Evidence from Kenya.".

URL: *https://tinyurl.com/2r48s4j5*

Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo and Matthew O Jackson. 2013. "The diffusion of microfinance." *Science* 341(6144):1236498.

Bansak, Kirk, Jeremy Ferwerda, Jens Hainmueller, Andrea Dillon, Dominik Hangartner, Duncan Lawrence and Jeremy Weinstein. 2018. "Improving refugee integration through data-driven algorithmic assignment." *Science* 359(6373):325–329.

Blair, Graeme, Rebecca Littman, Elizabeth R. Nugent, Rebecca Wolfe, Mohammed Bukar, Benjamin Crisman, Anthony Etim, Chad Hazlett and Jiyoung Kim. 2021. "Trusted authorities can change minds and shift norms during conflict." *Proceedings of the National Academy of Sciences* 118(42).

Broockman, David and Joshua Kalla. 2016. "Durably reducing transphobia: A field experiment on door-to-door canvassing." *Science* 352(6282):220–224.

Burns, Justine, Lucia Corno and Eliana La Ferrara. 2018. "Interaction, prejudice and performance: Evidence from South Africa." *Working Paper* . Available at:.
URL: *https://tinyurl.com/u2jf6z8p*

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.

Chiang, Mung, Steven H. Low, A. Robert Calderbank and John C. Doyle. 2007. "Layering as Optimization Decomposition: A Mathematical Theory of Network Architectures." *Proceedings of the IEEE* 95(1):255–312.

Choi, Donghyun Danny, Mathias Poertner and Nicholas Sambanis. 2019. "Parochialism, social norms, and discrimination against immigrants." *Proceedings of the National Academy of Sciences* 116(33):16274–16279.

Chwe, Michael Suk-Young. 1999. "Structure and strategy in collective action." *American journal of sociology* 105(1):128–156.

Chwe, Michael Suk-Young. 2000. "Communication and coordination in social networks." *The Review of Economic Studies* 67(1):1–16.

Dodds, Peter Sheridan and Duncan J Watts. 2004. "Universal behavior in a generalized model of contagion." *Physical review letters* 92(21):218701.

Eubank, Nicholas et al. 2019. "Social networks and the political salience of ethnicity." *Quarterly Journal of Political Science* 14(1):1–39.

Eubank, Nick, Guy Grossman, Melina Platas and Jonathan Rodden. 2021. "Voting: Social Networks and Political Participation." *Quarterly Journal of Political Science* .

Fearon, James D and David D Laitin. 1996. "Explaining interethnic cooperation." *American political science review* 90(4):715–735.

Ferrali, Romain, Guy Grossman, Melina R Platas and Jonathan Rodden. 2020. "It takes a village: peer effects and externalities in technology adoption." *American Journal of Political Science* 64(3):536–553.

Gould, Roger V. 1993. "Collective Action and Network Structure." *American Sociological Review* 58(2):182–196.

Greenhill, Kelly M and Ben Oppenheim. 2017. "Rumor has it: The adoption of unverified information in conflict zones." *International Studies Quarterly* 61(3):660–676.

Hainmueller, Jens and Daniel J Hopkins. 2015. "The hidden American immigration consensus: A conjoint analysis of attitudes toward immigrants." *American Journal of Political Science* 59(3):529–548.

Hopkins, Daniel J, John Sides and Jack Citrin. 2019*a*. "The muted consequences of correct information about immigration." *The Journal of Politics* 81(1):315–320.

Hopkins, Daniel J, John Sides and Jack Citrin. 2019*b*. "The muted consequences of correct information about immigration." *The Journal of Politics* 81(1):315–320.

Jackson, Matthew O, Brian W Rogers et al. 2007. "Relating network structure to diffusion properties through stochastic dominance." *The BE Journal of Theoretical Economics* 7(1):1–13.

Kalla, Joshua L and David E Broockman. 2020. "Reducing exclusionary attitudes through interpersonal conversation: Evidence from three field experiments." *American Political Science Review* 114(2):410–425.

Kalla, Joshua L and David E Broockman. forthcoming. "Which Narrative Strategies Durably Reduce Prejudice? Evidence from Field and Survey Experiments Supporting the Efficacy of Perspective-Getting." *American Journal of Political Science* .

Larson, Jennifer M. 2017. “Networks and interethnic cooperation.” *The 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. 2020. “Measuring Networks in the Field.” *Political Science Research and Methods* 8(1):123–135.

Larson, Jennifer M., Janet I. Lewis and Pedro Rodríguez. 2022. “From Chatter to Action: How Social Networks Inform and Motivate in Rural Uganda.” *British Journal of Political Science* 52(4):1769–1789.

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

Marwell, Gerald, Pamela E. Oliver and Ralph Prahl. 1988. “Social Networks and Collective Action: A Theory of the Critical Mass III.” *American Journal of Sociology* pp. 502–534.

Newman, Mark EJ. 2000. “Models of the small world.” *Journal of Statistical Physics* 101(3):819–841.

Paluck, E.L. 2010. “Is It Better Not to Talk? Group Polarization, Extended Contact, and Perspective Taking in Eastern Democratic Republic of Congo.” *Personality and Social Psychology Bulletin* 36(9):1170–1185.

Rosenzweig, Leah and Yang Yang Zhou. 2021. “Team and Nation: Sports, Nationalism, and Attitudes toward Refugees.” *Comparative Political Studies* 54(12):2123–2154.

Schon, Justin. 2021. “How Narratives and Evidence Influence Rumor Belief in Conflict Zones: Evidence from Syria.” *Perspectives on Politics* 19(2):539–552.

Search For Common Ground. 2021. “Uganda Conflict Scan.”. Available at:.
URL: *https://www.sfcg.org/wp-content/uploads/2021/10/SFCG-Uganda-Third-Conflict-Snapshot-September-2021.pdf*

Shaver, Andrew and Yang-Yang Zhou. 2021. "Reexamining the Effect of Refugees on Civil Conflict: A Global Subnational Analysis." *American Political Science Review* 115(4):1175–1196.

Shesterinina, Anastasia. 2021. *Mobilizing in uncertainty: collective identities and war in Abkhazia*. Cornell University Press.

Siegel, David A. 2009. "Social networks and collective action." *American journal of political science* 53(1):122–138.

Simonovits, Gábor, Gabor Kezdi and Peter Kardos. 2018. "Seeing the world through the other's eye: An online intervention reducing ethnic prejudice." *American Political Science Review* 112(1):186–193.

UNHCR. 2018. "Uganda Refugee Response Plan: Livelihoods Sector Technical Working Group Response Plan.". Available at:.
URL: *https://reliefweb.int/report/uganda/uganda-refugee-response-plan-livelihoods-sector-technical-working-group-response-plan.*

UNHCR. 2022. "High Commissioner for Refugees (UNHCR) Refugee Data Finder.". Accessed January 10, 2022. Available at:.
URL: *https://www.unhcr.org/refugee-statistics/*

Valente, Thomas W. 1996. "Social network thresholds in the diffusion of innovations." *Social networks* 18(1):69–89.

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

Vazquez-Bare, Gonzalo. 2022. "Identification and estimation of spillover effects in randomized experiments." *Journal of Econometrics* .

Williamson, Scott, Claire L Adida, Adeline Lo, Melina R Platas, Lauren Prather and

Seth H Werfel. 2021. "Family matters: How immigrant histories can promote inclusion." *American Political Science Review* 115(2):686–693.

World Vision. 2018. "Uganda West Nile Refugee Response Situation Report 19.". Available at:.
URL: *https://reliefweb.int/sites/reliefweb.int/files/resources/Uganda%20SITREP19-final-updated.pdf*

Zhou, Yang-Yang. 2019. "How Refugee Resentment Shapes National Identity and Citizen Participation in Africa.". Working paper available at.
URL: *https://tinyurl.com/2p99wzed*

Zhou, Yang-Yang and Guy Grossman. 2022. "When Refugee Presence Increases Incumbent Support through Development: Evidence from Uganda.". Working paper available at:.
URL: *https://osf.io/94tpy/*