

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

Jennifer M. Larson[†] Janet I. Lewis[‡]

February 15, 2024

Word Count: 11,694

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 Rebecca Wai, Scovia Aweko, Alison Craig, John Marshall, the Impact Evaluation Network at Makerere University, participants of the Overcoming Prejudice Against Immigrant Minorities conference hosted by the Identity & Conflict Lab at the University of Pennsylvania, and participants of the Empirical Models of Political Economy conference at Caltech. We also thank the members of Vanderbilt's Research on Conflict and Collective Action Lab for their valuable research assistance. We gratefully acknowledge funding for this project from George Washington University and UK International Development, awarded through Innovation for Poverty Action's Peace & Recovery Program.

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

[‡]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 strong majority 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.

Guided by new theory, 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

¹In 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).

better explained 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, as they were also involved in post-treatment conversations. Control attitudes warmed on average, and also moved towards the attitudes of their social network ties.

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 the spread of pro-refugee attitudes than others? We discuss promising ways to continue to advance this agenda in the conclu-

sion.

2 Theory of Prejudice Reduction through Social Networks

Our theory of prejudice reduction connects 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.

The intervention seeks to warm attitudes towards refugees through non-judgmental conversations that encourage taking the perspective of an out-group member (described in greater detail below). A robust literature elucidates the psychological mechanisms that underly the efficacy of this treatment in other settings: such conversations can reduce individuals' natural resistance to persuasion by avoiding self-image concerns, boost empathy, and actively engage the respondent in considering sources of his or her views (e.g. Galinsky and Moskowitz, 2000; Kalla and Broockman, 2020, 2023). We consider mechanisms such as

²Note 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.

these to be instances of “individual processing” and hypothesize that they will also operate to immediately reduce prejudice towards refugees in our setting:

Hypothesis 1 *Attitudes of individuals who engage in the perspective-taking intervention will warm in the short-term.*

An important though less-studied question is: what happens after the conversation warms attitudes? After all, the effect of interventions, perspective-taking included, tends to change over time. Are ultimate attitudes strictly the result of more individual processing, or do other people play a role?

Although it is possible that the treated continue to process exclusively on their own, individuals are embedded in a rich web of social relationships to which they could turn to further process their experiences with the intervention. Here we draw on insights from a broad range of work. When people are presented with new information, a natural reaction is to turn to social contacts to discuss or vet it (Atwell and Nathan, 2021; Larson, Lewis and Rodríguez, 2022). In general, people like to be accepted by and comply with the norms of their core set of social contacts (Falk and Scholz, 2018; Sinclair, 2012), so if their attitude changes, it would be natural to suss out social reactions. Information spreading through a network of social contacts can be persuasive, in some contexts “infecting” a person with motivation to do something that her peers are planning to do (Marwell, Oliver and Prahel, 1988; Gould, 1993; Centola and Macy, 2007; Centola, 2013).

Moreover, people regularly learn from their neighbors in social networks. If they hear what their social contacts know, they can update their own understanding in light of it (DeGroot, 1974; Golub and Jackson, 2010; Mobius and Rosenblat, 2014; Tian and Wang, 2023). People may also change their beliefs based, in part, on how valuable they are to hold socially (Sharot et al., 2023; Bénabou and Tirole, 2016). Information from social contacts can matter in even the highest stakes contexts. For instance, as people decide what to do amidst the uncertainty of conflict, trusted social ties appear to play an essential role in vetting narratives 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).

Putting these insights together, we hold that a person’s long-term attitudes may be shaped by not only individual processing, but social processing as well. We argue that individuals in our study had access to an additional, social source of information after their personal experience with the treatment– information about what social contacts think about their updated attitude, how others who received treatment reacted, what others who have heard about new attitudes in the village are think about that, and so on.

In Appendix B, we present a simple model that allows us to be precise about how individual and social processing could act together on long-term attitudes. It represents a treated person i ’s attitude in the short term as a weighted average of the treatment, y^* , and i ’s prior, baseline attitude $y_{i,bl}$, weighted by how sure i is about her baseline attitude ($0 \leq s_i \leq 1$).³

$$y_{i,st} = (1 - s_i)y^* + s_i y_{i,bl}$$

It considers treated i ’s long-term attitude, $y_{i,lt}$ to be a function of her short-term attitude $y_{i,st}$ that resulted from the treatment as well as two possible additional forces: environmental factors that contributed to her baseline attitude $y_{i,bl}$ in the first place, and possibly also social processing that recommends an attitude $y_{i,nw}$ based on her network. We represent this as:

$$y_{i,lt} = (1 - w_i^e - w_i^{nw})y_{i,st} + w_i^e y_{i,bl} + w_i^{nw} y_{i,nw}$$

where w_i^e and w_i^{nw} are the weights placed on the environment and the network, respectively

³This setup is meant to capture as simply as possible the intuition of a standard Bayesian learning model (Bartels, 2002; Clinton and Grissom, 2015; Anoll and Engelhardt, 2023). A person’s short-term attitude will respond more strongly to treatment when the treatment is a stronger signal (y^* is especially different from a person’s prior) and when a person is less sure about her prior (s_i is smaller). Our respondents in fact responded in a way consistent with such models; see Appendix B.2.

(such that $0 \leq w_i^{nw} \leq 1$, $0 \leq w_i^e \leq 1$, and $w_i^{nw} + w_i^e \leq 1$). This setup is meant to capture as simply as possible the key components of individual processing along with the possibility of social processing. Our key hypothesis about the long-term is that social processing is present:

Hypothesis 2 *The long-term attitudes of the treated are formed at least in part due to social processing ($w^{nw} \neq 0$).*

This general formulation leaves open the question of how social processing actually works to “recommend an attitude.” In principle, there are many ways a person could be influenced by others in her social network. We envision the process as one in which people become activated to share their reactions with others in their network. The treated can react to their experience with treatment. Neighbors of someone treated can react to their neighbor holding a new attitude. Others can react to the presence of attitude changes occurring in their community. If these people share their reactions with others in the network, then long-term attitudes may move not just as a function of the treatment and priors, but also of these activated reactions.

Appendix B presents a formal version of this process, which leads to a few testable implications useful for detecting social processing. First, we show that if no social processing were present, we should expect an individual’s response to treatment to either hold durably or to attenuate back towards her prior, baseline view. Social processing admits other options. A person’s response can accelerate, becoming even warmer or colder over time, or can flip, for instance becoming colder than baseline in the long-term when the treatment warmed the attitude in the short-term.

Testable Implication 1 *If no social processing were present, long-term attitudes of the treated would feature either effect durability or attenuation. The presence of other individual effect trajectories, for instance acceleration or flipping, are evidence of social processing.*

We also show that if social processing results in a treated person’s network neighbors

telling her their reactions, they can pull her attitude towards theirs.

Testable Implication 2 *Social processing can pull long-term attitudes of the treated towards the attitudes of their network neighbors.*

Figure 1 illustrates the logic for a node i , whose response to treatment could be pulled warmer if his three neighbors are quite warm and colder if his neighbors are colder.

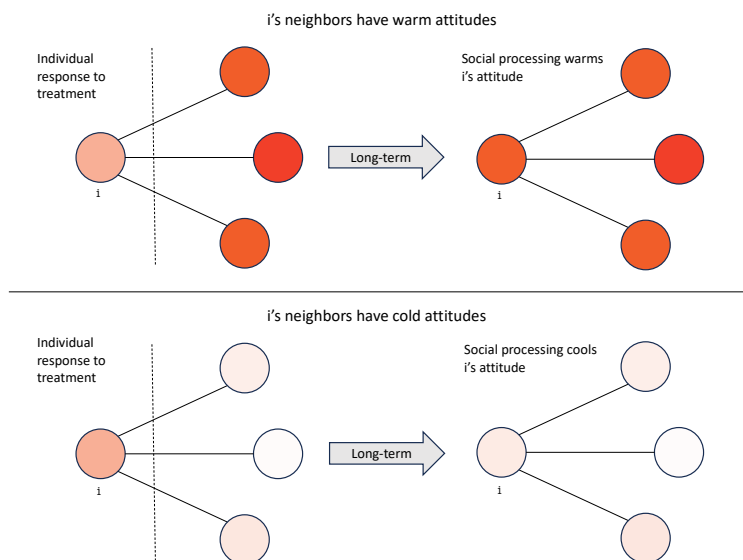


Figure 1: Social processing in network neighborhoods.

Furthermore, we note that the process of hearing reactions and updating based on them is not confined to the ears of the treated. If social processing is present, individuals in the control condition can learn reactions of their network neighbors too— either due to the neighbors' own treatment, or their reactions to others who were treated, or their reactions to conversations initiated by others who are reacting— and may also move towards the attitudes of their network neighbors in the long term.

Testable Implication 3 *Social processing can pull long-term attitudes of the control towards the attitudes of their network neighbors.*

Finally, we show that long-term attitudes can be pulled towards those of people farther

away in the network through social processing if people are highly motivated to share their reactions— perhaps because they were especially persuaded by the treatment, or especially concerned by it. Both the treated and control may respond, since again these reactions could be heard by either.

Testable Implication 4 *Social processing can pull long-term attitudes of the treated and the control towards the attitudes of activated individuals farther away in the network, though with diminishing impact by network distance.*

Figure 2 conveys the logic of Testable Implication 4: an activated respondent whose attitude was especially warmed by treatment, say, could spread warmth through the network, with reduced potency the farther it travels.




Figure 2: Social processing over network distances.

Each of these testable implications requires careful measurement of the social network and an operationalization of its features. We describe our measurement strategy below.

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).⁴ 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 (Nambuya, Okumu and Pagnucco, 2018; Ebere and Mwesigwa, 2021; 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.⁵

⁴Refugee 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.

⁵Author conversation via Skype with Pax Sakari, Director of Rural Initiative for Community Empowerment (RICE)-Uganda (January 2022).

The most severe recent case of intergroup violence was a 2020 attack on South Sudanese refugees that left over 10 dead and 15 homes destroyed, resulting in police and military deployment to the area in order to prevent escalation.⁶

3.2 Study Design

We carried out our study from February to August 2021. In each of the four study villages, a randomly selected set⁷ 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. Because we were interested in and anticipated spillovers, we completed baseline surveys for the control households first before beginning any treated ones in each village.⁸

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.⁹ 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

⁶Samuel Okiror, "Uganda calls in troops as violence flares between refugees and locals." *The Guardian* September 15, 2020.

⁷50% of households in villages 1 and 2, about 60% in 3 and 4.

⁸It appears that spillovers started quickly after treatment began in each village; see SI Figure 10.

⁹Kalla and Broockman (2023) distinguishes among and tests the component pieces of the Broockman and Kalla (2016) intervention. The authors 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.

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 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 Appendix A.¹⁰

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 village’s population receives. We expect that this design would lead to similar results in any context where local word-of-mouth networks are a key source of vetting unverified news (including “gossip” and “rumors”) and of shaping perceptions of social issues –especially attitudes towards outgroups. We expect this to be the case in many rural contexts in low-income countries, especially in Sub-Saharan Africa. We conjecture that cell phone penetration (which was low in our study villages) and frequency of contact with the outgroup (also fairly low in our villages) could attenuate the effectiveness of the intervention and likelihood of local spillover and hope these issues of external validity will be the subject of future research.

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.

¹⁰This study was not pre-registered. Our intent was to use it as a proof of concept, probing whether social processing appeared to be occurring and whether we could detect it in these networks. As we found yes to both, we have preregistered a follow-up study that includes the analyses below.

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 (VU #202053, GW #202995), from Uganda's National Council on Science and Technology (SS662ES), from a local Ugandan IRB (Mildmay Uganda Research Center, 0210-2020) and from the relevant district-level officials.¹¹ 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 in the West Nile region.¹² We selected four villages from the population of villages in West Nile using three criteria: size of village (aiming for an average-sized village of roughly 100-150 households); distance of at least 10 kilometers from the nearest peri-urban or urban area (to help ensure similarity of average wealth and education, as well as relevance of word-of-mouth networks); and distance of between 40-60 minutes via public transit to the nearest refugee settlement (to help ensure similarity of contact frequency with refugees).¹³ We also sought geographic diversity within West Nile (hence our villages come from three districts) and variation in village-level social heterogeneity (religion); allowing for this variation enables us to probe whether either factor strongly obstructs spillover.¹⁴

¹¹Certificates can be found with the replication materials (Larson and Lewis, 2024).

¹²Village 1 is in Arua district; village 2 is in Maracha district; villages 3 and 4 are in Yumbe district. Uganda has over 130 districts and a population of over 45 million.

¹³Consent of village leadership was a fourth criterion, but we did not face any refusals.

¹⁴We use religion to measure social homogeneity since in much of West Nile, while most people speak the same language (Lugbara) as their "mother tongue," the most socially-salient cleavage is religion, which also tends to overlap with kinship networks and which dialect (of Lugbara) is spoken in the home.

Our survey data show that the villages we selected 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; who report farmer or trader as their occupation; who report receiving no education or at least some primary, secondary, or college education; and 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 are a salient issue. Table 2 shows that many of our respondents were once refugees themselves.¹⁵ Most of the respondents have personally met a refugee, with the highest frequency in Village 1 where 76% of respondents have done so.¹⁶

The topic of refugees also comes up regularly for many of our respondents. In villages

¹⁵These 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.

¹⁶Although most refugees live in settlements separate from host communities, interactions with refugees are relatively common at shared water collection areas, markets, and sometimes in hospitals and schools.

	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.

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 that 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.

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 (2019b) 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 3 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.¹⁷ Baseline attitudes contain a fair bit of variation on all constituent questions.

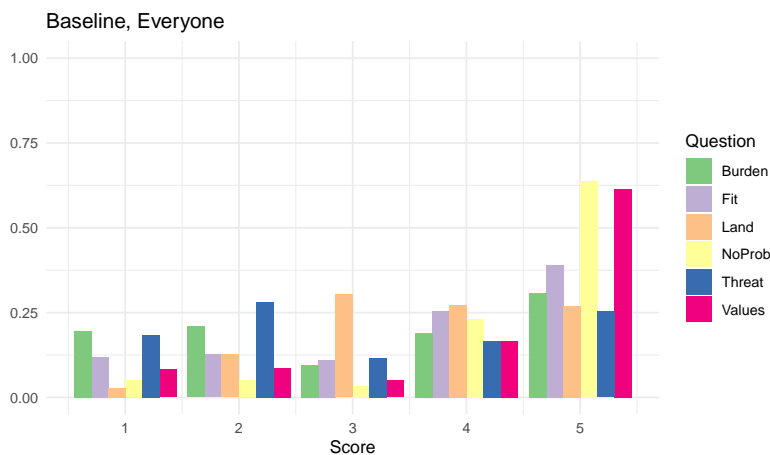


Figure 3: Baseline attitudes broken down by question for all villagers. Answers to each question range from least pro-refugee on the left to most pro-refugee on the right.

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

¹⁷Appendix C shows these baseline attitudes also broken apart by village.

4.2 Individual Short-Term Response to Treatment

We test Hypothesis 1 by comparing the treated respondents’ pro-refugee score before and after treatment. On average, respondents’ pro-refugee score increased 2.5 points in immediate response to treatment.¹⁸

	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.¹⁹ 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

¹⁸We also calculate the Average Treatment Effect (2.9), which requires accounting for some subtle SUTVA concerns; see Appendix C.

¹⁹We note that some respondents did respond by moving more negative in their attitudes towards refugees. Only 14% of respondents did so.

remained more positive towards refugees compared to their baseline attitudes, but the average increase was somewhat attenuated.²⁰ Figure 4 shows this pattern by village and 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 (immediately post-treatment) and endline points to indicate the statistical precision of a difference in means t-test when compared with the baseline.



Figure 4: 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 $-p < .1$; $*p < .05$; $**p < .01$; $***p < .001$.

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 (see Appendix C for more).

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

²⁰488 of the 521 households remained in our study through the endline. The 6% who attrited do not systematically differ from those who remained in the study in terms of demographic or network attributes (see Appendix D).

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 5 presents the same information as Figure 4, this time in terms of the *change* in score. The short-term change is the difference between the treated respondents' pro-refugee score at the end of the baseline survey (after treatment) and their score measured earlier in that survey (before treatment). 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. The center of the horizontal axis indicates no change in pro-refugee score. Respondents to the right saw warming in their score, and to the left saw cooling.

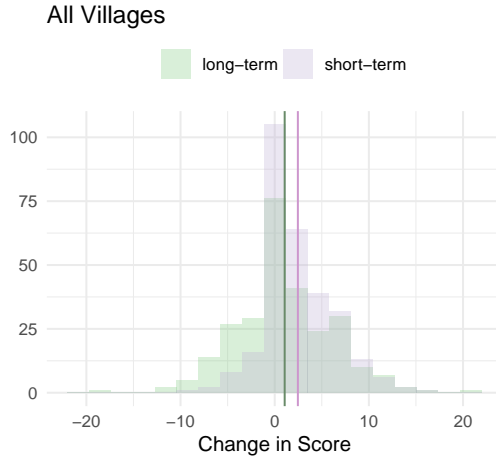


Figure 5: Change in attitude score of the treated in the short- and long-term, pooled across villages.

Figure 5 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 both the short- and long-term, by on average 2.5 and 1.1 points respectively, but the range in responses is wide, some moving more than ten points. Moreover, Testable Implication 1 guides us to examine individual effect trajectories. When we do, we see that although warming and attenuation are present on average, not all respondents follow this trajectory. In fact, all combinations of score changes are present in the data.

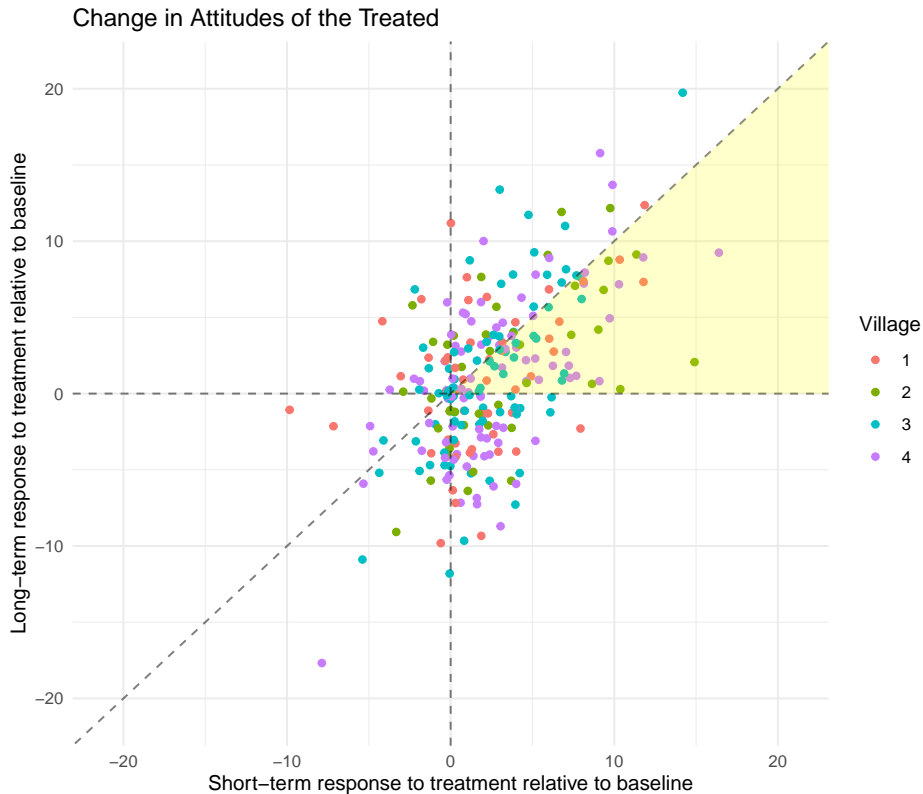


Figure 6: Change in attitude score of the treated; short-term change on the horizontal axis, long-term change on the vertical. If individual attitude changes were simply attenuating or wearing off, we should observe respondents primarily in the highlighted wedge, where short-term change is positive and long-term change is as well but with smaller magnitude.

Figure 6 displays the responses of the treated in a way that reveals these effect trajectories. It plots 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. If the effect were fully durable, respondents would fall on the 45-degree line. Some of our respondents are represented there too. But most lie elsewhere on the plot. Some became warm and then got warmer in the long-term (acceleration). Some became cooler immediately but then moved warmer in the long-term (flipping). Some got

warm but then cooled substantially (flipping). According to Testable Implication 1, these effect trajectories of the treated are consistent with the presence of social processing, and are the first set of evidence in support of Hypothesis 2.

4.4 Control Respondents and Spillover Effects

Were treated individuals the only ones who were ultimately affected by this treatment? In the short-term the answer is yes, by design. Control attitudes were measured before any treatment was administered in the village, and treatment was administered privately, without any other respondents present. Because the second baseline measure of attitudes was collected immediately after the treatment, there was no chance to talk to anyone other than the enumerator between the two measures.

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.²¹ Testable Implications 3 and 4 point out that long-term changes in the attitudes of the control are consistent with social processing.²²

Figure 7 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 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

²¹In other words, the stable unit treatment value assumption (SUTVA) might not hold in the long term, again by design.

²²This argument assumes that nothing substantial occurred external to our study to warm attitudes. In the language of the theory, it assumes that the day to day environment continued unchanged. We believe this to be the case during our study; see Appendix F for evidence.

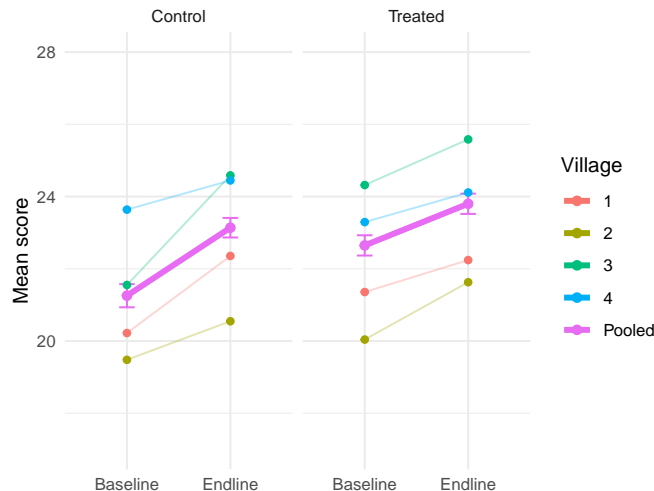


Figure 7: Change in mean attitude score of the treated compared to control in the long-term, pooled and separated by village. Bars show standard error of the pooled mean.

opens the possibility that the treated did not keep the treatment to themselves.

5 Social Processing

We consider the possibility that treatment kicked off the sharing of social information which contributed to the ultimate endline attitudes of both the treated and the control. In the two to three week interim between the baseline and endline surveys, respondents could reach out to people they trust to discuss their reaction, 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— to make use of Testable Implications 2 through 4—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 8 shows the resulting networks measured for each village.

5.2 Spillovers through Social Networks

Our approach assesses the case for social processing (Hypothesis 2) 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.

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




Figure 8: 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.

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.²³

²³This also provides a view of homophily, the extent to which people are linked to others with similar views on refugees. That people are on average quite different— about five points different— than their neighbors at the start hints that these networks were not formed primarily *because of* shared refugee views. We consider this point more fully below.

	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.

The first two rows show that network difference shrank from the baseline to the endline in all four villages. That is, people end the study with refugee attitudes that are 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 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 scores at the maximum value of 30, the network differences would shrink mechanically rather than due to any social processing. In Appendix E.1, we show that hitting the cap can only explain a tiny portion of the decrease in network difference for these data.²⁴ The bottom four rows decompose the change in network difference by treatment condition and show that in three of the four villages, the control respondents became even more similar to their network neighbors than the treated respondents did.

We next investigate the role that the network may have played in ultimately determining endline scores. Our theory points to individual-level measures of a respondent’s social network position that might be relevant. A key reference set for any respondent is 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

²⁴In Appendix E.1 we also show that if observed changes in scores were shuffled at random in our observed networks, we would not see a decrease in network differences of these sizes by chance.

connected socially— and not in the geographic sense. Testable Implication 2 suggests that social processing could lead a person’s long-term attitudes to move towards those of their neighbors’ baseline attitudes.

To test this logic, we count how many neighbors a respondent has (`# 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`).²⁵

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).²⁶ 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).

²⁵In Appendix E.3, we show that the same results hold if we use a count of the number of treated neighbors instead of an indicator for the existence of a treated neighbor.

²⁶Our data and qualitative follow-up suggest a process which 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.

The relationship between network neighbors’ views and endline views is more evidence in support of Hypothesis 2 (the presence of social processing), and persists through a variety of specifications and added demographic and network controls.²⁷

The next five columns in Table 5 incorporate Testable Implication 4 by adding 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 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. Appendix E.2 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.²⁸

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

²⁷See Appendix E.3, which demonstrates that the same conclusions about network neighbors’ attitudes also hold in simpler specifications than this flexible spillover model.

²⁸Four households are in components of size two in village 1. They are not connected to the reference households (all in the giant component) by any paths of finite length, so are dropped from these analyses which reduces the sample size by four more.

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

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 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.²⁹ Note that all specifications control for one’s own baseline attitudes, which helps to alleviate concerns about selection into the networks.³⁰ These results add support for Hypothesis 2. They are also consistent with an interpretation that spillovers can be positive or negative, depending on the source.

A virtue of this paper’s approach is its 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 we observe between network characteristics and attitudes was truly due to active social

²⁹Appendix E.3 shows the results also hold without `Neighbors B1 Atts` included, and with indicators for the reference categories included.

³⁰This 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, or that determine how far away they are from others with extreme views in the network, a person’s baseline score should account for this.

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.

Placebo DV: Baseline 2 Pro-Refugee Score	
Treated Neighbs	0.432 (0.651)
# Neighbs	0.020 (0.033)
Baseline Atts	0.603*** (0.041)
Neighb 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 6: Placebo test, re-running the analysis using the second baseline attitude score for the treated as the outcome. Since respondents had no chance to engage in social processing between the measurement of Baseline Atts and that attitude measurement, we should not see a relationship between network features and this outcome, as is indeed the case (see highlighted rows).

Table 6 shows the results of this placebo test.³¹ 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

³¹We use the full specification, including indicators for reference categories. **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). **Most Persuaded** and **Most Backlash** are indicators for respondents who responded most warmly and most coldly to the treatment at the end of the baseline. See Appendix E.3, Table 11, for the identical regression using endline pro-refugee scores.

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 neighbors' attitudes, distances to those with extreme baseline views, and distances to those with extreme reactions to treatment, highlighted in the table. 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.³² 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 network neighbors was impossible.

Taken together, these analyses show consistent support for the presence of social processing. The 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.

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.

³²Social 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.

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.

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 7 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 was usual before our study, 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. In sum, villagers (both treatment and control) were talking about refugees after our intervention, in many cases more than was typical before the study, and were hearing a mix of views on refugees in these conversations.

	<i>V1</i>		<i>V2</i>		<i>V3</i>		<i>V4</i>		<i>Pooled</i>		
	T	C	T	C	T	C	T	C	T	C	All
Had Ref Convo	0.80	0.58	0.72	0.69	0.44	0.28	0.46	0.46	0.56	0.49	0.53
More Often	0.40	0.36	0.61	0.31	0.54	0.69	0.61	0.42	0.53	0.41	0.48
Mostly Positive	0.37	0.56	0.23	0.28	0.57	0.88	0.71	0.54	0.48	0.52	0.50
Mostly Supportive	0.51	0.56	0.32	0.21	0.57	0.81	0.66	0.65	0.53	0.52	0.53

Table 7: 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.

A second piece of evidence also comes from a followup question in the endline asked 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 8: Average absolute difference in network neighborhoods where the network is who conversed about refugees with whom since the baseline. Compares the network difference in conversation partners’ baseline scores and endline scores. Conversation partners became much more similar after their conversations.

Table 8 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.

Furthermore, consistent with the impressions of our survey team during the study, none noted relevance of outside factors – unrelated to our study – that could have caused a warming of attitudes towards refugees among our treatment and control groups.³³

7 Conclusion and Questions for Future Work

This paper shows 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.

³³For further detail on how we attempt to rule out time trends or other events unrelated to our study that warmed attitudes in our study villages, see Appendix F.

This research raises many more questions than it answers, which opens a broad, pressing research agenda with potential importance for both theories of prejudice and belief-formation as well as practical implications for improving social cohesion. Researchers directly control but a small part of the bundle of new information and experiences that appears to ultimately shape attitudes following an intervention. Whether researchers can indirectly control the social reactions that follow— who is activated and what their reactions are— is one of many important, open questions. Ensuring that attitudes move in the intended direction and maximizing the effect of an intervention depend on better understanding how this works. The number of studies measuring social networks carefully enough to potentially detect this kind of social processing has grown in recent years (e.g. Ferrali et al., 2020; Arias et al., 2019; Eubank et al., 2019; Atwell and Nathan, 2021), making it all the more possible for this agenda to come to fruition.

Numerous questions remain about whether social processing works differently across communities. Although we detect something social happening across the board, the four study villages in this paper are quite different in composition of occupation, level of education, religious affiliation, and, shown starkly in Figure 8, 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? Of course, the context of host communities could influence the results. For example, does the social processing look different in villages that view refugees as economic competitors compared to those who see them as economic partners? Or in contexts where baseline prejudice levels are higher and more widely-held? Answering these questions will require more expansive theory and data collection from more villages, within and beyond West Nile region of Uganda, to allow for comparative analysis.

Much is left to explore within communities and their networks as well. Among individuals, average short-term reactions to the treatment were positive. This average includes most who responded positively and a few who responded negatively. Average long-term reactions were also positive, but this aggregate is also comprised of some positive and some

negative reactions. Future work could build on the substantial reservoir of social science about prejudice to theorize and then identify who the backsliders are likely to be in advance and understand how networks can dampen negative spillovers that appear to originate with them. Ideally, future work will also explain who ultimately becomes more positive and who moves negative in response to a treatment, who is more susceptible to attitude shift from discussions within the network versus from the external stimulus of an intervention, and how the social network functions in these processes. Distinguishing how and why these in-person dynamics may differ from behavior in on-line networks is yet another promising avenue of inquiry. This paper lays the foundation for future research that can expand the theory and build new tests of these processes.

The authors declare the human subjects research in this article was reviewed and approved by George Washington University's IRB, Vanderbilt University's IRB, Uganda's National Council on Science and Technology, and Mildmay Uganda Research Center and certificate numbers are provided in the text.

The authors affirm that this article adheres to the APSA's Principles and Guidance on Human Subject Research.

The authors declare no ethical issues or conflicts of interest in this research.

This research was funded by George Washington University and UK International Development, awarded through Innovation for Poverty Action's Peace & Recovery Program.

Research documentation and data that support the findings of this study are openly available in the APSR Dataverse at <https://doi.org/10.7910/DVN/YZGMQJ>.

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.
- Anoll, Allison P and Andrew M Engelhardt. 2023. “A Drop in the Ocean: How Priors Anchor Attitudes Toward the American Carceral State.” *British Journal of Political Science* pp. 1–20.
- 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):1236–498.
- 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.

- Bartels, Larry M. 2002. “Beyond the running tally: Partisan bias in political perceptions.” *Political behavior* 24:117–150.
- Bénabou, Roland and Jean Tirole. 2016. “Mindful economics: The production, consumption, and value of beliefs.” *Journal of Economic Perspectives* 30(3):141–164.
- 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.
- 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.
- Clinton, Joshua D and Jason A Grissom. 2015. “Public information, public learning and public opinion: democratic accountability in education policy.” *Journal of Public Policy* 35(3):355–385.
- DeGroot, Morris H. 1974. “Reaching a consensus.” *Journal of the American Statistical association* 69(345):118–121.

- Ebere, Ruth Nnadi and David Mwesigwa. 2021. "The effectiveness of social inclusion of refugees in Kiryandongo refugee settlement community of mid-western Uganda." *Asian Journal of Sustainable Business Research* 2(2):76–94.
- 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* .
- Falk, Emily and Christin Scholz. 2018. "Persuasion, influence, and value: Perspectives from communication and social neuroscience." *Annual review of psychology* 69:329–356.
- 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.
- Galinsky, Adam D and Gordon B Moskowitz. 2000. "Perspective-taking: decreasing stereotype expression, stereotype accessibility, and in-group favoritism." *Journal of personality and social psychology* 78(4):708.
- Golub, Benjamin and Matthew O Jackson. 2010. "Naive learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics* 2(1):112–149.
- 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. 2019a. “The muted consequences of correct information about immigration.” *The Journal of Politics* 81(1):315–320.
- Hopkins, Daniel J, John Sides and Jack Citrin. 2019b. “The muted consequences of correct information about immigration.” *The Journal of Politics* 81(1):315–320.
- 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. 2023. “Which Narrative Strategies Durably Reduce Prejudice? Evidence from Field and Survey Experiments Supporting the Efficacy of Perspective-Getting.” *American Journal of Political Science* 67(1):185–204.
- 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. and Janet I. Lewis. 2024. “Reducing Prejudice towards Refugees in Uganda: Evidence that Social Networks Influence Attitude Change.”
URL: <https://doi.org/10.7910/DVN/YZGMQJ>
- 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.

- Mobius, Markus and Tanya Rosenblat. 2014. "Social Learning in Economics." *Annual Review of Economics* 6(1):827–847.
URL: <https://doi.org/10.1146/annurev-economics-120213-012609>
- Nambuya, Sharon Sylvia, Joseph Okumu and Ronald Pagnucco. 2018. "Refugee Socio-Cultural Integration and Peaceful Co-Existence in Uganda." *The Journal of Social Encounters* 2(1):81–92.
- 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>
- Sharot, Tali, Max Rollwage, Cass R Sunstein and Stephen M Fleming. 2023. "Why and when beliefs change." *Perspectives on Psychological Science* 18(1):142–151.
- 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.
- 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.
- Sinclair, Betsy. 2012. *The social citizen: Peer networks and political behavior*. University of Chicago Press.
- Tian, Ye and Long Wang. 2023. “Dynamics of opinion formation, social power evolution, and naïve learning in social networks.” *Annual Reviews in Control* .
- 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/>
- 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/>