Aid Cuts and the Politics of Refugee Hosting

Mae MacDonald*

July 15, 2026

Word count: 10,504

Abstract

In 2025, U.S. cuts to USAID led to the fastest contraction of foreign aid on record. Does exposure to these cuts erode host citizens' support for hosting refugees? I develop a theory in which exposure pulls in two directions: losses of aid-related jobs and services generate backlash, while proximity to refugees suffering under the cuts can sustain support through empathy. Testing this in Kenya, I combine two nationwide surveys, a high-frequency panel in a major hosting region, a vignette experiment, and 52 refugee-leader interviews. Across the evidence, local support for hosting falls among exposed citizens, driven by those who personally lost economic benefits. Yet there is no nationwide backlash, humanitarian concern is widespread, and high-empathy individuals maintain support even under a hypothetical complete aid withdrawal. These results suggest that continued cuts could erode local support for hosting in the low- and middle-income countries that host most of the world's refugees.

*Department of Political Science, Stanford University. Email: maemac@stanford.edu. ORCID: 0000-0002-9055-4810.

I am grateful to Jeremy Weinstein, Jens Hainmueller, James D. Fearon, Michael Allen, and Judith Goldstein for their invaluable advising and feedback. I thank Claire Adida, Sigrid Weber, Shelby Carvalho, Allison Grossman, Paige Hill, and participants in the 4th UNHCR–World Bank Joint Data Center Research Conference on Forced Displacement, 2025 American Political Science Association Annual Conference, U-Penn PDRI Forced Displacement Working Group, Stanford Immigration Policy Lab workshop, and Stanford IPE Workshop for helpful comments. I thank Mary Achieng, Michael Mwarange, Maggie Ireri, and the team at TIFA Research for outstanding work implementing the surveys, as well as Tolossa Asrat for excellent research assistance. I sincerely thank the participants in this research for sharing their experiences, particularly the refugee leaders who spoke with me during interviews. I used Claude (Anthropic) for copyediting, to assist with code writing and de-bugging, and to categorize open-ended responses. All analyses, interpretations, and any remaining errors are my own. This project was supported by funding from the J-PAL/IPA Displaced Livelihoods Initiative and the Stanford King Center on Global Development.

1 Introduction

In 2025, foreign aid contracted at the fastest rate on record: aid from OECD countries fell by 23.1 percent, largely driven by U.S. cuts to USAID spending (OECD, 2026). These cuts are expected to have severe consequences for aid-recipient populations, including an estimated 1.2 to 2.6 million excess deaths (Kenny and Sandefur, 2025). The world's 42.5 million refugees are exposed through two channels. First, cuts to major refugee organizations reduce access to food, medical care, schooling, and employment, worsening welfare and increasing preventable mortality.[1] Second, aid cuts may alter the politics of refugee hosting. 71% of refugees live in low- and middle-income countries (LMICs), where hosting often depends on both local communities' willingness to accept refugees and donor financing to sustain services (Betts, Loescher and Milner, 2012; Dempster and Ginn, 2026). When that funding is reduced, citizens and governments in asylum countries may become less willing to host refugees, increasing pressure for them to return to unsafe origin countries or move onward to third countries.

Motivated by such political consequences, this paper asks whether exposure to refugee aid cuts produces attitudinal backlash against refugees, defined as reduced support among locals for hosting refugees in their area. Existing research pulls in two directions. Unlike in many high-income countries, those most exposed to refugees in LMICs are often the most supportive of hosting them, in part because refugee presence can bring aid, jobs, and infrastructure (Zhou et al., 2025; Baseler et al., 2025; Demir, Mayda and Maystadt, 2025; Kadigo and Maystadt, 2023; MacDonald and Lichtenheld, 2025). On this account, withdrawing aid may lead to backlash where the aid economy is most concentrated, as citizens lose jobs and services or anticipate costs from hosting refugees with less international support (Hainmueller and Hopkins, 2014; Weber et al., 2025). A second literature suggests the opposite: refugee hosting is often supported because of humanitarian concerns, particularly in LMICs (Newman et al., 2015; Pettigrew and Tropp, 2006; Ghosn, Braithwaite and Chu, 2019). Aid cuts impose severe harms on refugees, particularly acute hunger. Those most exposed have the most contact with refugees and witness these harms firsthand, which may sustain or deepen support.

I build a formal model to show that, given these opposing dynamics, the net effect of aid cuts on attitudes is not obvious *ex ante*. The model predicts that backlash concentrates where exposure to the aid economy is highest and that it is driven by households who personally lose benefits, but that humanitarian concern pulls in the opposite direction, limiting backlash and, among the most empathetic, potentially increasing support. I

[1] The United States funded roughly 40% of the UN refugee agency (UNHCR) budget, contributing over $2 billion in 2024. Between 2024 and 2025, UNHCR's available funding fell by $1.2 billion (24 percent) to $3.9 billion, covering only 37 percent of assessed needs, and the agency cut roughly one third of its workforce (UNHCR, 2026*a*). WFP funding was projected to fall 34%, affecting up to 16.7 million people (WFP, 2025*a*). Recent studies document the harms that cuts to WFP rations, predating the 2025 cuts, impose on refugees (Grossman, Zhou and Carvalho, 2025; Sterck and Bruni, 2025).

test the theory against the 2025 U.S. aid cuts in Kenya, a major hosting country and one of the largest pre-cut U.S. aid recipients. I draw on two nationally representative phone surveys bracketing the cuts, a high-frequency face-to-face panel in the hosting region of Kakuma refugee camp during the height of the cuts, a pre-registered vignette experiment, and 52 refugee-leader interviews.

The analysis yields three findings. First, support for refugees declined in the regions most exposed to aid cuts. The literature's predictions about the economic losses hosts face are borne out: those closest to refugees show declines in support across the nationwide, panel, experimental, and qualitative data. Yet there was no nationwide backlash, with support even modestly increasing between 2023 and 2026. The net result is convergence, with hosting regions starting higher and moving toward the national average. Second, this regional backlash is driven by personal egocentric loss, not sociotropic economic concerns. Using a difference-in-differences design to isolate personal loss, I find the effect is concentrated among households that personally lose aid-related benefits rather than among their neighbors who do not lose jobs or services.

Third, suggestive evidence indicates that empathy can limit backlash, helping explain why nationwide support did not decline. Citizens often discuss their support for refugees in terms of humanitarian concerns, both at the nationwide and local level. In Kakuma, the main determinant of support for hosting is whether conflict is ongoing in the country of origin, rather than aid conditions, and high-empathy respondents do not change their support for hosting even in a hypothetical scenario where aid is entirely cut. The key takeaway of the results is that aid cuts can weaken support for refugee hosting where LMIC hosts depend on the aid economy. In the case of Kenya, the decline remains modest and localized, with empathy limiting its spread. Whether localized backlash can nonetheless reshape national hosting policy is a question I return to in the conclusion.

This study makes three contributions. First, it contributes to long-standing debates over foreign aid and to pressing policy questions raised by the U.S.' large-scale withdrawal of assistance. Given that existing research has focused on the provision of aid (e.g., Bermeo, 2017; Bearce and Tirone, 2010), less is known about the political consequences of sudden retrenchment in aid-recipient societies. I provide evidence on these questions, including interviews with refugee leaders documenting the severe, distressing consequences of aid cuts as they unfolded. Although the study is based in Kenya, the theory predicts backlash wherever refugees are tied to a visible local aid economy, especially in camps. Since one in five refugees worldwide lives in a camp, including across sub-Saharan Africa, Bangladesh, Jordan, Thailand, and elsewhere (Anti et al., 2025; UNHCR, 2026*b*), the findings raise broader questions about alternatives to camp-based hosting, particularly local integration policies that allow refugees to move and work (MacDonald and Lichtenheld, 2025; Hovil and Maple, 2022).

Second, the paper makes a theoretical contribution by explaining how exposure shapes attitudes when a shock harms hosts and refugees at the same time. Existing research typically treats exposure either as an economic channel or as a source of contact and empathy. Aid cuts move both channels at once: they impose losses on hosts while worsening conditions for the refugees they live alongside. I bring these channels into a single framework and use the aid-cut shock to examine whether sociotropic, egocentric, or humanitarian concerns drive attitudes. Against the prevailing view that sociotropic concern outweighs personal economic loss (Hainmueller and Hopkins, 2014; Weber et al., 2025), I find that egocentric loss and humanitarian concern matter most, and that insecurity matters more than economic cost. The framework applies beyond aid cuts to any shock that jointly harms a host population and a marginalized group with whom it lives in close contact.

Third, the paper provides real-time evidence on a major international shock. I designed the study as the cuts were announced but before their trajectory was known, fielding a high-frequency panel that ultimately captured both the sharpest contraction and partial improvements from stabilization. This was a setting with high uncertainty, potential insecurity, and restrictions on access, making data collection difficult. I therefore use a range of quantitative and qualitative evidence to triangulate the findings: no single design identifies the effect of the cuts on its own, but a panel with open-ended questions, two nationally representative surveys, a vignette experiment, and refugee-leader interviews converge on the same pattern. The design illustrates how fast-moving policy shocks can be studied as they unfold, and the value of triangulating across a mixed-method design.

The paper proceeds as follows. I first develop the theory and hypotheses, then describe the Kenyan setting, the sequence of aid cuts, and the research design. I next present results across the national and regional levels before discussing robustness checks and concluding.

2 Theory

Two literatures generate opposite predictions about how aid cuts might affect host attitudes toward refugees. The first anticipates backlash: aid cuts impose economic costs on hosts, both personally and on their communities, and where refugee support relies on the material returns of hosting, withdrawing aid may reduce it. By contrast, a second literature suggests that empathy may dampen backlash: the same cuts that impose costs on hosts inflict acute suffering on refugees, which may instead sustain or increase support.

2.1 Economic costs and backlash

In LMICs, refugee hosting can bring material benefits to host communities. Camps often bring international organizations and aid to remote border regions, generating jobs, business opportunities, and access to services for local residents (Alix-Garcia et al., 2018; International Finance Corporation, 2018; World Bank and UNHCR, 2021; Zhou, Grossman and Ge, 2023; Kadigo and Maystadt, 2023). Zhou, Grossman and Ge (2023) find that refugee settlements in Uganda improve local infrastructure, including schools, roads, and health clinics, while Alix-Garcia et al. (2018) show that refugee inflows in Kenya increased local economic activity, measured using nighttime lights. While most present in camp settings, there can be broader economic gains to hosting: aid can provide services for refugees and citizens in urban areas, support national public services and infrastructure, and become linked to wider development financing (Betts, Loescher and Milner, 2012; Dempster and Ginn, 2026). Consistent with this logic, greater exposure to the material returns of hosting is associated with higher support for refugees and incumbent politicians (Baseler et al., 2025; Zhou et al., 2025; MacDonald and Lichtenheld, 2025; Demir, Mayda and Maystadt, 2025).

Aid cuts threaten these benefits through two channels, both of which increase with exposure. The first is egocentric, concerning losses to a citizen's own household: locals employed by NGOs or refugee households may lose work, those using camp-area services may lose access, businesses selling to refugees may lose customers, and rising refugee deprivation may generate theft or resource conflict.[2] The second is sociotropic, concerning costs to the wider community or country: citizens who are not personally harmed may nonetheless worry that the state will assume more of the fiscal cost of hosting, that refugees will compete for public services, or that camp closures will displace people into urban areas. Evidence on immigration suggests that sociotropic concerns often outweigh personal economic vulnerability in shaping attitudes toward immigrants and refugees, but whether this extends to aid cuts is unknown (Hainmueller and Hopkins, 2014; Bansak, Hainmueller and Hangartner, 2023; Weber et al., 2025; Aviña et al., 2026).

Material loss does not necessarily imply backlash against refugees. Since hosting initially brought these benefits, affected citizens might favor both continued hosting and renewed aid. Yet backlash may occur for two reasons. First, hosts may not expect aid to return, and so may prefer that refugees leave rather than wait for employment and services to be reinstated. Second, losses may breed resentment toward refugees, particularly amid rising insecurity and crime, even when that response is not instrumentally rational. Loss aversion suggests that people react more sharply to losing benefits they already hold than to forgoing equivalent gains (e.g., Kahneman and Tversky, 1979). Hosts who lose out may therefore blame refugees directly and want them to leave, while those not personally harmed may still fault refugees for perceived costs to the wider

[2]This channel differs from the labor-market competition concern usually studied in the context of egocentric migration concerns, where the threat is that refugees take jobs or drive down wages. Here egocentric concerns are related to the loss of an existing refugee-related benefit, a reframing that reflects how refugees are associated with different economic costs and benefits in LMICs than in the high-income settings.

community.

2.2 Humanitarian concern and support

A second literature points in the opposite direction, suggesting that empathy and humanitarian concern can offset economic concerns about hosting refugees. Citizens are consistently more welcoming toward refugees, who have suffered severe physical or mental distress, than toward those fleeing economic hardship (Bansak, Hainmueller and Hangartner, 2016, 2023). Humanitarian concern reduces support for restrictive immigration policy, an effect that is stronger among more empathetic individuals (Newman et al., 2015). People who have themselves endured violence host more refugees and respond more strongly to refugee distress (Hartman and Morse, 2020; Hartman, Morse and Weber, 2021), and perspective-taking and empathy-priming exercises raise willingness to support refugees and immigrants (Adida, Lo and Platas, 2018; Williamson et al., 2021).

These studies suggest that humanitarian concern can increase support for refugees, and aid cuts are likely to activate this response. Reductions in food, healthcare, education, and cash assistance impose immediate and severe harms on refugees, particularly where they lack the legal right to work, move freely, or access state services. Aid cuts therefore differ from many policy shocks studied in the immigration literature: they impose economic costs on hosts while simultaneously severely worsening conditions for refugees, potentially generating greater humanitarian concern. This response may be strongest among the most exposed citizens. A large literature on intergroup contact finds that proximity and sustained interaction foster empathy and more favorable attitudes toward refugees (Pettigrew and Tropp, 2006; Ghosn, Braithwaite and Chu, 2019). Hosts in high-exposure regions are most likely to have contact with refugees and witness their suffering, and therefore may be most likely to respond with empathy.

2.3 Formalization

I formalize this theory as follows. Throughout, I use "exposure" to refer to geographic variation in communities' connection to the refugee aid economy. Household-level differences in aid-related benefits are described as "personal benefits" or "losses". Each host i's latent utility from refugee hosting in period t, H_{it}, combines three channels – egocentric, sociotropic, and empathy – capturing how much the host's own household, the community, and the refugees each gain or lose from hosting, weighted by how much i values them. Observed support is an increasing function of H_{it}.

$$H_{it} = \underbrace{\alpha_i \theta_{it}}_{\text{egocentric (household)}} + \underbrace{\beta_i E_i\big[\sigma A_t - C(A_t)\big]}_{\text{sociotropic (community)}} + \underbrace{\gamma_i\big[E_i S(A_t) + \tfrac{1}{2} S(A_t)\big]}_{\text{empathy (refugees)}} \tag{1}$$

The egocentric channel captures the host's household welfare: α_i is how much i values it, and $\theta_{it} \in \{-1, 0, 1\}$ indicates whether i's household holds an aid-related benefit or bears an aid-related cost in period t. Aid cuts enter through discrete transitions ($\theta_{i1} = \theta_{i0} - 1$) as a household loses a benefit or absorbs a new cost; households not directly affected retain $\theta_{i1} = \theta_{i0}$. $E_i \in [0, 1]$ is community exposure to the refugee aid economy, equal to one near camps and lower but strictly positive elsewhere. Personal loss and community exposure are not independent: I denote by $p(E_i)$, increasing in E_i, the probability a household undergoes $\theta_{i1} = \theta_{i0} - 1$ given a cut, so residents of high-E_i areas are more likely to lose out personally.

The sociotropic channel captures community welfare. β_i is how much i values it, E_i is community exposure, and A_t is total aid. The community-level payoff is $\sigma A_t - C(A_t)$, where $\sigma \in [0, 1]$ is the share of aid reaching host communities and $C(A_t) = \bar{c} - A_t$ is the cost of hosting ($\bar{c}$ the maximum, at $A_t = 0$), defined for aid low enough that cost and suffering stay non-negative. The community receives net benefits as long as $\sigma A_t > C(A_t)$, with crossover at $A^* = \bar{c}/(\sigma + 1)$: above A^* it still gains from a contracting aid economy, but below it incurs net costs.

The empathy channel captures the host's response to refugee suffering. Whereas the household and community weights α_i and β_i are non-negative, γ_i, the weight i places on refugee welfare, can take either sign: $\gamma_i > 0$ indicates empathy, $\gamma_i < 0$ discrimination or hostility. I do not require the weights to sum to one, since γ_i can be negative and only their relative size shapes the results ($\beta_i(\sigma + 1)$ against γ_i). E_i enters again, capturing contact with refugees and awareness of their conditions, and $S(A_t) = \bar{s} - A_t$ is the level of refugee suffering ($\bar{s}$ the maximum, at $A_t = 0$). The additive $\frac{1}{2}S(A_t)$ term lets even unexposed hosts respond through media coverage and indirect social ties, though less strongly than the exposed.

The cut works through aid, which falls from a pre-cut level A_0 to a lower level A_1, while the weights and exposure describing each host and their area stay fixed over the short study window. The object of interest is therefore the change in latent utility across the cut, $\Delta H_i = H_{i1} - H_{i0}$, and the payoff of the formalization is that community exposure, E_i, enters all three channels. It shapes the egocentric channel indirectly, through $p(E_i)$: a household is more likely to lose its benefit to a cut where exposure is high. It enters the sociotropic and empathy channels directly, and with opposing effects. Through the sociotropic channel, the decline from A_0 to A_1 reduces support by shrinking host benefits and raising host costs. Through the empathy channel, the same decline raises support by worsening refugee suffering and strengthening empathic concern. Community exposure amplifies all three, so the net response depends on whether the host suffers a personal loss and on whether $\beta_i(\sigma + 1)$ or γ_i is larger.

2.3.1 Empirical implications

The model yields three empirical implications when aid declines, which motivate the hypotheses tested below. Comparative statics and full derivations appear in the Appendix.

First, aid cuts should reduce hosting support more in high-exposure communities. In these communities, cuts are more likely to produce personal losses, generating an egocentric decrement of $-\alpha_i p(E_i)$. Greater exposure also amplifies both the sociotropic costs of reduced aid and the countervailing effect of empathy. On balance, the egocentric and sociotropic channels make backlash more likely in high-exposure than in low-exposure communities, unless empathy is sufficiently strong to offset them.

H1 (Community exposure): *Aid cuts reduce hosting support more in high-exposure than in low-exposure populations.*

Second, within high-exposure communities, backlash should be greater among hosts who personally lose aid-related benefits. Holding community exposure and other characteristics constant, a household that absorbs a decrement in its aid-related position ($\Delta\theta_i = -1$) makes the attitudinal response more negative by α_i relative to a household with no change.

H2 (Personal loss): *Within an exposed community, hosts who lose aid-related benefits show a steeper decline in hosting support than otherwise-similar neighbors.*

Third, empathy should moderate backlash. Holding community exposure and personal loss constant, the response to an aid cut becomes more favorable as γ_i increases. More empathetic hosts should therefore exhibit less backlash, and sufficiently strong empathy may offset or reverse the decline in support.

H3 (Empathy): *Holding community exposure and personal loss constant, aid cuts reduce hosting support less among more empathetic hosts. Where empathy is sufficiently strong, support may increase.*[3]

3 Context

3.1 Refugee hosting in Kenya

I examine the effects of aid cuts in Kenya, a strong setting to test the theory for three reasons. First, Kenya hosted approximately 800,000 refugees and was among the largest recipients of USAID assistance before the cuts (U.S. Department of State, 2025). Second, exposure to refugee assistance is geographically concentrated, with

[3]H3 was not pre-registered; see Section 4 and the appendix.

Figure 1: Refugee-hosting areas and survey sample

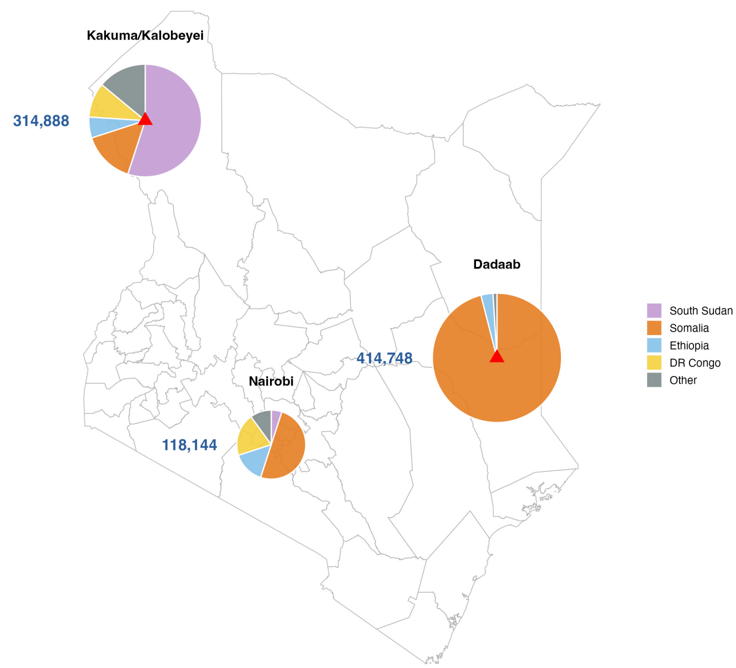

(a) Refugee populations in Kenya

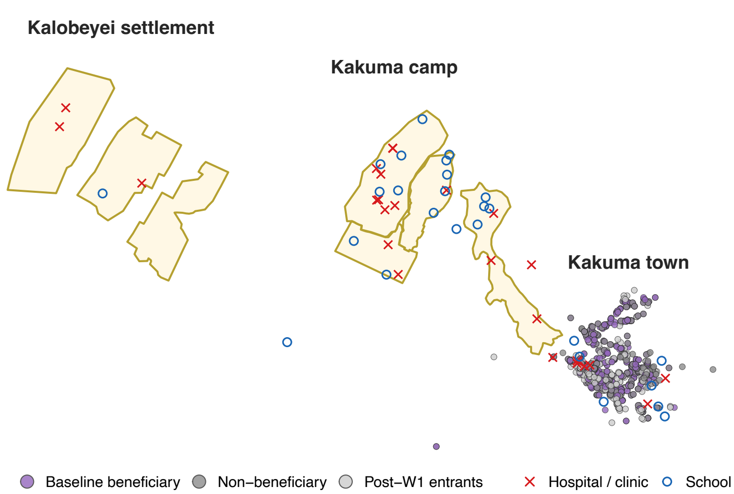

(b) Sample around Kakuma, by exposure status

Notes: Panel (a) shows refugee populations by hosting location and origin country. Panel (b) plots survey respondent GPS locations by baseline-beneficiary status (purple = baseline aid recipient or NGO/aid employee; grey = not baseline aid/employed; light grey = post-W1 entrants to the survey in RCS). Yellow polygons show camp and settlement boundaries.

the vast majority of refugees, mostly from South Sudan and Somalia, living in the Kakuma and Dadaab camps (see Figure 1a; UNHCR, 2025). Camp-based refugees depend heavily on international assistance, given restrictions on movement outside the camps and limited employment opportunities (MacDonald and Lichtenheld, 2025).[4]

Third, existing research shows that the area has benefited economically from refugees. Turkana is a semi-arid border county among Kenya's poorest and most politically marginalized regions: 72% of residents live in poverty, against 68% among refugees and 37% nationwide (World Bank and UNHCR, 2021). Before the camp was established in 1992, the surrounding area was sparsely populated; refugees, camp infrastructure, and international aid contributed to the growth of an urban center that is more developed than the rest of the county (International Finance Corporation, 2018; Alix-Garcia et al., 2018; Sanghi, Onder and Vemuru, 2016; World Bank and UNHCR, 2021).

Hosts in Kakuma benefit from the camp economy through three channels, which I describe using evidence from the panel survey (see Section 4) and World Bank and UNHCR (2021). First, hosts access refugee services, particularly health facilities and schools in the camp, shown in Figure 1b. 69% of Turkana hosts have no formal education (compared to 47% of refugees), and only 40% can read and write (78% for refugees). Hosts account for an estimated 10–15% of beneficiaries of camp healthcare services. As one respondent put it, "Since the refugees came, there has been help. Schools and hospitals were built" (Woman, 21).[5]

Second, hosts find employment with NGOs and with refugee households, particularly as casual domestic laborers. The Turkana hosts are "asked to fetch water, wash the dishes, do the laundry, and after doing the work [they] get something to go and eat at home with [their] kids" (Woman, 48).[6] Third, hosts run businesses that depend on refugee demand, including shops and the sale of local resources. One respondent described: "Most of [us] take firewood to the camps and sell it to the refugees, and some take charcoal. After that, [we] are given food [by the refugees] and feed [our] kids" (Woman, 23).[7] In May 2025, a large majority of panel respondents (71%) agreed that locals benefit from the camp.

Table 1: Aid shocks and data collection

Date	Event	Description	Captured by	Trajectory
2023–2025	WFP reductions	General food basket cut in stages from 80% to 40% of minimum daily needs: 80→60% Jul 2023; 60→50% Feb 2024; 50→45% May 2024; 45→40% Feb 2025.	Natrep 2023 → 2026 (bundled)	Decline
Jan–Feb 2025	USAID cuts	USAID-funded NGOs reduced services and laid off staff across the camp.	Natrep 2023 → 2026 (bundled)	Decline
Jun 2025	Cash assistance suspended	Cash grants fully suspended; 40→30% food basket for all residents.	W1 → W2; Natrep	Decline
Jul 2025	DA announced	Vulnerability-tiered Differentiated Assistance model announced; sparked protests and delayed food distribution.	W1 → W2; Natrep	Trough
Aug 2025	DA takes effect	Tiered rations begin: Cat 1 at 40%, Cat 2 at 20%, Cat 3–4 receive nothing; cash still suspended; increased police presence.	W2 → W3; Natrep	Stabilization
Sep–Oct 2025	Cash resumption announced	Resumption of cash for Oct–Nov announced. Reduced insecurity.	W2 → W3; Natrep	Stabilization
Oct–Nov 2025	Rations increased, cash resumes	Rations raised (Cat 1 55%, Cat 2 35%, Cat 3 20%); cash resumes for Categories 1–3.	W3 → W4; Natrep	Partial reinstatement

3.2 Aid cuts

Refugee assistance in Kenya was cut in three stages (see Table 1 for details). First, WFP reduced food assistance from 80 percent to 40 percent of minimum daily needs between July 2023 and February 2025 (Anyadike, 2024).[8] Second, after the U.S. cut foreign assistance worldwide in 2025, rations fell again and organizations reduced staff and services. These cuts affected the main NGOs serving Kakuma, including UNHCR,

[4]Thirteen percent of refugees live in urban areas, particularly Nairobi. This group is more self-sufficient and less dependent on aid.

[5]Additional quotes: refugees enabled hosts to "access basic needs like health from the UN" (Man, 31); "[we] get free treatments from the camp hospitals and free schooling for [our] children" (Man, 32).

[6]Additional quote: refugees "boosted the economy by introducing NGOs, creating job opportunities" (Man, 27).

[7]Additional quote: "Kakuma's revenue has increased over time as compared to Lodwar due to a lot of businesses going on about and a lot of people duly paying their taxes" (Woman, 34).

[8]The U.S. previously funded a large share of this food assistance: in fiscal year 2024, the U.S. provided 47.0 percent of WFP's 2023–2027 contributions in Kenya and 45.8 percent of UNHCR's contributions (WFP, 2025*b*; UNHCR, 2024). Sterck and Bruni (2025) show that a 20 percent cash reduction during this period harmed refugees, though they do not estimate effects on hosts.

which funds many other camp organizations, IRC's healthcare centers, and LWF's education programs.[9]

The third stage, captured by the panel, marked the trough of the cuts. In June, WFP suspended cash transfers, eliminating the \$3.09–\$3.98 per refugee that had helped cover fresh food, hygiene products, firewood, school fees, and other costs. In July, WFP announced a Differentiated Assistance (DA) model that classified refugees into four vulnerability tiers.[10] When the model was rolled out in August, almost half of registered refugees received no rations or cash. Refugee leaders reported that many vulnerable households were misclassified as self-sufficient, producing acute hunger and malnutrition. The announcement triggered protests and delayed food distribution by two weeks, leading to increased police presence. By late August, there was stabilization among the refugee community, with reduced insecurity and the announcement of aid being partially reinstated for some groups in October. This was not a full reinstatement of aid, however. Cash assistance and rations resumed at lower levels than before, Category 4 refugees still received no assistance, and NGO staffing did not return to prior levels.

Interviews with refugee leaders provide qualitative data on the harm these cuts caused, particularly in July 2025. The central problem was widespread hunger, with many refugees eating "once per day". To cope, they skipped meals, took on debt, resorted to crime, and sold their assets. One mother was forced to "sell the roof of the house . . . so that she can get food for her kids"; adolescent girls "[slept] with a man to get [sanitary] pads"; and "people have been cut with panga [knives] just because of a smart phone". Schools and clinic access became more strained, the main issue being a lack of medication: "Now we have no drugs in the hospital. They're just giving you painkillers. . . Malaria is the biggest killer". Refugees in Kakuma cannot move elsewhere in Kenya for work without formal permission, nor grow their own food, because "it's arid, dry, [and] there's no water." Some returned to South Sudan despite the risk of conflict, reasoning that it is "better to die in your country with a bullet than hunger". By late September, leaders noted the situation had improved, though not fully: "Comparing August and now, things are much better. August was really really bad. People had no food, there was a lot of tension and theft also. [Now I] feel like things are not good all time, but at least, you're not hearing stories of someone starving and rushed to the hospital".

Host citizens were also affected through lost employment, services, and the food previously exchanged between refugees and hosts. As one host citizen said, "life in the host

[9] In fiscal year 2024, prior to the cuts, the U.S. disbursed \$116.7 million to WFP, \$36.6 million to UNHCR, \$4.8 million to IRC, \$0.7 million to LWF, and \$6.0 million to other organizations serving refugees in Kakuma (U.S. Department of State, 2025; OCHA, 2025). Author calculated these statistics using U.S. government data.

[10] For evidence on the impact of WFP's prioritization model on refugees in Uganda, see Grossman, Zhou and Carvalho (2025).

Figure 2: Host livelihood outcomes across the Kakuma panel

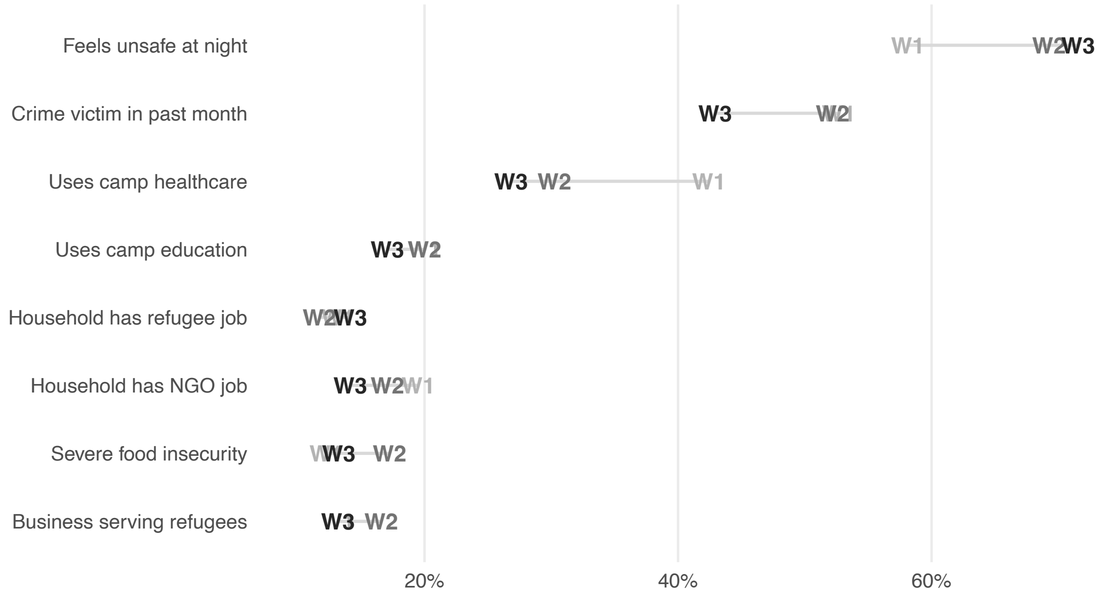

Notes: Percentage of Kakuma panel households reporting each outcome at each wave. Services, employment, and crime are asked for the previous month. Severe food insecurity: missed meals on 5–7 days in the past week. Feels unsafe at night: disagrees or strongly disagrees with feeling safe. Rows ordered by W3 prevalence.

community has started to become difficult, there is no food or work," and a refugee leader echoed that "the situation is becoming worse on both sides". Figure 2 summarizes the household-level changes across waves.[11] The clearest shift was a sharp rise in feeling unsafe at night, reflecting insecurity from the protests, rising crime and theft, and conflict as refugees collect charcoal and firewood they previously bought from the locals. There was also a sustained decline in camp healthcare access, with smaller decreases in education access. On economic outcomes, hosts reported lost income, increased food insecurity, lost NGO employment, and lost businesses serving refugees. Even though there was an announcement of more aid for some refugees between W2 and W3, this largely did not improve outcomes for hosts: of the panel households that lost access to camp-area education, healthcare, refugee employment, or NGO employment between W1 and W2, only 8% had regained access in that domain by W3. The main improvement between W2 and W3 for hosts was in reduced crime victimization, showing stabilization in the wake of the W2 protests and insecurity.

[11] Tests are linear probability models of each outcome on wave indicators, with standard errors clustered by respondent; see Appendix. Of these changes, only the rise in feelings of insecurity, the decline in camp healthcare access, and the W3 decrease in crime victimization reach statistical significance ($p < 0.05$). Given the small per-wave sample, the remaining estimates should be read as imprecise rather than as evidence of no change.

4 Research Design

4.1 Data sources

I draw on three data sources to examine the effects of the aid cuts in Kenya on attitudes toward refugees. First, I fielded two original nationally representative phone surveys. The baseline was conducted in September–October 2023, with 3,326 respondents: a nationally representative sample of 2,432 and oversamples near Kakuma and Dadaab camps (441 in Garissa and 453 in Turkana).[12] The follow-up was fielded in February 2026, after the major 2025 cuts, with a nationally representative sample of 2,042 and a Turkana oversample of 289.[13] In both surveys, I construct post-stratification weights using entropy balancing, matching each sample to the contemporaneous Afrobarometer distribution of age, gender, education, and region (Hainmueller, 2012).[14]

Second, I fielded three waves of a face-to-face panel survey with citizens living next to Kakuma camp.[15] Although designed before the trajectory of aid was known, the panel tracks citizens through the main aid cut shock, from pre-trough conditions in May to the trough in July and minor improvements during stabilization in September. The main analysis uses 205 respondents present in all three waves. The appendix reports parallel results using the repeated cross-section, refreshed to at least 300 citizens per wave. A fourth phone wave, fielded alongside the 2026 nationwide survey, captures longer-term change and partial aid reinstatement but has higher attrition and is reported in the appendix. Figure 1b maps the camp, Kakuma town, panel respondents, and schools and health facilities.

Third, I conducted 52 semi-structured interviews with a panel of refugee leaders alongside the Kakuma survey waves. The interviews provide qualitative evidence on refugees' suffering, host-community spillovers, and the evolution of the cuts between waves. Seventeen leaders remained in the panel across all three waves (18 in Wave 1 and 17 in Waves 2 and 3).[16] This evidence is informed by previous in-person interviews in Kakuma in 2022 and fieldwork across multiple years in Kenya.

[12]Because the 2023 survey experimentally varied whether respondents were asked about "refugees" in general or specific refugee groups (for a separate paper), over-time comparisons with the 2026 survey use only the generic-"refugees" arm for the *Nation Hosting* and *Local Hosting* variables ($N = 810$). See appendix.

[13]A May 2026 follow-up corrected a technical issue in the nationwide hosting items only ($N = 1{,}776$ nationally representative).

[14]Round 9 (2021) for the 2023 baseline and Round 10 (2024) for the 2026 follow-up.

[15]I use "Kakuma" as shorthand for Kakuma camp and neighboring Kalobeyei settlement. I was unable to study refugees in Dadaab due to security issues at the time of the study.

[16]Leaders were purposively sampled to vary by nationality, gender, and leadership role, including healthcare workers, local news broadcasters, leaders of refugee-led organizations, and employees of international organizations. Respondents were recruited with the help of a local journalist and research assistant with strong networks in Kakuma. Details are in the appendix.

4.2 Empirical strategy

This section lays out the empirical strategy for each hypothesis. I estimate all regressions by OLS, using a linear probability model for binary outcomes. Standard errors are HC1 for cross-sectional specifications and clustered by respondent for panel specifications. Full regression equations appear in the Appendix.

4.2.1 H1: Community exposure

Nationwide difference-in-differences. I begin by examining whether attitudes toward refugees in the country's major hosting regions shifted after the cuts relative to the country as a whole. Using the two nationwide surveys, I examine change over a multi-year window, from 2023 to 2026, that spans the majority of the aid cuts, including most of the WFP ration reductions and the major 2025 cuts. The design exploits geographic variation in exposure to the aid economy: because community exposure is concentrated near the camps, comparing over-time change in the hosting regions with change nationwide isolates the differential effect of the cuts on the most exposed areas. Under H1, the hosting regions should exhibit more negative change than the rest of the country, given their higher rates of personal loss and community exposure, unless empathy is strong enough to offset the egocentric and sociotropic losses.

The outcome is the *Hosting Index*, an inverse-covariance-weighted index (Anderson, 2008) combining three items. *Nation Hosting* and *Local Hosting* are five-point scales measuring support for hosting refugees in Kenya and in the respondent's district, respectively. *Not Repat.* is an indicator equal to one if the respondent selects any option other than repatriation – the most restrictive option, defined as returning refugees to their countries of origin – from among five refugee-policy alternatives: full rights, settlements, camps, repatriation, or third-country resettlement. All components are coded so that higher values indicate more supportive attitudes. The index is standardized to the 2023 national baseline (mean 0, SD 1), so coefficients represent standard-deviation changes relative to pre-cut attitudes. I also examine these items separately.

I estimate a pooled difference-in-differences comparing pre-cut (2023) and post-cut (2026) attitudes in the main refugee-hosting counties, Turkana and Garissa, with the rest of the country:[17]

$$Y_i = \alpha + \beta \, \mathrm{Post}_i + \rho \, \mathrm{Host}_i + \tau \, (\mathrm{Post}_i \times \mathrm{Host}_i) + X_i'\Gamma + \varepsilon_i, \tag{2}$$

where Host_i indicates a refugee-hosting county, and X_i controls for age, gender, and education. The coefficient β captures the adjusted change outside hosting counties,

[17] Note that the sample size is larger for Turkana residents in the 2026 sample than Garissa residents, given that there was no Garissa oversample in this wave. The difference-in-differences is estimated unweighted.

where community exposure is lower, while τ captures the differential change in hosting counties. H1 predicts $\tau < 0$ on the *Hosting Index*.

Panel over-time change. Because the nationwide difference-in-differences spans a long window, I next narrow the focus to the hosting region itself, using the high-frequency panel to ask whether support for refugees falls within a single high-exposure community as aid is cut. Here identification comes from time variation rather than cross-regional comparison: I track within-person attitudes across three closely spaced waves, with baseline conditions at W1, the trough of the contraction at W2, and continued (though partially eased) cuts at W3. Under H1, attitudes in the hosting region should worsen once aid is cut. The design does not support a formal causal test, but the high frequency of the panel lends confidence that attitudinal shifts track the specific aid changes documented in Section 3.2 rather than slower-moving confounders.

I focus on two outcomes specific to local hosting, both of which capture how residents feel about continuing to host refugees in their area. *Oppose Camp Closure* measures disagreement with "the Kenyan government closing Kakuma camp and moving the refugees in Kakuma back to their countries," recorded on a five-point scale and coded so that higher values indicate pro-refugee attitudes. *Local Hosting* is the same district-level hosting item described above.

I estimate an event-study specification on the balanced panel,

$$Y_{it} = \mu + \nu_i + \delta_2 D_{2t} + \delta_3 D_{3t} + \varepsilon_{it}, \tag{3}$$

with respondent fixed effects ν_i, wave indicators D_{2t} and D_{3t} (W1 omitted), and standard errors clustered on the respondent. The coefficient δ_2 is the cumulative change from baseline to the trough of the contraction, and the linear contrast $\delta_3 - \delta_2$ measures the change from the trough to the stabilization period. H1 predicts $\delta_2 < 0$ for both *Local Hosting* and *Oppose Camp Closure*.

Vignette experiment. The two analyses above are descriptive. To test causally whether aid cuts reduce support for refugees, I embed a factorial vignette experiment in W3 of the Kakuma panel. The experiment asks whether residents of the hosting region withdraw support when the aid attached to hosting changes. Respondents rate their support for Kenya hosting refugees (the *Nation Hosting* measure) under scenarios that randomly vary four attributes: whether international organizations provide 100, 40, or 0 percent of food rations and cash support (*Aid*); whether healthcare and schools are shared with the host community or reserved for refugees only (*Services*); whether international organizations hire local residents (*Hiring*); and whether conflict is ongoing or peace has been restored in refugees' country of origin (*Conflict*). The $3 \times 2 \times 2 \times 2$ design yields 24 possible profiles. To limit respondent fatigue, each respondent rates six profiles drawn at random. I estimate average marginal component effects (AMCEs) with respondent fixed effects and standard errors clustered by respondent. Respondents are asked a comprehension check concerning the level of aid in the last scenario they rate.

The attribute of primary interest for H1 is *Aid*: support for hosting should fall as aid drops from 100 percent to lower levels. I also expect support to decline when services are restricted to refugees (*Services*) and when organizations stop hiring locally (*Hiring*). The *Conflict* attribute tests H3 and is discussed below. I complement the experiment with a descriptive item, fielded in W1 and W2, on support for hosting if aid were withdrawn entirely (details in Appendix).

4.2.2 H2: Personal loss

Where H1 concerns whether community exposure is associated with backlash, H2 asks whether the backlash within high-exposure communities is driven by the households that personally lose aid-related benefits. I test this with an intent-to-treat difference-in-differences design comparing households vulnerable to personal loss with those that are not. I define the treated group by baseline attributes: a household is a "baseline-beneficiary" if, at W1, the respondent accessed camp healthcare or education services or was employed by an NGO or a refugee in the past 30 days. Because status is fixed at baseline, before the severe cuts, the design avoids post-treatment bias and is likely conservative, since not every baseline beneficiary lost benefits. The main outcomes are *Oppose Camp Closure* and *Local Hosting*, the two local hosting measures, and I also report *Nation Hosting* (defined above). H2 predicts that baseline beneficiaries backlash, moving toward support for closure, while non-beneficiaries do not.

I interact a baseline-beneficiary indicator with wave dummies on the balanced panel, with respondent fixed effects:

$$Y_{it} = \mu + \nu_i + \delta_2 D_{2t} + \delta_3 D_{3t} + \tau_2 \left(B_i \times D_{2t}\right) + \tau_3 \left(B_i \times D_{3t}\right) + \varepsilon_{it}, \tag{4}$$

where B_i equals one if the household received camp healthcare, education, or employment in the 30 days before W1. The coefficients τ_2 and τ_3 capture the differential change among baseline beneficiaries at the trough and stabilization waves, and H2 predicts $\tau_2 < 0$ and $\tau_3 < 0$.

Selection into beneficiary status is non-random: households receiving aid-related benefits before the cuts were more supportive of refugees at baseline than those that were not. Respondent fixed effects absorb this baseline gap, but cannot rule out that the two groups were already on different trajectories, so I read these estimates as quasi-causal and triangulate with two further designs, detailed in the Appendix and Section 6. The first defines treatment as households that actually lost a benefit, comparing those that lost access between W2 and W3 with those that never lost access, using the W1–W2 window to check for parallel pre-trends. The second pools all loss events, estimating how attitudes change when a household switches into or out of benefit access between any adjacent waves.

4.2.3 H3: Empathy

H3 holds that empathy can dampen backlash, or even raise support for refugees, in the wake of an aid cut. I examine this in two ways: first, by testing whether more empathetic hosts respond less negatively to cuts; and second, by assessing whether humanitarian concern is sufficiently salient among Kenyans more broadly to shape attitudes toward refugee hosting.

Empathy heterogeneity in the vignette experiment. I test whether responses to the *Aid* attribute in the vignette experiment differ by respondents' empathy for refugees. I measure empathy with a standardized index averaging two five-point items: concern about refugees and their suffering, and worry that refugees lack enough food because of the cuts. I interact *Aid* with this empathy index in the vignette regression, again with respondent fixed effects and standard errors clustered on the respondent. H3 predicts a positive interaction: high-empathy respondents should have a less negative aid slope than low-empathy respondents, withdrawing support less as aid falls. Three limitations apply. Empathy is not randomly assigned, so the interaction does not identify a causal effect of empathy; empathy is measured post-treatment at W3, alongside the experiment, rather than before it; and H3 was not pre-registered, having emerged from the open-ended responses and refugee-leader interviews, so I test it as an exploratory analysis. I therefore read the results as suggestive.

Levels of humanitarian concern. I next gather evidence on whether humanitarian concern is salient enough to shape refugee attitudes in Kenya more generally. This does not test how empathy moderates responses to aid cuts, as the theory specifies, but it establishes whether humanitarian concern is widespread enough to plausibly matter. I do this in two ways. First, I use the *Conflict* attribute in the vignette experiment, which varies whether conflict is ongoing or peace has been restored in refugees' country of origin. If there is empathy for refugees, respondents should be more supportive of hosting when conflict is ongoing, as it would be unsafe for refugees to return under these conditions. By comparing the magnitude of the *Conflict* effect to that of the *Aid* effect, I examine whether hosts respond more to refugee need or to the aid environment.

Second, I draw on three open-ended questions to gauge the salience of humanitarian concern. In the Kakuma panel, respondents were asked what changes they had seen from the cuts (W1–W3) as well as why they oppose or support closing the camps (W3). Nationwide, the 2023 survey asked why respondents support or oppose hosting.[18] H3 predicts that there will be frequent references to refugee suffering and protection concerns in Kakuma, as well as humanitarian reasons to support refugee hosting nationwide.

[18] I create a set of categories for these questions by examining 10% of responses and then use Claude to code the remaining responses into these categories, which I subsequently check. See appendix.

5 Results

5.1 H1: Community exposure

H1 predicts that as aid is cut, support for hosting falls further in high-exposure communities than in low-exposure ones, unless empathy offsets the egocentric and sociotropic losses. The evidence supports H1 consistently across the nationwide, panel, experimental, and descriptive analyses, with the nationwide difference-in-differences and vignette experiment providing the strongest support.

Nationwide difference-in-differences. The nationwide difference-in-differences analysis finds strong support for H1. Table 2 shows that hosts in the refugee-hosting counties moved against refugees relative to the rest of the country by roughly a third of a standard deviation ($\hat{\tau} = -0.365$, $p < 0.001$). The decline appears across all three components: it is largest for local hosting ($\hat{\tau} = -0.659$, $p < 0.001$) but also significant on nationwide hosting ($\hat{\tau} = -0.418$, $p < 0.05$) and not selecting repatriation as preferred policy ($\hat{\tau} = -0.099$, $p < 0.01$). Figure 3 shows regional shifts in support descriptively.

Because hosting regions began above the national average ($\hat{\rho} = 0.326$, $p < 0.001$), the shift represents convergence toward the national mean from a higher baseline rather than hosting regions becoming more hostile than the rest of the country. The coefficient for the non-hosting regions $\hat{\beta} = 0.114$ ($p < 0.01$) aligns with the prediction that where exposure is low and empathy is high, there is no backlash and even a slight increase in support, though the increase is substantively small and largely driven by the repatriation-choice measure.

Panel over-time change. The Kakuma panel traces within-person attitudes as aid contracted (W1 to W2) and conditions stabilized (W2 to W3). This analysis finds some support for H1. Table 3 (Panel A) shows a decline in opposition to camp closure from W1 to W2 ($\hat{\delta}_2 = -0.260$, $p < 0.10$), consistent with backlash. This is only weakly significant, though this may be due to the small sample size. There is no significant decrease in support for local hosting over the aid cuts period. Neither measure changes significantly between W2 and W3, meaning that elevated support for closure persisted through to W3.

Vignette experiment. The experiment finds support for H1. Figure 4 reports the AMCEs for the vignette experiment and shows that all three reductions in aid for hosts move attitudes in the predicted direction: cutting aid to zero lowers hosting support by 0.19 points ($p < 0.05$), removing local hiring by 0.20 ($p < 0.01$), and restricting services to refugees by 0.18 ($p < 0.05$). Reducing aid to 40 percent produces a smaller decline of 0.14 that is weakly significant ($p < 0.10$). This is supported by an additional descriptive question asking support for hosting under a full aid cut scenario, detailed in the appendix. In W2, support for hosting refugees drops by 0.84 points (on a 1–5

Figure 3: Hosting Index across Kenya, 2023–2026

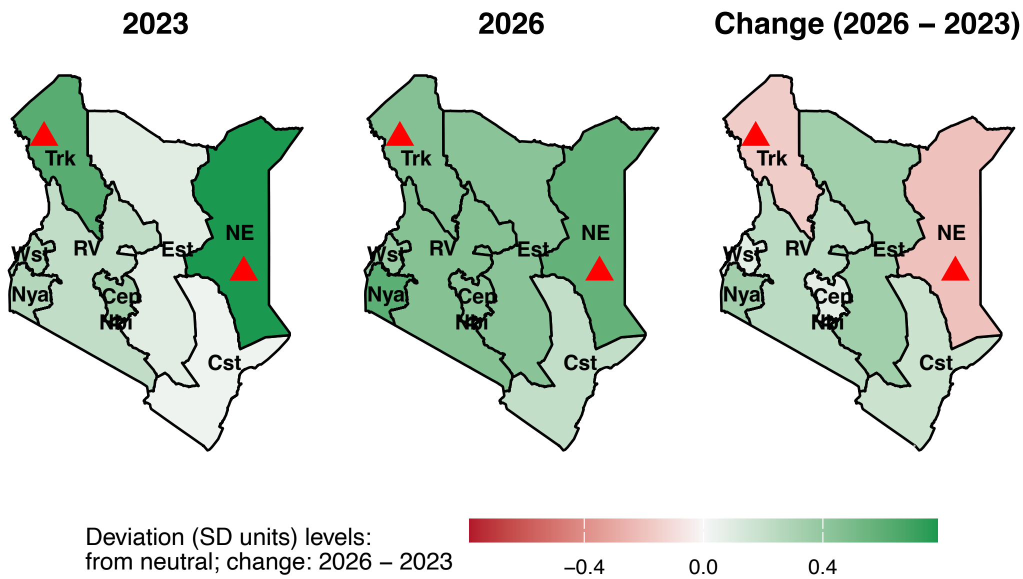

Notes: Signed deviations in the *Hosting Index* by region. The 2023 and 2026 panels show deviations from the substantive neutral point (all items at their midpoints). The third panel shows the 2026–2023 change. Provinces are weighted by entropy-balancing weights matched to the contemporaneous Afrobarometer distribution on age, gender, education, and region. The Turkana overlay is unweighted and includes the Turkana oversamples. Red triangles mark Kakuma and Dadaab camps.

Table 2: Pooled difference-in-differences: nationwide refugee attitudes, 2023–2026

	Hosting Index	Nation Hosting	Local Hosting	Not Repat.
Post	0.114**	-0.027	0.099	0.079***
	(0.040)	(0.066)	(0.072)	(0.019)
Hosting county	0.326***	0.562***	0.381**	0.109***
	(0.063)	(0.095)	(0.120)	(0.028)
Post × Hosting county	-0.365***	-0.418*	-0.659***	-0.099**
	(0.081)	(0.196)	(0.155)	(0.036)
Controls	Yes	Yes	Yes	Yes
Observations	3,401	2,841	3,381	3,357
R^2	0.040	0.023	0.034	0.033

Notes: OLS with HC1 standard errors. Post coefficient gives nationwide change outside hosting counties; Post × Hosting county gives differential change in Turkana and Garissa. All outcomes coded so higher = more pro-refugee. Demographic controls: age, gender, education. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

Figure 4: Vignette experiment

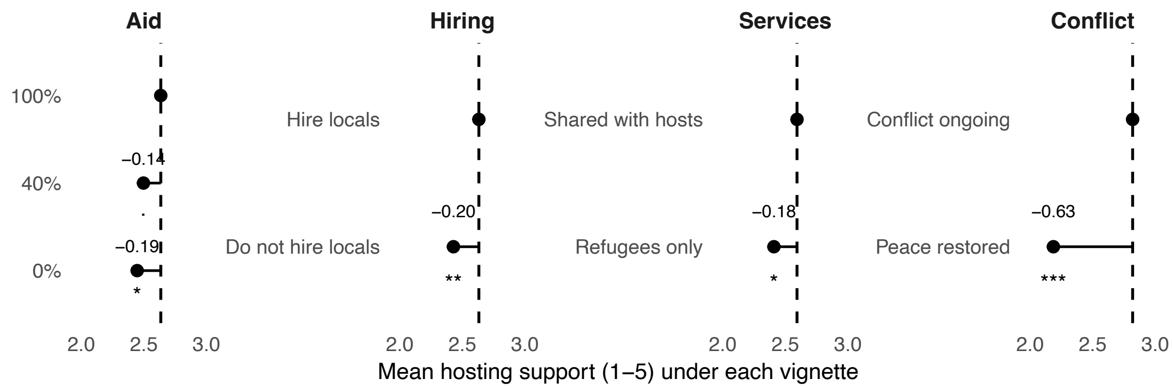

Notes: AMCEs from Wave 3 factorial vignette (Kakuma, $N = 299$ respondents, 1,883 profiles). Points show mean *Nation Hosting* at each attribute level; dashed line marks reference. Linear model with respondent fixed effects, SEs clustered by respondent. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

scale) under conditions of no international aid, with only 29.4% of respondents willing to continue hosting under such conditions.

5.2 H2: Personal loss

H2 predicts that within the high-exposure community, hosts who personally lose aid-related benefits show an additional decline in hosting support relative to otherwise-similar neighbors who share the same community exposure but face no direct household losses. The evidence supports H2: the decline is concentrated among households that personally lose aid-related benefits, corroborated by two further specifications in the Appendix.

Panel B of Table 3 supports H2. Attitudes among non-beneficiaries changed little across Waves 2 and 3, suggesting that hosts without direct ties to the aid economy did not backlash against refugees as the cuts unfolded. Beneficiary households, by contrast, were substantially more supportive before the cuts but became less so over time. On the camp-closure outcome, beneficiaries shifted toward closure at both W2 ($\hat{\tau}_2 = -0.697$, $p < 0.05$) and W3 ($\hat{\tau}_3 = -0.699$, $p < 0.05$). By W3, similar declines appeared for *Local Hosting* ($\hat{\tau}_3 = -0.669$, $p < 0.10$) and *Nation Hosting* ($\hat{\tau}_3 = -0.598$, $p < 0.10$). These shifts nearly eliminated the pre-cut gap in opposition to closure between beneficiaries and non-beneficiaries. The backlash documented under H1 was therefore concentrated among households that personally lost from the cuts, while attitudes among non-beneficiaries remained stable. Two additional treatment specifications in the Appendix support the same conclusion.

Table 3: Within-Kakuma hosting attitudes over time and by household loss

	Oppose Camp Closure	Local Hosting	Nation Hosting
Panel A: Over-time change (event study)			
Wave 2 vs Wave 1	-0.260†	-0.010	-0.166
	(0.151)	(0.128)	(0.136)
Wave 3 vs Wave 2	-0.039	-0.059	0.441**
	(0.147)	(0.167)	(0.135)
Respondent-wave obs.	612	613	614
Respondents	204	205	205
Panel B: Intent-to-treat, baseline beneficiary status			
Wave 2	0.179	-0.051	-0.103
	(0.251)	(0.219)	(0.237)
Wave 3	0.141	0.332	0.645*
	(0.266)	(0.288)	(0.270)
Baseline beneficiary (W1)	0.953***	0.926***	0.735**
	(0.258)	(0.247)	(0.252)
Baseline beneficiary × Wave 2	-0.697*	0.067	-0.102
	(0.312)	(0.269)	(0.289)
Baseline beneficiary × Wave 3	-0.699*	-0.669†	-0.598†
	(0.336)	(0.352)	(0.334)
Respondent-wave obs.	613	613	614
Respondents	205	205	205

Notes: Panel A: OLS with respondent fixed effects, standard errors clustered by respondent; event-study coefficients (W2 vs. W1; W3 vs. W2). Panel B: pooled OLS with standard errors clustered by respondent, so that the time-invariant baseline-beneficiary difference is identified; ITT by baseline beneficiary status ($N = 127$ beneficiary, $N = 78$ non-beneficiary). Higher values indicate more pro-refugee attitudes. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

Table 4: Vignette experiment by empathy

	Nation Hosting	
	AMCE at mean empathy	× Empathy (interaction)
Aid: 40% (vs 100%)	-0.137† (0.082)	0.230** (0.080)
Aid: 0% (vs 100%)	-0.180* (0.080)	0.131 (0.080)
Hiring: Do not hire locals (vs hire)	-0.211** (0.073)	0.046 (0.074)
Access: Refugees only (vs shared)	-0.187* (0.076)	-0.136† (0.077)
Conflict: Peace restored (vs ongoing)	-0.623*** (0.080)	0.030 (0.084)

Notes: Linear model with respondent fixed effects; standard errors (in parentheses) clustered by respondent. Empathy is the two-item index, standardized (mean 0, SD 1), so the AMCE column gives each contrast at mean empathy and the interaction column gives the change in that contrast per one-standard-deviation increase in empathy. A positive interaction indicates that more empathetic hosts withdraw support less as conditions worsen. $N = 1{,}883$ profiles from 299 respondents. † $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

5.3 H3: Empathy

H3 predicts that when hosts have high empathy, they are less likely to backlash against hosting in response to aid cuts, and may even become more supportive. The evidence is consistent with H3, though it is suggestive and exploratory: empathy dampens the response to aid cuts, and the open-ended responses show that humanitarian concern for refugees is widespread.

Empathy heterogeneity in the vignette experiment. Table 4 provides some support for H3. The interaction between empathy and the 40 percent aid contrast is positive and significant (0.230, $p < 0.01$): at mean empathy, reducing aid to 40 percent lowers hosting support (-0.137), but a one-standard-deviation increase in empathy offsets essentially all of this decline. The interaction for the 0 percent contrast runs in the same direction but is smaller and weaker (0.131, $p \approx 0.10$), so the dampening effect of empathy is clearest for a partial cut. The descriptive pattern mirrors this: among low-empathy respondents, support for hosting falls from 40.9 percent under full aid to 32.5 percent at 40 percent and 33.2 percent at 0 percent, whereas high-empathy respondents do not withdraw support at all (40.3 percent under full aid, 40.3 percent at 40 percent, and 39.8 percent at 0 percent). This does not extend to the other aid attributes: empathy does not condition the response to local hiring, and for *Services* the interaction is negative and marginally significant, though this reflects the higher starting point from which high-empathy respondents support hosting when services are shared with locals (see appendix).

Levels of humanitarian concern. Both the experimental and qualitative evidence suggest high levels of humanitarian concern for refugees. Figure 4 shows that ongoing origin-country conflict raises hosting support by 0.63 points on the five-point scale relative to restored peace ($p < 0.001$). This is the largest effect in the experiment, exceeding each of the aid contrasts. The result suggests that respondents place weight on refugees' need for protection, independent of the aid-related benefits they receive.

The open-ended questions provide complementary descriptive evidence that humanitarian concern is high on average. At the nationwide level, when respondents were asked why they support or oppose hosting, over 40 percent cited humanitarian reasons, more than any other category, including economic considerations (Figure 5). In Kakuma, when respondents are asked what changes they have seen from the cuts, refugee suffering is the most commonly mentioned category, cited by 34.1 percent of respondent-waves, against 23.3 percent for crime, 23.1 percent for food, and just 2.4 percent for impacts on the host community itself (Figure 5). Hosts describe the cuts primarily by reference to what is happening to refugees, not themselves.

However, responses to a third open-ended question – asking Kakuma residents in Wave 3 why they supported or opposed closing the camps – suggest that humanitarian concern can be outweighed by economic and security considerations. Among respondents who oppose closure, the most common rationale is the camps' economic benefits, cited in 63.8 percent of coded responses, compared to 22.7 percent instead discussing humanitarian concern. Among the minority who support closure, the modal rationale is not economic cost but insecurity, cited in 58.1 percent of responses. This helps explain why aggregate backlash may emerge even when some members of the exposed community express high empathy toward refugees.

6 Robustness

I conduct a series of robustness checks, detailed in the appendix. For the nationwide difference-in-differences, I confirm the samples are nationally representative and similar to each other. Because only two counties are treated, conventional clustered standard errors are unreliable, so I use a wild cluster bootstrap (Cameron, Gelbach and Miller, 2008). The results hold under entropy-balancing weights, alternative index constructions, binary recodings of the outcomes, ordered logit in place of OLS, including enumerator fixed effects, and county leave-one-out analysis. The appendix reports further outcomes, including an Afrobarometer refugee-numbers item fielded in 2024 and 2026, and cross-sectional differences in 2026 attitudes towards refugee integration. One concern is that changes over the three-year window between surveys could confound the estimates, most notably the Shirika Plan, a refugee policy announced in 2021. I show that the Shirika Plan is unlikely to account for the shift due to low awareness; nine in ten respondents in 2026 were unaware of the Plan.

For the panel, I assess attrition, showing that no Wave 1 observable predicts retention and that inverse-probability-of-retention weights leave the estimates essentially unchanged. I show that the results replicate on the full repeated cross-section of at least 300 respondents per wave, which includes refresher respondents. Two additional personal-loss specifications, a cohort difference-in-differences with a Wave 2 placebo and a switching design pooling loss events by domain, corroborate the intent-to-treat result. The appendix reports the tests behind the description of household changes in Figure 2, examines differences by baseline-beneficiary status, and confirms the intent-to-treat estimates hold under enumerator fixed effects. I examine longer-run changes with a phone survey fielded alongside the February 2026 national survey. For the vignette experiment, I show that treatment assignment is random, report marginal means by empathy group alongside the interaction coefficients (Leeper, Hobolt and Tilley, 2020), and present a descriptive question on support under a full aid cut. 97.7% of respondents pass the comprehension check, and the results hold excluding those who fail.

Given the small sample, I report minimum detectable effects for the main tests. Social desirability is unlikely to drive the results, since the quantities of interest are over-time changes rather than levels, so desirability pressures present at every wave difference out. Because both the hosting counties and the baseline-beneficiary households begin

Figure 5: Open-ended responses on camp closure, the cuts (Kakuma), and reasons for hosting refugees (national)

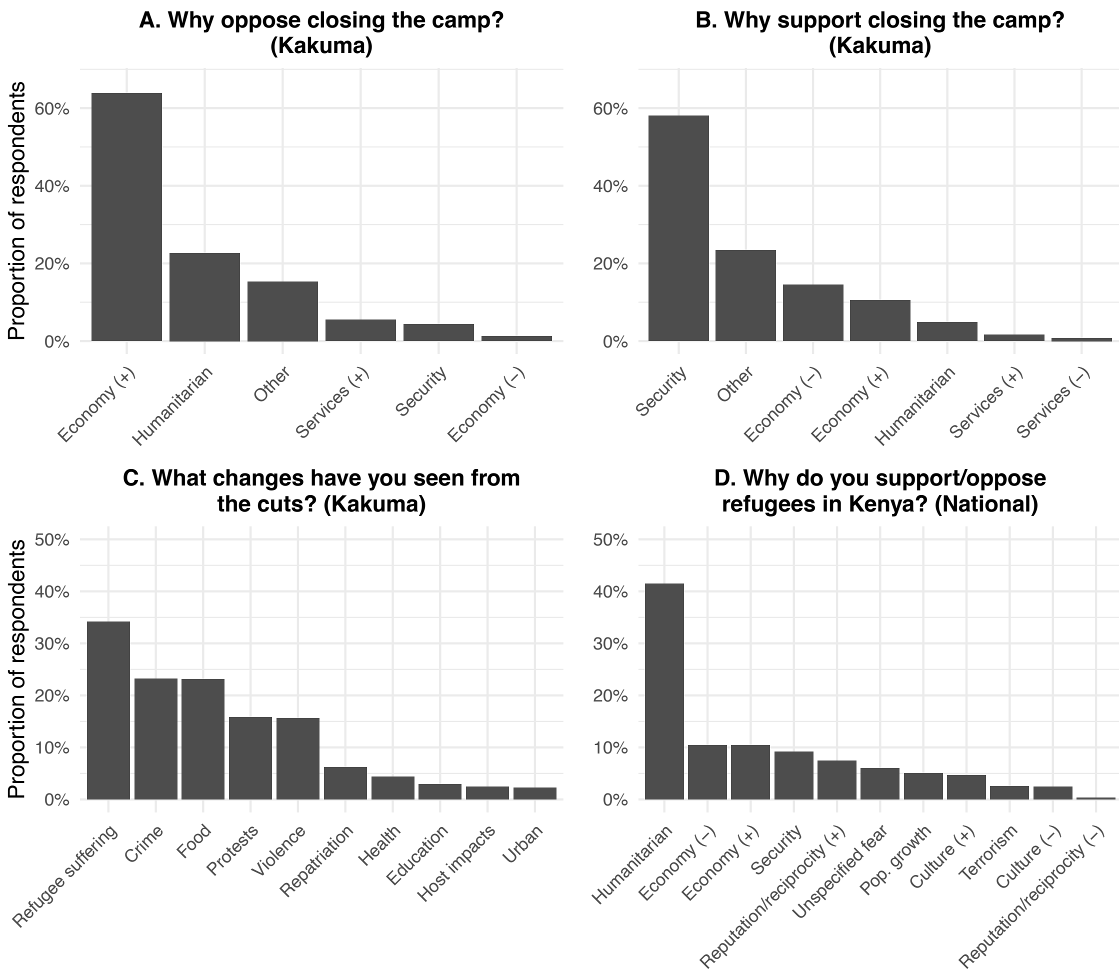

Notes: Share of respondents mentioning each theme (multiple mentions allowed). Panels A–B: W3 Kakuma respondents supporting ($N = 124$) and opposing ($N = 163$) camp closure. Panel C: pooled W1–W3 Kakuma ($N = 615$ respondent-waves). Panel D: 2023 national sample, control arm. See appendix.

with more favorable attitudes, I also address the concern that the estimated declines reflect reversions to the mean.

Finally, I show that despite the decline in local hosting support in Kakuma, support for Kenya hosting refugees more broadly does not fall: Table 3 shows that nationwide hosting support holds steady and rises with the minor improvements from stabilization, and the same pattern appears for trust in refugees and multiple refugee-integration measures (see appendix). The backlash that aid cuts produce is therefore specific to local hosting rather than to the principle of hosting refugees in Kenya more generally, consistent with the open-ended evidence that respondents retain support for refugees' broader protection.

7 Conclusion

In 2025, global humanitarian aid contracted at the fastest rate on record. This paper is concerned with the domestic political consequences of such aid cuts on countries providing asylum to refugees, specifically whether host citizens exposed to refugee aid cuts withdraw support for hosting in their area. Drawing on two nationwide surveys, a high-frequency panel in a refugee-hosting region, a vignette experiment, and refugee-leader interviews in Kenya, I find localized but not nationwide backlash. Support fell in refugee-hosting regions while holding steady elsewhere, producing convergence toward the national mean. Within the hosting region, the decline was concentrated among households that personally lost aid-related benefits. Empathy seems to dampen backlash: the main reason that citizens support refugee hosting nationwide is because of humanitarian concern, and high-empathy hosts in the hosting region do not withdraw support even under a hypothetical no-aid scenario.

This study revises how we understand the relationship between exposure and attitudes toward immigration, particularly when a policy shock harms hosts and migrants together. Existing research often treats exposure either as a source of economic threat or benefit, or as a source of contact. This paper brings these channels into the same framework. In aid-dependent hosting areas, exposure to aid cuts makes citizens vulnerable to the loss of jobs, services, and other benefits tied to the refugee-aid economy, but it also increases the likelihood that citizens witness and empathize with refugees' suffering. A single shock that harms hosts and refugees at once can therefore move attitudes in opposing directions. Personal, egocentric losses push some hosts toward opposition, while humanitarian concern pulls in the opposite direction. This framework can be applied to any policy shock that jointly harms a host population and a marginalized group with whom it lives in close contact.

Although the empirical evidence comes from a single context, the theory applies more broadly and specifies where the dynamics documented here should recur. Backlash should be strongest where refugees are tied to a visible local aid economy, as in camp settings where assistance is geographically concentrated, locals receive benefits, and hosts and refugees interact. Such settings are common: one in five refugees worldwide lives in a camp, most often in sub-Saharan Africa, where 74% of refugees are encamped, but also in Bangladesh, Jordan, Thailand, and elsewhere (Anti et al., 2025; UNHCR, 2026*b*). The dynamics should be weaker where refugees are dispersed, urban, or not clearly associated with aid benefits. The pattern of stasis nationwide, in turn, is most likely where empathy toward refugees is greater than hostility. Where anti-refugee rhetoric dominates, aid cuts may produce broader backlash even among citizens who are not personally harmed.

A question that follows from these findings concerns the broader implications of localized backlash. Can such backlash affect national refugee-hosting policy, especially when attitudes elsewhere in the country remain supportive? I suggest that it can, either when local politicians bring host-community concerns into national asylum debates or when insecurity in hosting areas justifies camp closures, relocation, or limits on future hosting. In the 1990s, for example, the Kenyan government closed refugee camps along the coast after violence between refugees and local residents killed eight people and displaced more than 2,000 others, requiring refugees to move to other camps or return to Somalia (MacDonald, 2026). Whether and how local concerns aggregate into national policy is a question that warrants further research as aid cuts continue.

The central takeaway of this study is that aid cuts have political consequences beyond their direct welfare effects. They do not only reduce food, healthcare, education, and employment for refugees; they can also weaken the local support that makes refugee hosting politically sustainable in low- and middle-income countries, where most refugees live and where governments often rely on donor financing to maintain asylum. The backlash documented here remained modest and local, in part because some aid continued and material benefits remained, humanitarian concern was widespread, and the cuts did not become a highly politicized national issue. These conditions may not persist under sustained aid reductions. Continued cuts may further reduce the benefits hosts receive and their empathy for refugees, raising pressure for them to return to unsafe origin countries or move onward. Whether hostility spreads to citizens outside hosting regions, as the secondary economic and security consequences of the cuts reach them, is a question for future research. Understanding how host populations respond as the aid economy contracts will be increasingly important for explaining the future politics of refugee hosting.

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.

Alix-Garcia, Jennifer, Sarah Walker, Anne Bartlett, Harun Onder and Apurva Sanghi. 2018. "Do refugee camps help or hurt hosts? The case of Kakuma, Kenya." *Journal of Development Economics* 130:66–83.

Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103(484):1481–1495.

Anti, Sebastian, Colette Salemi, Jonathan Rigberg, Karishma Silva and Johannes Hoogeveen. 2025. "A New Expanded Dataset to Study Refugee Camps in Sub-Saharan Africa 1999–2024." Working paper.
URL: *https://colettesalemi.ca/files/Anti_etal_2025.pdf*

Anyadike, Obi. 2024. "Dadaab Voices: Despair as Refugee Food Rations in Kenya Slashed by 60%." *The New Humanitarian* .
URL: *https://www.thenewhumanitarian.org/news-feature/2024/06/13/dadaab-voices-despair-refugee-food-rations-kenya-cut-60*

Aviña, Marco M., Taeku Lee, Mashail Malik, Reed Rasband, Marcel F. Roman and Priyanka Sethy. 2026. "Which Immigrants Do Citizens Prefer? A Meta-Reanalysis of 100 Conjoint Experiments." *Science Advances* 12(22).

Bansak, Kirk, Jens Hainmueller and Dominik Hangartner. 2016. "How economic, humanitarian, and religious concerns shape European attitudes toward asylum seekers." *Science* 354(6309):217–222.

Bansak, Kirk, Jens Hainmueller and Dominik Hangartner. 2023. "Europeans' support for refugees of varying background is stable over time." *Nature* 620(7975):849–854.

Baseler, Travis, Thomas Ginn, Robert Hakiza, Helidah Ogude-Chambert and Olivia Woldemikael. 2025. "Can redistribution change policy views? Aid and attitudes toward refugees." *Journal of Political Economy* 133(9):2874–2917.

Bearce, David H. and Daniel C. Tirone. 2010. "Foreign Aid Effectiveness and the Strategic Goals of Donor Governments." *The Journal of Politics* 72(3):837–851.

Bermeo, Sarah Blodgett. 2017. "Aid Allocation and Targeted Development in an Increasingly Connected World." *International Organization* 71(4):735–766.

Betts, Alexander, Gil Loescher and James Milner. 2012. *UNHCR: The Politics and Practice of Refugee Protection*. Second ed. London: Routledge.

Cameron, A. Colin, Jonah B. Gelbach and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *The Review of Economics and Statistics* 90(3):414–427.

Demir, Cansu, Anna Maria Mayda and Jean-François Maystadt. 2025. "The Political Impact of Refugees in Africa." Working paper.
URL: *https://www.jointdatacenter.org/wp-content/uploads/2026/06/Jean-Francois-Maystadt.pdf*

Dempster, Helen and Thomas Ginn. 2026. The World Bank's Window for Host Communities and Refugees: Successes, Challenges, and the Way Forward. Policy Paper 390 Center for Global Development Washington, DC: .
URL: *https://www.cgdev.org/publication/world-banks-window-host-communities-and-r*

Ghosn, Faten, Alex Braithwaite and Tiffany S Chu. 2019. "Violence, displacement, contact, and attitudes toward hosting refugees." *Journal of Peace Research* 56(1):118–133.

Grossman, Guy, Yang-Yang Zhou and Shelby Carvalho. 2025. "Disrupted Aid, Displaced Lives: Unraveling the Impact of Refugee Funding Cuts through a Regression Discontinuity Study in Uganda." Working paper.
URL: *https://pdri-devlab.upenn.edu/project/disrupted-aid-displaced-lives-unraveling-the-impact-of-refugee-funding-cuts-through*

Hainmueller, Jens. 2012. "Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies." *Political Analysis* 20(1):25–46.

Hainmueller, Jens and Daniel J Hopkins. 2014. "Public attitudes toward immigration." *Annual Review of Political Science* 17(1):225–249.

Hartman, Alexandra C and Benjamin S Morse. 2020. "Violence, empathy and altruism: Evidence from the Ivorian refugee crisis in Liberia." *British Journal of Political Science* 50(2):731–755.

Hartman, Alexandra C, Benjamin S Morse and Sigrid Weber. 2021. "Violence, displacement, and support for internally displaced persons: Evidence from Syria." *Journal of Conflict Resolution* 65(10):1791–1819.

Hovil, Lucy and Nicholas Maple. 2022. "Local Integration: A Durable Solution in Need of Restoration?" *Refugee Survey Quarterly* 41(2):238–266.

International Finance Corporation. 2018. Kakuma as a Marketplace: A Consumer and Market Study of a Refugee Camp and Town in Northwest Kenya. Technical report International Finance Corporation Washington, DC: . License: All Rights Reserved.
URL: *https://openknowledge.worldbank.org/handle/10986/34645*

Kadigo, Mark Marvin and Jean-Francois Maystadt. 2023. "How to cope with a refugee population? Evidence from Uganda." *World Development* 169:106293.

Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2):263–291.

Kenny, Charles and Justin Sandefur. 2025. Update on Lives Lost from USAID Cuts. CGD blog post Center for Global Development.
URL: *https://www.cgdev.org/blog/update-lives-lost-usaid-cuts*

Leeper, Thomas J., Sara B. Hobolt and James Tilley. 2020. "Measuring Subgroup Preferences in Conjoint Experiments." *Political Analysis* 28(2):207–221.

MacDonald, Mae. 2026. "Push, Not Peace: Reconsidering the Drivers of Refugee Return." Working paper.
URL: *https://www.maemacdonald.com/papers/MMacDonald_Return.pdf*

MacDonald, Mae and Adam Lichtenheld. 2025. "Support for Refugee Integration in a Major Refugee-Hosting Country: Evidence from Kenya." Working paper.
URL: *https://osf.io/preprints/socarxiv/92qy5_v3*

Newman, Benjamin J., Todd K. Hartman, Patrick L. Lown and Stanley Feldman. 2015. "Easing the Heavy Hand: Humanitarian Concern, Empathy, and Opinion on Immigration." *British Journal of Political Science* 45(3):583–607.

OCHA. 2025. "Financial Tracking Service (FTS).". Accessed: 20 March 2025.
URL: *https://fts.unocha.org*

OECD. 2026. "A Historic Decline in Foreign Aid: Preliminary 2025 ODA Data." https://www.oecd.org/en/data/insights/data-explainers/2026/04/a-historic-decline-in-foreign-aid-preliminary-2025-oda-data.html. Accessed 23 May 2026.

Pettigrew, Thomas F. and Linda R. Tropp. 2006. "A Meta-Analytic Test of Intergroup Contact Theory." *Journal of Personality and Social Psychology* 90(5):751–783.

Sanghi, Apurva, Harun Onder and Varalakshmi Vemuru. 2016. "Yes" in My Backyard? The Economics of Refugees and Their Social Dynamics in Kakuma, Kenya. Technical report World Bank Washington, DC: .
URL: *https://documents1.worldbank.org/curated/en/308011482417763778/pdf/111303-WP-Kakuma-Report-Yes-in-my-backyard-December-2016-PU pdf*

Sterck, Olivier and Vittorio Bruni. 2025. "The Welfare and General-Equilibrium Impacts of Aid Cuts." Working Paper.
URL: *https://ssrn.com/abstract=5270846*

UNHCR. 2024. "Kenya Funding Update." https://reporting.unhcr.org/kenya-funding-update. Accessed: 2025-03-21.

UNHCR. 2025. "Kenya Operational Data Portal.".
URL: *https://data.unhcr.org/en/country/KEN/796*

UNHCR. 2026*a*. Global Report 2025. Technical report United Nations High Commissioner for Refugees Geneva: . Accessed July 2026.
URL: *https://www.unhcr.org/publications/global-report-2025*

UNHCR. 2026*b*. "Operational Data Portal: Refugee Situations." https://data.unhcr.org. Accessed July 2026.

U.S. Department of State. 2025. "Foreign Assistance Data.". Accessed: 20 March 2025.
URL: *https://www.foreignassistance.gov*

Weber, Sigrid, Nik Stoop, Peter Van der Windt and Haoyu Zhai. 2025. "A meta-analysis of attitudes towards migrants and displaced persons." *British Journal of Political Science* 55:e168.

WFP. 2025*a*. "Food Security Impact of Reduction in WFP Funding." World Food Programme.
URL: *https://www.wfp.org/publications/food-security-impact-reduction-wfp-funding*

WFP. 2025*b*. "Kenya Country Portfolio Needs (2023–2027)." World Food Programme, Operation KE02 Resourcing Update. Accessed March 20, 2025.
URL: *https://one.wfp.org/operations/current_operations/ResUpdates/KE02.pdf*

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 Bank and UNHCR. 2021. Understanding the Socioeconomic Conditions of Refugees in Kenya: Volume B: Kakuma Camp. Results from the 2019 Kakuma Socioeconomic Survey. Technical report World Bank and UNHCR. Accessed November 14, 2024.
URL: *https://documents1.worldbank.org/curated/en/443431613628051180/pdf/Understanding-the-Socio-Economic-Conditions-of-Refugees-in-Kenya-Volume-B-Kakuma-Camp-Results-from-the-2019-Kakuma-Socioeconomic-Survey.pdf*

Zhou, Yang-Yang, Guy Grossman and Shuning Ge. 2023. "Inclusive refugee-hosting can improve local development and prevent public backlash." *World Development* 166:106203.

Zhou, Yang-Yang, Naijia Liu, Shuning Ge and Guy Grossman. 2025. "Liberalizing Refugee Hosting Policies without Losing the Vote." OSF Preprints, Version 4.
URL: *https://doi.org/10.31219/osf.io/94tpy*

Aid Cuts and the Politics of Refugee Hosting

Supplementary Materials

Contents

1 Study design and measurement

1.1 Enumeration and recruitment

1.1.1 Nationwide

I worked with the Nairobi-based survey firm, TIFA Research, to conduct this study. TIFA maintains a database of previously profiled citizens, from which it recruited for this project. Respondents were drawn by stratified random sampling on gender, age, and location so that the sample matched the demographic distribution of the 2019 Kenya census. The sample is distributed across counties in proportion to the adult population using probability proportional to size, with the 2019 Kenya Population and Housing Census as the reference. Citizens were eligible if they were Kenyan voters aged 18 to 65.

Prior to data collection, I conducted extensive in-person piloting and training with enumerators. TIFA implemented rigorous quality control checks, including a quality control team listening to every interview to check for enumerator miscoding and to ensure that the qualitative open-ended responses had been captured correctly. If an interview was incorrectly administered, it was removed and not counted in the total survey number. For the qualitative open-ends, enumerators asked these questions in Swahili or English (or Somali in the 2023 survey, because of the Garissa oversample) depending on the preferences of the respondent, and then typed the response into TIFA's survey software in that language. TIFA's quality control team checked the correctness and quality of these responses against their recordings before employing translators to translate the responses into English for analysis. Ethics approvals, consent procedures, incentives, and data protections are described in Section 1.2.

In the 2023 baseline, the *Nation Hosting* and *Local Hosting* measures were embedded in a survey experiment fielded for a separate study, in which respondents were randomized across arms that varied refugee nationality. In the current study, I use only the generic-refugees control arm for these outcomes, which reduces the 2023 sample from the full nationally representative sample of 2510 to the 816 respondents in the control arm. A technical error in the February 2026 survey affected the *Nation Hosting* measure only. The fault for this error is mine, not TIFA's. Because of this, I fielded a recontact survey in May 2026, which asked respondents the *Nation Hosting* measure alongside questions for a separate study. Of the 2,042 February respondents, 1,776 (87.0 percent) completed this survey. Section 1.3.1 compares these different samples. Enumerator identifiers are available for the February and May 2026 surveys but not the 2023 baseline. Section 5.1 shows the 2026 estimates are robust to enumerator fixed effects.

1.1.2 Panel

The face-to-face panel was administered by TIFA enumerators in Kakuma across three waves in 2025, with a fourth phone wave in February 2026 as part of the nationally representative survey, described in Section 3.4. Respondents were recruited through a multi-stage design. Sub-locations, the smallest administrative units, were treated as clusters and allocated sample using probability proportional to the number of households; within each cluster, households were selected at random; and within each household, a single respondent was drawn from the eligible members at random, with a further random selection where the first-drawn member was unavailable. Kenyan voters aged 18 to 65 resident in the study area were eligible. Where a panel respondent could not be re-interviewed, refresher respondents were added to maintain the cross-sectional sample. Enumerators were recruited from the Kakuma host community, which helped establish trust with participants, and interviewed the same respondents at each wave where possible. Section 5.1 verifies that the panel results are robust to enumerator fixed effects. Consent procedures, incentives, and data protections are described in Section 1.2.

1.1.3 Refugee-leader interviews

Alongside each face-to-face survey wave, I conducted semi-structured interviews with a panel of refugee leaders in Kakuma. Leaders were purposively sampled to vary by nationality, gender, and leadership role, including healthcare workers, local news broadcasters, leaders of refugee-led organizations, and employees of international organizations, and were recruited with the help of a local journalist and research assistant with strong networks in Kakuma. Table 1 summarizes the composition of the interview panel, showing that it is broadly representative of the gender and nationality of the Kakuma/Kalobeyei refugee population.

Because UNHCR restrictions prevented in-person research in the camp during the cuts, interviews were conducted remotely over WhatsApp and lasted 30 to 60 minutes. Interviews were conducted in English. We first discussed the leader's role in the camp. The topic guide then covered three areas. First, the cuts: which organizations and services were affected, the main impacts on refugees, coping strategies, and attribution of responsibility for the cuts. Second, integration and the Shirika Plan. Third, relations with the host community. The guide was held constant across waves so that changes in leaders' accounts track changes in conditions, and each follow-up opened by asking what had changed since the previous interview. Consent procedures, compensation, and data protections are described in Section 1.2.

I reviewed the full set of interviews at each wave and selected quotations that illustrate themes recurring across multiple leaders, prioritizing accounts corroborated by other interviews or by the survey data. The interviews were not formally coded. Quotations are reproduced verbatim except for light edits for readability.

Table 1: Composition of refugee leader interview sample

Characteristic	Category	N	Sample (%)	Population (%)
Gender	Female	9	50.0	46.7
Gender	Male	9	50.0	53.3
Nationality	South Sudanese	10	55.6	61.7
Nationality	Somali	4	22.2	12.9
Nationality	Congolese	2	11.1	8.2
Nationality	Burundian	2	11.1	7.7

Note: Population percentages refer to the Kakuma/Kalobeyei refugee population (UNHCR operational data portal, March/April 2026).

1.2 Ethics

This research was conducted in accordance with APSA's Principles and Guidance for Human Subjects Research. The study was approved by my university's Institutional Review Board, and received Kenyan ethics approval from the Amref Health Africa Ethics and Scientific Review Committee (ESRC) and a research license from the National Commission for Science, Technology and Innovation (NACOSTI).

Phone-survey respondents gave verbal informed consent at the start of each call under a waiver of documentation, were told they could decline any question, and were asked to take the survey in a private place. Face-to-face panel respondents gave written consent before each interview, in English, Swahili, or Turkana according to their preference, and consent was re-obtained at each follow-up with a modified script stating that responses would be linked across waves. Refugee-leader interviewees gave verbal consent, could decline any question or end the interview at any time, and consented to being audio-recorded. Survey participants in the phone and face-to-face surveys received small phone airtime incentives, set by TIFA Research based on standard payments in the sector; refugee leaders received a payment of 1,000 KSh (approximately 8 USD) at each survey wave in mobile money, an amount advised by the research assistant on the basis of local norms. All survey data were anonymized by TIFA before being transferred to me. Refugee leaders are not identified in the paper, and their quotations were reviewed to remove potentially identifying details.

The nationwide surveys are representative of the Kenyan adult population and not composed mainly of vulnerable or marginalized groups. The Kakuma panel and refugee-leader interviews were conducted with individuals from relatively vulnerable populations, especially at the time of the cuts. Given this, the study was designed so that participation posed no more than minimal risk and questions about sensitive experiences could be declined. In particular, the interview questions with refugee leaders were designed to discuss the community as a whole, rather than their individual experiences. The research did not involve deception.

1.3 Sample

1.3.1 Nationwide surveys

Both nationwide surveys were conducted by phone, which is likely to skew toward respondents who are younger and more educated than the adult population. I therefore construct post-stratification weights for each survey using entropy balancing (Hainmueller 2012), which reweights each sample so that its distribution of age, gender, education, and region matches a target population distribution. Each survey is weighted to the Afrobarometer Kenya face-to-face survey conducted prior to that survey wave: the 2023 baseline to Round 9 (fielded 2021) and the 2026 follow-up to Round 10 (fielded 2024). I use Afrobarometer data rather than the census because the census does not include a measure of education. Weights are constructed separately for the 2023 and 2026 surveys, with the 2026 weights also applied to the May follow-up. All descriptive statistics and maps in the paper are weighted, and the difference-in-differences estimates are unweighted with HC1 standard errors.

I make three sample comparisons, reported in Table 2 descriptively and Tables 3 and 4 as formal tests. First, because the 2023 hosting items were asked only of the control arm, I compare that arm against the other arms to confirm the smaller control sample is not demographically distinct. The 'Full nat-rep' and 'Control arm' columns of Table 2 show the two are similar, and Panel A of Table 3 finds no covariate differs significantly between them.

Second, I compare the 2023 and 2026 nationally representative samples. The '2023 Full nat-rep' and '2026 Full nat-rep' columns of Table 2 show the two are similar, and Panel B confirms they match on gender and education; the one difference is age, with 2026 respondents on average 1.2 years older ($p < .01$). I account for this by including age, along with gender and education, as a control in all cross-survey specifications, and by weighting all descriptive statistics and maps. Third, I compare the May recontacted respondents against the full February sample. Of the 2,042 February respondents, 87.0 percent responded in May. The 'Full nat-rep' and 'Follow-up' columns of Table 2 show the recontact subsample is nearly identical to the February sample on every demographic. Table 4 regresses response to the May survey on demographics from the February survey wave. The only predictor of response is age, and its effect is small, with each additional year raising the response probability by 0.2 percentage points. As above, to account for this, I include age as a control in every specification and the weighting strategy. Importantly, there is no evidence that those who were more or less supportive of refugees were more likely to attrit from the study.

Table 2: Nationwide survey samples and external benchmark

	2023		2026		Afrobarometer
	Full nat-rep	Control arm	Full nat-rep	Follow-up	R10 (2024)
N	2510	816	2042	1776	2400
Age	36.92 (11.84)	36.67 (11.67)	36.07 (12.30)	36.34 (12.43)	36.88 (14.99)
Female	50.0%	47.2%	49.9%	50.3%	49.9%
Education: secondary or higher	48.3%	47.1%	58.3%	59.3%	55.8%
Region					
Central	10.3%	9.7%	11.7%	12.3%	12.6%
Coast	10.1%	10.1%	10.7%	10.1%	9.0%
Eastern	13.3%	13.3%	15.2%	14.0%	14.5%
Nairobi	9.8%	8.1%	9.0%	9.6%	10.2%
North Eastern	6.2%	7.4%	4.7%	4.9%	4.6%
Nyanza	12.2%	11.8%	12.4%	12.1%	12.7%
Rift Valley	25.5%	26.3%	26.4%	26.8%	26.0%
Western	12.6%	13.3%	9.9%	10.3%	10.4%

Note: Continuous variables: weighted mean (SD). Other rows: weighted percentages. Region percentages are weighted shares of each province and sum to 100 within each column.

Table 3: Tests of nationwide sample differences

Covariate	Diff	SE
Panel A: 2023 framing arms (other vs. refugees arm); joint F = 0.81, p = 0.488		
Age	-0.285	(0.492)
Female	-0.026	(0.021)
Education: secondary or higher	-0.010	(0.019)
Panel B: 2023 vs. 2026 nationally representative; joint F = 4.10, p = 0.006		
Age	1.164**	(0.361)
Female	0.007	(0.015)
Education: secondary or higher	0.001	(0.013)

Note: Each row reports the difference in means between the two samples with an HC1 standard error in parentheses. The joint test is an HC1-robust Wald test that all covariates are unrelated to sample membership. Tests are unweighted. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 4: Correlates of May 2026 follow-up survey retention

Predictor	Coef.	SE
Age	0.002***	(0.001)
Female	-0.000	(0.015)
Education: secondary or higher	0.029	(0.019)
Approval of national govt (Feb.)	-0.007	(0.005)
Integration index (Feb.)	0.013	(0.011)

Note: Linear probability model of responding to the May 2026 recontact on February 2026 observables, HC1 standard errors. Joint Wald test that all predictors are zero: F = 3.19, p = 0.007. N = 2,000. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

1.3.2 Kakuma

Table 5: Kakuma sample: balanced panel vs. repeated cross-section

		Repeated cross-section		
	Panel	W1	W2	W3
N	205	300	301	301
Age	30.17 (9.14)	30.04 (8.88)	29.85 (8.92)	29.72 (8.84)
Female	48.8%	47.0%	47.2%	47.2%
Somali	0.5%	0.3%	0.0%	0.7%
Turkana (ethnicity)	95.1%	94.0%	95.0%	94.3%
Education: None	14.6%	14.3%	13.0%	14.6%
Education: Primary	17.6%	16.7%	18.6%	18.3%
Education: Secondary	38.5%	38.0%	40.5%	36.9%
Education: Tertiary	28.8%	30.7%	27.6%	25.9%
Monthly income (KES)	12,399 (26,120)	14,191 (21,299)	10,453 (20,681)	13,943 (35,079)
Food insecure	84.5%	86.9%	81.3%	79.7%
Approval of national govt (1–5)	2.98 (1.73)	2.84 (1.71)	3.00 (1.78)	3.13 (1.74)

Note: Continuous variables: mean (SD). Other rows: %. The Panel column is the balanced three-wave panel (respondents observed in all three waves), with time-invariant traits measured at Wave 1 and time-varying measures pooled across waves; W1–W3 are the full face-to-face waves, including refreshers added in each wave.

Table 6: Attrition selectivity on Wave 1 observables

	(1) LPM	(2) Logit
Constant	0.352**	-0.575
	(0.124)	(0.704)
Age	-0.003	-0.015
	(0.003)	(0.017)
Female	-0.038	-0.192
	(0.054)	(0.274)
Secondary or higher	0.013	0.071
	(0.060)	(0.313)
W1 hosting support	0.015	0.075
	(0.015)	(0.077)
Baseline beneficiary	-0.046	-0.228
	(0.055)	(0.270)
N	300	300
Joint test	$F = 0.76$	$\chi^2 = 3.44$
Joint p	0.576	0.633

Note: The dependent variable indicates attriting after Wave 1. Column (1) is a linear probability model with HC1 robust SEs; column (2) a logit reporting raw coefficients with conventional SEs (in parentheses). The LPM joint test is a robust Wald F; the logit joint test a likelihood-ratio χ^2. Significance: *** $p < .001$, ** $p < .01$, * $p < .05$, † $p < .10$.

The Kakuma panel was conducted face-to-face, with the intention of surveying the same 300 host-community respondents at roughly two-month intervals. The Wave 1 cohort comprised 300 tracked respondents, of whom 217 (72%) completed the survey at Wave 2 and 205 at Wave 3; 83 attrited after Wave 1 and a further 12 between Waves 2 and 3. Waves 2 and 3 comprised 301 and 301 interviews respectively, the balance made up of refresher respondents recruited to maintain the repeated cross-section.

Attrition was concentrated between Waves 1 and 2, which coincided with a period of protests and heightened insecurity around Kakuma. Enumerators struggled to reach some households, and some respondents declined to be interviewed because of the unrest, producing a higher attrition rate than anticipated. Table 6 shows the results of regressing an indicator for attriting after Wave 1 on Wave 1 age, gender, education, hosting support, and baseline beneficiary status, the H2 treatment. None of the five predictors is individually significant, and the joint tests cannot reject that all five are simultaneously zero, so those who left the panel do not differ observably from those who stayed. However, note that these tests may be underpowered to detect changes.

Because no Wave 1 observable predicts three-wave retention (Table 6), inverse-probability-of-retention weighting should leave the intent-to-treat estimates essentially unchanged, and it does (Table 25). I estimate each retained respondent's probability of appearing in all three waves from a logit on age, gender, education, hosting support, and beneficiary status, then weight the fixed-effects intent-to-treat regressions by the normalized inverse of the fitted probabilities. Adjusting for differences between respondents who remained in the panel and those who left does not change the main result. Note that this analysis corrects for selection on observables only. Attrition correlated with unobserved trajectories is addressed instead by the repeated cross-section replication (Section 3.3), which finds the same pattern on a sample refreshed at every wave.

1.4 Survey instruments

This section reports the English wording of every survey item used in the analysis. Response scales are named once here and referenced by shorthand after each item. Unless noted otherwise, every item carries a do-not-read "don't know or refused" option, and in cleaning I reorient every outcome so that higher values denote more pro-refugee or pro-host responses.

- **[SO]** — five-point support scale: strongly support / somewhat support / neither support nor oppose / somewhat oppose / strongly oppose.
- **[AD]** — five-point agreement scale: strongly disagree (1) / somewhat disagree / neither agree nor disagree / somewhat agree / strongly agree (5).
- **[YN]** — yes / no.

1.4.1 Nationwide outcome items

Nation Hosting [SO]. "We now want to ask you more broadly about refugees. To what extent do you support or oppose Kenya hosting refugees?"

Local Hosting [SO]. 2023 Natrep: "Suppose the Kenyan government is going to enable refugees to move to different areas throughout the country. Would you support or oppose refugees moving into your constituency?" 2026 Natrep: "Please tell me how strongly you support or oppose the following policies: Refugees moving to live in your district." The item was reworded slightly between rounds so that it read naturally alongside the 2026 integration battery, into which it was folded. The shift from "constituency" to "district" does not introduce a meaningful difference: "district" is a colloquial term for what became the electoral constituency after the 2010 constitutional reforms, and the terms are used interchangeably in Kenya. TIFA Research, the survey firm, confirmed the two are treated as equivalent. Also, because the difference-in-differences compares each round to itself across counties, any level shift common to all respondents from the rewording is absorbed by the Post main effect and does not enter the treatment interaction. Panel: Introduced with "Please tell me how strongly you support or oppose the following policies," followed by the text: "Refugees continuing to live in your district". The phrasing differs slightly to reflect the hosting-area context. Because the analysis compares changes within the same panel, this wording difference does not affect the estimates.

Policy preference (source of *Not Repat.*). "Which of these five statements is closest to your own opinion towards refugees living in Kenya?" Five options read in the following order: "(1) they should live in camps and not be able to legally work or move; (2) they should live in settlements and be able to legally work and move in Turkana and Garissa; (3) they should live throughout Kenya and be able to legally work and move freely; (4) they should be resettled to other countries; (5) they should be sent back to their country of origin." *Not Repat.* is coded 1 for any response other than option (5) and 0 for option (5).

Refugee numbers (replicating Afrobarometer Round 10 Q73C verbatim). "Do you think Kenya should allow more or fewer refugees from other countries to come and live in this country?" [more refugees / about the same / fewer refugees / no refugees at all / don't know or refused]. The indicator used in the interim-benchmark analysis (Section 2.5.2) equals 1 for "more refugees" or "about the same" and 0 otherwise.

Integration items [SO]. Introduced with "Please tell me how strongly you support or oppose the following policies," followed by the text:

- Healthcare access: "Refugees having access to Kenyan doctors and hospitals and paying the same price as Kenyans."
- Right to work: "Refugees being allowed to work in your community."
- Freedom of movement: "Refugees being able to move freely and settle throughout Kenya."

Camp closure (reversed as *Oppose Camp Closure*) **[SO]**. "Please tell me how strongly you support or oppose the following policies," followed by: "The Kenyan government closing Kakuma camp and moving the refugees in Kakuma back to their countries."

Hypothetical full-cut scenario [SO]. "Currently, Kenya receives financial support from international donors to host refugees. Imagine a situation where donors completely cut off/end their financial support for refugees in Kenya. This would mean that they no longer provide any of the aid that is used for services for refugees, such as healthcare and education. In this hypothetical situation, to what extent would you support Kenya continuing to host refugees?"

Empathy items [AD]. Index construction is described in Section 4.3. Introduced with "I'm going to now say a few statements. Please indicate how much you agree with each statement when you think about refugees living in Kenya." Presented in random order:

- "You are very concerned about the refugees and their suffering."
- "You are worried that the refugees do not have enough food to eat because of the cuts to cash assistance and food rations."

Trust in refugees [AD]. "To what extent do you agree or disagree with the following statement? 'Generally speaking, most people in the refugee community are trustworthy.' "

Exposure items defining baseline-beneficiary status [YN].

- Camp healthcare use: "Have you or anyone in your household accessed hospitals, clinics, or dispensaries in Kakuma camp or Kalobeyei for any health issue in the past month, including the issue we asked about in the previous question or any other medical issue?"
- Camp education use: "Do you have any children in your household who have accessed schooling or education services in Kakuma camp or Kalobeyei in the past month?"
- NGO/aid employment: "Have you (or anyone in your household) been employed by an NGO in Kakuma camp or Kalobeyei in the past month, including casual work?"
- Employment by a refugee household: "Have you (or anyone in your household) been employed by a refugee (for things like housekeeping, cooking, childcare) in Kakuma camp or Kalobeyei in the past month?"
- Business serving refugees: "Do you or anyone in your household have a business that has sold goods or products to refugees in the past month?"

1.4.2 Open-ended questions and coding

Changes observed from the cuts (Kakuma, W1–W3). Respondents were first asked "Have you heard about recent cuts and reductions to services for refugees in Kakuma camp?" Those who answered yes were then asked, "What changes have you seen as a result of these cuts?"

Reasons for supporting or opposing camp closure (Kakuma, W3). Asked after the camp-closure question: "Why are you supportive/opposed/neutral toward the Kenyan government closing Kakuma camp and moving refugees back to their countries?" The wording was piped from the respondent's previous answer.

Reasons for supporting or opposing hosting refugees in Kenya (Nationwide, 2023). Asked after the *Nation Hosting* question: "Why do you support/oppose Kenya hosting refugees?" The wording was piped from the respondent's previous answer.

The unit of coding is the respondent-wave, and responses can receive multiple codes, so the category shares in Figure 5 of the paper sum to more than 100 percent. The three open-ended questions map to the figure as follows: the camp-closure reasons (Kakuma, W3) are shown in Panel A, for respondents who opposed closure (N = 163), and Panel B, for those who supported it (N = 124), split by the respondent's answer to the closure item; the changes observed from the cuts (Kakuma, pooled W1–W3, N = 615 respondent-waves) shown in Panel C; and the reasons for hosting refugees (2023 national sample, control arm) shown in Panel D.

Coding proceeded in two stages. First, I developed the codebook inductively: for each question, I read the first 10 percent of responses and defined a category scheme, with labels matching those in Figure 5. Second, the remaining responses were coded against the fixed codebook by a large language model (Claude Opus 4.7, Anthropic). Before coding the full set, I provided the model with the category definitions and worked examples, then piloted the codebook on a subset of responses. I then revised the definitions where the model's codes diverged from my own, iterating until the codebook produced reliable classifications. Responses of "don't know," refusals, and blank entries were excluded from the coded denominators.

1.4.3 Vignette experiment (W3)

Each respondent evaluated six scenarios drawn from the full $3 \times 2 \times 2 \times 2 = 24$-profile design, presented in random order.

Introduction. "Next you will see several possible scenarios. For each scenario, please tell us how much you support or oppose Kenya hosting refugees. Use a scale from 1 (strongly oppose) to 5 (strongly support)."

Scenario template. "In this scenario, international organizations provide [AID] to refugees. Health care and schools are [ACCESS]. International organizations [JOBS]. In the countries refugees come from, [CONFLICT]."

The four attributes took the following level texts, each randomized independently:

Attribute	Levels (text shown to respondents)
Aid level	"100% of food rations and cash support ('bamba chakula')" — "40% of food rations and cash support ('bamba chakula')" — "0% (no food rations or cash support ('bamba chakula'))"
Service access	"shared with the host community (refugees and hosts can both use them)" — "for refugees only"
Local jobs	"hire local residents" — "do not hire local residents"
Origin context	"conflict is ongoing" — "peace has been restored"

Example profile. "In this scenario, international organizations provide 40% of food rations and cash support ('bamba chakula'). Health care and schools are for refugees only. International organizations hire local residents. In the countries refugees come from, conflict is ongoing."

Outcome question (after each scenario). "On a scale from 1 (strongly oppose) to 5 (strongly support), how much do you support Kenya hosting refugees under this scenario?" [1 = strongly oppose / 2 = somewhat oppose / 3 = neither support nor oppose / 4 = somewhat support / 5 = strongly support].

Manipulation check (asked once, after the final scenario). "In the scenario you just saw, what was the level of food/cash support?" [100% / 40% / 0%].

1.5 Estimation equations

This section reports the estimating equations for analyses in the main paper and appendix. All models are estimated by ordinary least squares, with linear probability models for binary outcomes. Cross-sectional specifications use heteroskedasticity-robust (HC1) standard errors; panel specifications cluster standard errors by respondent.

1.5.1 Nationwide difference-in-differences (H1)

The pooled difference-in-differences in Table 2 of the paper is

$$Y_i = \alpha + \beta\,\mathrm{Post}_i + \rho\,\mathrm{Host}_i + \tau\,(\mathrm{Post}_i \times \mathrm{Host}_i) + \mathbf{X}_i'\Gamma + \varepsilon_i, \tag{1}$$

where Post_i indicates the 2026 round, Host_i indicates residence in a refugee-hosting county (Turkana or Garissa), and $\mathbf{X}_i$ contains age, gender, and an education index. The coefficient of interest is τ. Section 2.2 re-estimates equation (1) with the entropy-balancing post-stratification weights, Section 2.3 replaces HC1 inference with a wild cluster bootstrap clustered on county, and Section 2.4 varies the composition of the treated group. The Afrobarometer difference-in-differences in Section 2.5.2 estimates equation (1) with the 2024 Afrobarometer as the pre-period and with a different outcome variable.

1.5.2 Kakuma panel event study (H1)

The over-time analysis on the balanced three-wave panel (Table 3, Panel A of the paper) is

$$Y_{it} = \mu + \nu_i + \delta_2 D_t^2 + \delta_3 D_t^3 + \varepsilon_{it}, \tag{2}$$

with respondent fixed effects ν_i, wave indicators D_t^2 and D_t^3 (Wave 1 omitted), and standard errors clustered by respondent. The Wave 3 versus Wave 2 contrast is read from the same specification re-referenced to Wave 2. The repeated cross-section version (Section 3.3) replaces ν_i, which is unavailable once refresher respondents are included, with the demographic controls $\mathbf{X}_{it}$ from equation (1), retaining respondent-clustered standard errors:

$$Y_{it} = \mu + \delta_2 D_t^2 + \delta_3 D_t^3 + \mathbf{X}_{it}'\Gamma + \varepsilon_{it}. \tag{3}$$

The integration and trust outcomes (Section 3.5) and the four-wave extension (Section 3.4) use equation (2) with the corresponding outcomes and wave indicators.

1.5.3 Personal loss (H2)

All three tests of personal loss (discussed in Section 3.6.2) share the same estimator, a respondent-fixed-effects difference-in-differences on the balanced three-wave Kakuma panel, with standard errors clustered by respondent. They differ only in how they define which households are exposed to personal loss. Each takes the form

$$Y_{it} = \mu + \nu_i + \delta_2 D_t^2 + \delta_3 D_t^3 + \tau_2 \left(G_i \times D_t^2\right) + \tau_3 \left(G_i \times D_t^3\right) + \varepsilon_{it}, \tag{4}$$

with respondent fixed effects ν_i, wave indicators D_t^2 and D_t^3 (Wave 1 omitted), and an exposure indicator G_i interacted with each post wave. The interactions τ_2 and τ_3 give the differential change in exposed households' attitudes at the trough and at the partial-recovery wave. What changes across the three designs is the definition of G_i.

The baseline-beneficiary specification (Panel A) sets $G_i = B_i$, where B_i equals one if the household accessed camp healthcare or education or was employed by an NGO or a refugee in the 30 days before Wave 1. Because B_i is fixed before the cuts, it cannot be contaminated by attitudes that formed later, which makes this the intent-to-treat design and the primary test of H2. Table 3, Panel B of the paper reports the equivalent pooled parameterization, which replaces the respondent fixed effects with the beneficiary main effect ρB_i so that the pre-cut level gap between beneficiary and non-beneficiary households is displayed alongside the interactions:

$$Y_{it} = \mu + \delta_2 D_t^2 + \delta_3 D_t^3 + \rho B_i + \tau_2 \left(B_i \times D_t^2\right) + \tau_3 \left(B_i \times D_t^3\right) + \varepsilon_{it}. \tag{5}$$

In a fully balanced panel the fixed-effects and pooled parameterizations return identical interaction estimates; here they differ only through item-level nonresponse, and the estimates are nearly identical (Section 3.6.2).

The cohort difference-in-differences (Panel B) sets $G_i = L_i$, an indicator for first losing access to a benefit between Waves 2 and 3, estimated on the sample that excludes households that lost access between Waves 1 and 2. Here τ_2 is a placebo: because the loss occurs after Wave 2, the Wave 2 interaction should be near zero if exposed and unexposed households were on parallel trajectories beforehand, and τ_3 is the effect of the realized loss.

The switching difference-in-differences (Panel C) defines exposure by within-person loss events rather than a fixed group, pooling loss events across waves and domains:

$$Y_{it} = \nu_i + \lambda_t + \tau P_{it} + \varepsilon_{it}, \tag{6}$$

where P_{it} switches to one from the wave after household i first loses access in a domain, with respondent fixed effects ν_i, wave fixed effects λ_t, and respondent-clustered standard errors. The four domains (education, healthcare, refugee employment, and NGO employment) are estimated separately and pooled as loss in any domain, so τ recovers a single within-person effect of losing access.

1.5.4 Vignette experiment (H1 and H3)

Each respondent i rates six profiles j; the average marginal component effects in Figure 4 of the paper come from a single linear model of rated support on all attribute-level indicators,

$$Y_{ij} = \nu_i + \sum_a \sum_{\ell \neq \ell_0(a)} \beta_{a\ell}\, T_{ij}^{a\ell} + \varepsilon_{ij}, \tag{7}$$

where $T_{ij}^{a\ell}$ indicates that profile j presents level ℓ of attribute a (reference levels ℓ_0: aid at 100 percent, services shared, hire locals, conflict ongoing), with respondent fixed effects and standard errors clustered by respondent. The empathy moderation (Table 4 of the paper) interacts every attribute with the standardized two-item empathy index E_i:

$$Y_{ij} = \nu_i + \sum_a \sum_{\ell \neq \ell_0(a)} \left[\beta_{a\ell}\, T_{ij}^{a\ell} + \pi_{a\ell} \left(T_{ij}^{a\ell} \times E_i\right)\right] + \varepsilon_{ij}, \tag{8}$$

so that $\beta_{a\ell}$ is the contrast at mean empathy and $\pi_{a\ell}$ its change per standard deviation of empathy. Because E_i is absorbed by the respondent fixed effects, its main effect is not separately estimated. Marginal means by empathy group appear in Section 4.4.

1.6 Hosting Index construction

I combine the hosting items into a single index rather than analyzing each separately to avoid the multiple-comparisons problem that arises from testing several correlated outcomes one at a time. A summary index controls the family-wise error rate and, by pooling the shared signal across items, typically yields a more powerful test than any single item (Anderson 2008). I use Anderson's inverse-covariance weighting, which down-weights near-redundant components so that highly correlated items do not dominate the index.

The three items measure related but distinct facets of support for hosting: support for hosting nationally, support for hosting locally, and rejection of repatriation. Combining them requires only that they share a common orientation, which holds: all three pairwise correlations are positive in both the Kakuma panel and the nationwide baseline, with mean absolute correlations of 0.43 and 0.51 (Table 8), and Cronbach's alpha is 0.69 and 0.67. The weighting choice makes little difference here: because the components are similarly and only moderately correlated, the implied ICW weights are nearly equal, and the index correlates with a simple standardized average at 1.00 in the panel and 0.997 nationwide. The main results are unchanged under the simple average (Section 2.5.1).

Table 8: Hosting Index diagnostics

Statistic	Kakuma panel	Nationwide
Mean |r| among items	0.43	0.51
Cronbach's alpha	0.69	0.67
ICW vs. simple-average correlation	1.000	0.997

Note: Mean absolute pairwise correlation among the three index components, Cronbach's alpha, and the correlation between the inverse-covariance-weighted index and a simple standardized average. Panel components: Nation Hosting, Local Hosting, Oppose Closure. Nationwide components: Nation Hosting, Local Hosting, Not Repat.

2 Nationwide results and robustness

2.1 Summary statistics

Table 9 reports the distribution of the four outcomes by round for the national sample and the Turkana subsample. Nationally, all four outcomes increase or hold steady between 2023 and 2026: the hosting index rises from −0.12 to 0.05, and the share not selecting repatriation rises from 64 to 76 percent. In Turkana, by contrast, the index falls from 0.24 to 0.09 and local hosting support falls from 3.31 to 2.98, while national hosting support and the repatriation measure hold roughly constant. The table also shows that Turkana was more pro-hosting than the national average on every measure in 2023, a gap that closes by 2026.

Table 9: Nationwide refugee attitudes: outcome distributions by round

Outcome	Year	N	Mean	SD	P25	Med.	P75	Min	Max
National									
Hosting Index	2023	814	-0.12	1.05	-0.80	0.13	0.86	-1.83	1.23
	2026	2,041	0.05	0.88	-0.38	0.08	0.81	-2.04	1.37
Nation Hosting	2023	809	3.55	1.59	2.00	4.00	5.00	1.00	5.00
	2026	1,776	3.61	1.61	2.00	4.00	5.00	1.00	5.00
Local Hosting	2023	810	2.99	1.76	1.00	4.00	5.00	1.00	5.00
	2026	2,028	3.14	1.77	1.00	4.00	5.00	1.00	5.00
Not Repat.	2023	801	0.64	0.48	0.00	1.00	1.00	0.00	1.00
	2026	2,018	0.76	0.43	1.00	1.00	1.00	0.00	1.00
Turkana									
Hosting Index	2023	150	0.24	0.91	-0.18	0.60	1.07	-1.83	1.07
	2026	349	0.09	0.94	-0.33	0.03	1.07	-1.90	1.37
Nation Hosting	2023	150	4.15	1.33	4.00	5.00	5.00	1.00	5.00
	2026	72	3.94	1.61	4.00	5.00	5.00	1.00	5.00
Local Hosting	2023	150	3.31	1.78	1.00	4.00	5.00	1.00	5.00
	2026	346	2.98	1.78	1.00	3.50	5.00	1.00	5.00
Not Repat.	2023	148	0.80	0.40	1.00	1.00	1.00	0.00	1.00
	2026	343	0.80	0.40	1.00	1.00	1.00	0.00	1.00

Note:
National means and standard deviations are weighted (entropy-balancing post-stratification weights); Turkana means and standard deviations are unweighted, since the post-stratification weights target the national population. Percentiles, minimum, and maximum are unweighted throughout.

2.2 Weighted difference-in-differences analysis

The difference-in-differences estimates in the main text are unweighted with HC1 standard errors, as pre-specified (PAP Section 4.4). Table 10 re-estimates all four specifications applying the entropy-balancing post-stratification weights described in Section 1.3. The weights are constructed for the nationally representative samples, so the 2026 Turkana oversample carries no post-stratification weight. Dropping these respondents would reduce the treated post-period cell from 383 to 118 and discard most of the treated-group information in the weighted specification. I therefore retain the oversample, assigning its respondents unit weight and normalizing weights to mean one within each period-by-hosting-county cell so that the two weight systems sit on a common scale. This keeps all treated respondents in the weighted estimation while weighting the nationally representative respondents by their post-stratification weights. Point estimates are stable across panels, and the *Hosting Index*, *Local Hosting*, and *Not Repat.* effects remain significant; only *Nation Hosting* loses significance under the weights, reflecting the larger standard errors that weighting introduces rather than a change in sign or magnitude.

2.3 Wild cluster bootstrap

Because treatment is assigned at the county level and only two counties are treated (Turkana and Garissa), the analytic cluster-robust standard errors risk overstating precision. Table 11 therefore recomputes inference for the interaction term using a wild cluster bootstrap with Rademacher weights, clustering on county and referencing a t distribution with $G - 1$ cluster degrees of freedom. The *Hosting Index*, *Local Hosting*, and *Not Repat.* effects all remain significant under this more demanding standard, with bootstrap confidence intervals that exclude zero. Only the *Nation Hosting* estimate attenuates, from conventional significance to marginal ($p = 0.078$). This is consistent with the panel's finding that there is most backlash to refugee hosting locally.

Table 10: Nationwide difference-in-differences, unweighted and weighted

V1	Hosting Index	Nation Hosting	Local Hosting	Not Repat.
Panel A: Unweighted (main specification)				
Post	0.114**	-0.027	0.099	0.079***
	(0.040)	(0.066)	(0.072)	(0.019)
Hosting county	0.326***	0.562***	0.381**	0.109***
	(0.063)	(0.095)	(0.120)	(0.028)
Post x Hosting county	-0.365***	-0.418*	-0.659***	-0.099**
	(0.081)	(0.196)	(0.155)	(0.036)
Observations	3,401	2,841	3,381	3,357
Panel B: Entropy-balancing weights				
Post	0.159**	0.034	0.167†	0.097***
	(0.056)	(0.090)	(0.095)	(0.026)
Hosting county	0.395***	0.671***	0.461***	0.134***
	(0.072)	(0.110)	(0.133)	(0.032)
Post x Hosting county	-0.343***	-0.130	-0.584**	-0.109*
	(0.098)	(0.252)	(0.191)	(0.046)
Observations	3,401	2,841	3,381	3,357

Note:
OLS with HC1 standard errors in parentheses; demographic controls (age, gender, education) included in all models. Panel B weights observations by the entropy-balancing post-stratification weight for each survey round. All outcomes coded so higher values are more pro-refugee. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 11: Wild cluster bootstrap inference for the nationwide difference-in-differences

	Hosting Index	Nation Hosting	Local Hosting	Not Repat.
Post × Hosting county	-0.365	-0.418	-0.659	-0.099
Analytic (HC1) p	0.000	0.033	0.000	0.006
Bootstrap SE	(0.092)	(0.232)	(0.209)	(0.020)
Bootstrap p	0.000***	0.078†	0.003**	0.000***
Bootstrap 95% CI	[-0.551, -0.180]	[-0.885, 0.049]	[-1.080, -0.239]	[-0.140, -0.057]
Clusters	49	49	49	49

Note: Wild cluster bootstrap standard errors (Rademacher weights, 9,999 replications) clustered on county. Bootstrap p-values and confidence intervals reference a t distribution with G-1 cluster degrees of freedom. Analytic p-values are two-sided from HC1 standard errors. Significance from the bootstrap p-value: *** < .001, ** < .01, * < .05, † < .10.

2.4 County leave-one-out analysis

The treated group in the nationwide difference-in-differences pools two counties whose exposure and measurement differs: Turkana hosts Kakuma and is measured through a 2026 oversample, while Garissa hosts Dadaab, has a majority-Somali host population co-ethnic with most Dadaab refugees, and is included in the 2026 round through the nationally representative sample only. Table 12 therefore re-estimates equation (1) varying the composition of the treated group: excluding Garissa (Turkana as the sole treated county) and excluding Turkana (Garissa as the sole treated county), dropping the excluded county's respondents from the estimation sample entirely. As a complementary check on the control group, I also re-estimate the *Hosting Index* specification dropping each control county one at a time and report the range of the resulting interaction estimates in the table note.

The *Hosting Index* and *Local Hosting* interactions remain negative and significant whether Turkana or Garissa is the sole treated county, so neither drives the pooled result. The Turkana-only estimates are close to the pooled ones, as expected given Turkana's larger share of the treated sample; the Garissa-only estimates rest on far fewer treated post-period observations, since Garissa received no 2026 oversample, and their magnitudes should be read with that lower precision in mind. Dropping each control county in turn leaves the *Hosting Index* interaction in a tight band around the pooled estimate, so no single control county drives the comparison group either.

Table 12: Nationwide difference-in-differences: leave-one-out by treated county

Post × Hosting county	Both (main)	Excl. Garissa	Excl. Turkana
Hosting Index	-0.365***	-0.252**	-0.640***
	(0.081)	(0.096)	(0.171)
Nation Hosting	-0.418*	-0.157	-0.925**
	(0.196)	(0.228)	(0.353)
Local Hosting	-0.659***	-0.424*	-1.388***
	(0.155)	(0.189)	(0.338)
Not Repat.	-0.099**	-0.071†	-0.126
	(0.036)	(0.043)	(0.081)
N (treated, 2026)	383	349	34
Observations	3,401	3,216	2,902

Note: Each cell reports the Post x Hosting county interaction from the nationwide difference-in-differences, OLS with HC1 standard errors in parentheses and demographic controls (age, gender, education). Columns vary the treated group: both hosting counties, Turkana only, and Garissa only. Sample sizes are for the hosting-index specification. Dropping each of the 47 control clusters one at a time leaves the hosting-index interaction between -0.375 and -0.358 (pooled estimate -0.365; the interaction is estimable for 46 of the 47 drops). Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

2.5 Additional dependent variables

2.5.1 Hosting index robustness

The *Hosting Index* used in the main text is an inverse-covariance-weighted (ICW) composite of three five-point and binary items that mean-imputes any missing component, so respondents who answered only some of the three items still contribute. I check that the nationwide result is not an artifact of these construction choices. Table 13 re-estimates the pooled difference-in-differences with three alternative dependent variables alongside the main index: a complete-case ICW index, restricted to respondents with all three components present; a simple average of the three standardized components in place of the ICW weighting; and an ICW index built from binarized components, recoding each five-point item as an indicator for support before constructing the index.

All three alternatives return the same conclusion as the main table: the Post × Hosting county coefficient stays negative, significant at the same level, and comparable in magnitude. If anything, the complete-case restriction yields a slightly steeper decline, so the result is not driven by respondents with partial responses.

2.5.2 Afrobarometer refugee-numbers measure

The nationwide difference-in-differences spans a long window, from 2023 to 2026, over which attitudes could have shifted for reasons unrelated to the cuts. As an interim benchmark, I use a refugee-numbers item from Afrobarometer Round 10, fielded in mid-2024, after the initial WFP ration reductions but before the major 2025 cuts, and replicated verbatim in the 2026 follow-up. The item asks whether Kenya should host more refugees, about the same number, fewer, or none; I code an indicator for preferring the same number or more. Table 14 reports response shares by round, and Table 15 estimates the pooled difference-in-differences of equation 1, with the 2024 Afrobarometer sample as the pre period. The specification includes age, gender, and education, harmonized to the survey coding as in the weighting procedure (Section 1.3).

Two findings emerge. First, the national trend matches the main analysis: between 2024 and 2026, the share of Kenyans preferring that the country host the same number of refugees or more nearly doubled, from 19.9 to 38.0 percent ($\hat{\beta} = 0.177$, $p < 0.001$). The rise in support outside the hosting regions documented in Table

Table 13: Nationwide difference-in-differences with alternative index constructions

	Hosting Index (ICW)	Hosting Index (ICW, complete cases)	Hosting Index (simple average)	Hosting Index (binary components)
Post	0.114**	0.096*	0.101*	0.107**
	(0.040)	(0.041)	(0.040)	(0.040)
Hosting county	0.326***	0.331***	0.337***	0.301***
	(0.063)	(0.063)	(0.063)	(0.062)
Post × Hosting county	-0.365***	-0.445***	-0.363***	-0.333***
	(0.081)	(0.107)	(0.081)	(0.079)
Controls	Yes	Yes	Yes	Yes
Observations	3,401	2,795	3,401	3,401
R^2	0.040	0.046	0.039	0.039

2 of the main paper therefore appears on an independently fielded measure over a shorter, post-2024 window, and is not an artifact of the 2023 baseline. Second, hosting counties moved in parallel with the rest of the country on this measure ($\hat{\tau} = 0.004$, s.e. 0.067). This pattern is consistent with the integration results below and with the main finding: exposed counties track the rest of the country on support for hosting refugees in Kenya broadly, while the decline documented in the main text is specific to willingness to host refugees locally. However, this null should be read with caution: Afrobarometer includes only 78 hosting-county respondents, so the minimum detectable effect for the interaction is roughly 18 percentage points.

Table 14: Preferred number of refugees hosted, 2024–2026

	National		Hosting counties	
	2024	2026	2024	2026
None	20.2	13.0	17.4	7.6
Fewer	59.9	48.9	66.4	60.1
About the same	9.6	19.1	4.8	13.7
More	10.3	18.9	11.3	18.6
Same or more (%)	19.9	38.0	16.1	32.3
N	2,369	1,757	78	100

Note: Weighted response shares. 2024: Afrobarometer Round 10; 2026: follow-up survey. Hosting counties are Turkana and Garissa.

2.5.3 Integration measures

The 2026 survey includes measures of support for refugee integration: refugees' access to healthcare, right to work, and freedom of movement, each on a five-point scale. Table 16 reports weighted national means and compares Turkana with the rest of the country. Nationwide support is moderately high, greatest for healthcare, then right to work, then freedom of movement. Turkana respondents are statistically indistinguishable from the rest of the country on all three items; the point estimates are small and negative and none approaches significance. That exposed and non-exposed counties now look alike is consistent with the convergence documented in the main text. MacDonald and Lichtenheld (2025) find that support for integration was higher in refugee-hosting counties in 2023, but by 2026 that difference between refugee-hosting counties and the rest of the country is no longer evident. I read this as suggestive rather than as a within-design estimate, since it compares a 2026 cross-section against a separately fielded 2023 study rather than tracking the same integration items over time. The broader pattern is the same across measures: exposure to refugees in these counties no longer predicts higher support than the rest of the country.

Table 15: Difference-in-differences: preferred refugee numbers, 2024–2026

	Same/More refugees
Post	0.177***
	(0.015)
Hosting county	-0.033
	(0.046)
Post × Hosting county	0.004
	(0.067)
Controls	Yes
N (hosting counties, 2024)	78
N (hosting counties, 2026)	100
Observations	4,112
R^2	0.040

Note: OLS with HC1 standard errors. Outcome is an indicator for preferring that Kenya host the same number of refugees or more. Pre period is the 2024 Afrobarometer; post is the 2026 follow-up. Controls: age, gender, education. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

Table 16: Support for refugee integration, 2026 follow-up

Outcome	National (wt.)	Rest of Kenya	Turkana	Difference	N
Access to healthcare	4.05	4.08	4.02	-0.07 (0.09)	2,315
Right to work	3.45	3.51	3.36	-0.15 (0.10)	2,320
Freedom of movement	3.13	3.13	3.00	-0.13 (0.11)	2,322

Note: All items are five-point scales, with higher values indicating more supportive attitudes. The National column is the weighted mean on the nationally representative sample (post-stratification weights). The remaining columns come from a separate OLS regression of each outcome on a Turkana indicator, estimated unweighted on all 2026 respondents including the Turkana oversample, with HC1 standard errors in parentheses. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

2.6 Alternative estimators

Beyond the index, I probe the two five-point outcomes individually with alternative estimators. Table 17 reports three specifications for each item: Panel A repeats the OLS estimate from the main table; Panel B recodes the outcome as a binary indicator for support, a rating of four or five, and estimates a linear probability model; and Panel C estimates an ordered logit on the full five-point scale, relaxing the assumption that the response categories are equally spaced.

The *Local Hosting* result is robust across all three. For *Nation Hosting*, the ordered logit matches the OLS result, but the binary recoding is no longer significant, indicating that the *Nation Hosting* movement occurs partly within the oppose side of the scale rather than across the support threshold. This is consistent with the wild cluster bootstrap results above: the *Nation Hosting* estimate is the least robust of the four outcomes.

3 Kakuma panel results and robustness

3.1 Summary statistics

Table 18 reports the distribution of the four panel outcomes at each wave for the three-wave sample. Opposition to camp closure falls from 3.39 at Wave 1 to 3.13 at the Wave 2 trough and does not recover by Wave 3 (3.09). Local hosting support is flat. Support for hosting refugees in Kenya nationally, by contrast, dips at Wave 2 (3.12 to 2.96) and then rises above its baseline by Wave 3 (3.40), and the hosting index correspondingly returns to its Wave 1 level (0.00, −0.10, −0.02).

Table 17: Nationwide difference-in-differences: binary recoding and ordered logit

	Nation Hosting	Local Hosting
Panel A: OLS (five-point outcome)		
Post × Hosting county	-0.418*	-0.659***
	(0.196)	(0.155)
Panel B: LPM, binary outcome (support ≥ 4)		
Post × Hosting county	-0.088	-0.174***
	(0.055)	(0.044)
Panel C: Ordered logit (latent scale)		
Post × Hosting county	-0.587*	-0.776***
	(0.243)	(0.166)
Observations	2,841	3,381

Table 18: Kakuma panel refugee attitudes: outcome distributions by wave

Outcome	Wave	N	Mean	SD	P25	Med.	P75	Min	Max
Hosting Index	W1	205	0.00	1.00	-0.84	0.12	1.02	-1.55	1.26
	W2	205	-0.10	1.02	-1.07	0.08	0.80	-1.55	1.26
	W3	205	-0.02	0.94	-0.61	0.08	0.79	-1.55	1.26
Nation Hosting	W1	205	3.12	1.77	1.00	4.00	5.00	1.00	5.00
	W2	205	2.96	1.76	1.00	4.00	5.00	1.00	5.00
	W3	204	3.40	1.74	1.00	4.00	5.00	1.00	5.00
Local Hosting	W1	205	3.11	1.76	1.00	4.00	5.00	1.00	5.00
	W2	205	3.10	1.76	1.00	4.00	5.00	1.00	5.00
	W3	203	3.03	1.81	1.00	4.00	5.00	1.00	5.00
Oppose Closure	W1	205	3.39	1.79	1.00	4.00	5.00	1.00	5.00
	W2	204	3.13	1.82	1.00	4.00	5.00	1.00	5.00
	W3	204	3.09	1.86	1.00	4.00	5.00	1.00	5.00

Note: Statistics are computed across the three face-to-face waves of the balanced Kakuma panel (respondents observed in all three waves) and are unweighted. All outcomes are oriented so that higher values indicate more pro-refugee attitudes. Nation Hosting, Local Hosting, and Oppose Closure are five-point items; Oppose Closure reverses the camp-closure item. The Hosting Index is a standardized inverse-covariance-weighted (Anderson 2008) combination of the three, fit on Wave 1 (mean 0, SD 1).

3.2 Over-time welfare tests

Figure 2 in the main text plots the share of Kakuma panel households reporting each livelihood, security, and service-access outcome at each wave. This section reports the formal tests behind that figure. For each of the eight outcomes I estimate a linear probability model of the binary outcome on wave indicators, with standard errors clustered by respondent, and record the wave-to-wave differences (W2−W1, W3−W1, W3−W2), a joint test that the three waves are equal, and the implied prevalence at each wave. Consistent with the small per-wave sample (roughly 200 balanced-panel households), most movements are imprecise: only the rise in feeling unsafe at night, the decline in camp healthcare access, and the Wave 3 fall in crime victimization are statistically significant. Rows are ordered by Wave 3 prevalence to match the figure.

3.3 Repeated cross-section

The panel results section in the main paper analyzes the 205 respondents interviewed in all three waves, which isolates within-person change but discards the refresher respondents recruited at Waves 2 and 3 to replace attriters. Here I reproduce the over-time analysis on the full repeated cross-section. This trades the

Table 19: Host livelihood outcomes across the Kakuma panel

Outcome	Prevalence (%) W1	W2	W3	Wave difference (pp) W2−W1	W3−W1	W3−W2	Joint p
Feels unsafe at night	58.1	69.3	71.6	+11.1** (4.1)	+13.4** (4.4)	+2.3 (4.3)	0.003
Crime victim in past month	52.7	52.2	42.9	-0.5 (3.9)	-9.8* (4.3)	-9.3* (4.1)	0.037
Uses camp healthcare	42.4	30.2	26.8	-12.2** (3.9)	-15.6*** (4.5)	-3.4 (4.1)	0.001
Uses camp education	20.0	20.0	17.1	+0.0 (3.1)	-2.9 (3.5)	-2.9 (3.3)	0.623
Household has refugee job	13.2	11.7	14.1	-1.5 (2.4)	+1.0 (3.0)	+2.4 (3.0)	0.683
Household has NGO job	19.5	17.1	14.1	-2.4 (3.1)	-5.4 (3.5)	-2.9 (3.2)	0.305
Severe food insecurity	12.3	17.2	13.2	+5.0 (3.2)	+1.0 (3.2)	-4.0 (3.1)	0.243
Business serving refugees	13.2	16.6	13.2	+3.4 (2.9)	+0.0 (3.1)	-3.4 (3.3)	0.444

Note: Each row is a linear probability model of a binary household outcome (coded 0/100) on wave indicators, estimated on the balanced Kakuma panel with standard errors clustered by respondent. Prevalence columns give the fitted share reporting each outcome at each wave; difference columns give the pairwise wave contrasts in percentage points, with the clustered standard error in parentheses. Joint p is a Wald test that the three waves are equal. Rows are ordered by Wave 3 prevalence. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

panel's within-person identification for a larger sample, and serves as a check that the panel results are not an artifact of who remained in the study.

Because the repeated cross-section changes composition across waves — refreshers recruited at Waves 2 and 3 replace attriters, and attrition, while unrelated to Wave 1 observables (Section 1.3.2), was concentrated in the W1–W2 window — raw wave means could conflate genuine attitude change with compositional drift. I therefore estimate covariate-adjusted wave contrasts, regressing each outcome on wave indicators and controls for age, gender, and education (equation (3)), with standard errors clustered on respondent to account for the repeat interviews of panel members who reappear across cross-sections. In place of the balanced panel's respondent fixed effects, the covariate adjustment holds the demographic composition of each wave fixed. Table 20 reports the results, using the same three outcomes and the same two contrasts (Wave 2 vs. Wave 1, and Wave 2 vs. Wave 3) as the balanced-panel event study in the main text (Table 3, Panel A of the paper).

The results mirror the balanced panel. Opposition to camp closure falls from Wave 1 to Wave 2 by 0.332 points ($p < 0.05$), a steeper and now conventionally significant version of the marginal decline in the panel, and does not recover by Wave 3. Local hosting support shows no significant over-time change. *Nation Hosting* support decreases significantly at Wave 2 (-0.275, $p < .05$) before rising over the stabilization period (0.629, $p < .001$ from Wave 2 to Wave 3) to end above its baseline. The main results therefore hold when including refresher respondents in the sample.

3.4 Fourth wave: longer-term changes

This subsection extends the panel to a fourth wave, fielded by phone in February 2026 alongside the nationwide follow-up survey, roughly five months after Wave 3. The design captures the period of partial aid reinstatement described in Section 3.2 of the main paper: between Waves 3 and 4, rations were raised for Categories 1 through 3 and cash assistance resumed, though not to pre-cut levels. Wave 4 therefore extends the observation window from the acute-crisis phase into a period of partial recovery, allowing me to ask whether the Wave 2–3 decrease in support persisted as conditions stabilized.

Two features of Wave 4 warrant caution, and I read its results as suggestive of longer-run trajectories rather than as confirmatory. First, the switch from face-to-face to phone administration introduces a change in survey mode that could shift responses independently of any real attitude change. Second, attrition is substantially higher than across the three in-person waves: of the 120 Wave 4 respondents, 83 are present in all four waves. To hold sample composition fixed, Table 21 restricts to these 83 respondents, tracing within-respondent change in support for local hosting and opposition to camp closure from May 2025 through February 2026.

Table 20: Within-Kakuma hosting attitudes over time, repeated cross-section

	Oppose Camp Closure	Local Hosting	Nation Hosting
Wave 2 vs Wave 1	-0.332*	-0.111	-0.275*
	(0.140)	(0.123)	(0.128)
Wave 3 vs Wave 2	0.057	0.042	0.629***
	(0.143)	(0.166)	(0.155)
Respondent-wave obs.	722	734	735
Respondents	301	302	302

Note: OLS on the full repeated cross-section of Kakuma host-community respondents (at least 300 per wave, including refreshers), controlling for age, gender, and education, with standard errors clustered by respondent in parentheses. Composition drift is absorbed through the demographic controls rather than the respondent fixed effects used in the balanced-panel event study (paper Table 3, Panel A). Wave 3 vs Wave 2 is estimated from a specification re-referenced to Wave 2. All outcomes are five-point scales oriented so that higher values indicate more pro-refugee attitudes. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

Table 21: Within-Kakuma hosting attitudes over four waves, balanced panel

	Oppose Camp Closure	Local Hosting	Nation Hosting
Wave 2 vs Wave 1	-0.481*	0.169	0.024
	(0.232)	(0.209)	(0.226)
Wave 3 vs Wave 2	0.073	-0.331	0.325
	(0.241)	(0.243)	(0.212)
Wave 4 vs Wave 3	0.336	-0.030	0.525
	(0.270)	(0.258)	(0.399)
Respondent-wave obs.	330	331	257
Respondents	83	83	83

Note: OLS with respondent fixed effects on the balanced four-wave panel (respondents interviewed in all of Waves 1–4); standard errors clustered by respondent, in parentheses. Each contrast is estimated from a specification re-referenced to the preceding wave, so rows give consecutive wave-on-wave change. Wave 4 is the February 2026 phone wave. Higher values indicate more pro-refugee attitudes. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

Across Waves 3 and 4 none of the changes are statistically significant. Descriptively, opposition to camp closure continues to recover, nation hosting rises further, and local hosting is essentially flat. The pattern is consistent with a partial recovery, or at least no further backlash, as aid was partly reinstated. However, with only 83 respondents these estimates are underpowered, and I read the trajectories as suggestive.

3.5 Integration and trust measures

The main text shows that, over the period of the cuts, support for closing the camp rose in Kakuma even as support for hosting refugees in Kenya more broadly did not fall. This points to a distinction between opposition to hosting refugees locally and opposition to their presence in the country more generally: a host can become less willing to have refugees in their own area without opposing their country providing asylum at all. To probe this distinction further, I examine four additional panel measures that capture attitudes toward refugees beyond local hosting: support for three integration rights (access to healthcare, the right to work, and freedom of movement) and a five-point measure of trust in refugees. Each is estimated with the same event-study specification and balanced three-wave sample as the main panel analysis.

Table 22 shows no decline in any of the four measures. None of the Wave 2 versus Wave 1 differences is distinguishable from zero, so neither support for integration nor trust in refugees fell during the trough of the aid contraction. If anything, both moved in the opposite direction over the stabilization period: support for freedom of movement rose between Waves 2 and 3 (+0.465, $p < 0.01$) and trust in refugees rose marginally (+0.229, $p < 0.10$), mirroring the recovery in nationwide hosting support documented in the main text. Consistent with the empathy channel and the open-ended evidence, the decline in willingness to host refugees locally was not accompanied by a broader decline in support for refugee rights or trust in refugees.

Table 22: Refugee integration and trust over the Kakuma panel

	Access to healthcare	Right to work	Freedom of movement	Trust in refugees
Wave 2 vs Wave 1	-0.198	-0.039	-0.093	-0.173
	(0.151)	(0.116)	(0.130)	(0.121)
Wave 3 vs Wave 2	-0.053	0.132	0.465**	0.229†
	(0.155)	(0.145)	(0.144)	(0.117)
Respondent-wave obs.	609	615	614	606
Respondents	205	205	205	205

Note:
OLS with respondent fixed effects on the balanced three-wave Kakuma panel; standard errors clustered by respondent, in parentheses. Wave 3 vs Wave 2 is estimated from a wave factor re-referenced to Wave 2. All outcomes are five-point scales, with higher values indicating more supportive attitudes. Trust in refugees is a five-point agreement item with the statement that refugees can be trusted. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

3.6 Personal-loss difference-in-differences

3.6.1 Comparing baseline-beneficiaries with non-baseline-beneficiaries

Because selection into aid benefits is not random, Table 23 compares baseline-beneficiaries and non-baseline-beneficiaries. Baseline-beneficiaries ($N = 127$) are demographically similar to non-baseline-beneficiaries ($N = 78$) on gender and education, and marginally older (+2.45 years, $p = 0.05$). They differ, however, on pre-shock attitudes: at Wave 1, baseline-beneficiaries were substantially more supportive of hosting nationally (+0.73, $p < 0.01$) and locally (+0.93, $p < 0.001$), and substantially more opposed to closing the camp (-0.95, $p < 0.001$).

This difference across the two populations was expected and in-line with the paper's theory. Households tied to the refugee-aid economy plausibly hold more pro-refugee views both because they benefit from that economy and because pre-existing support may have increased their ability or desire to work and access services alongside refugees. The identifying assumption of the ITT specification is not baseline balance but parallel trends: the respondent fixed effects absorb these level differences, and the estimates in Table 24 identify the differential change in exposed households' attitudes after the cuts, not the difference in their levels. The remaining threat is differential trends, that exposed households would have converged toward unexposed households' attitudes even absent the cuts. The Wave 2 placebo in the cohort specification (Table 24, Panel B) speaks against this, and the triangulation across the three H2 designs in Section 3.6.2 addresses it further.

3.6.2 Alternative personal loss measures

The main text reports the baseline-beneficiary (intent-to-treat) specification as the primary test of H2, chosen because its treatment is fixed before the cuts and therefore cannot be contaminated by later attitudes. Here I report all three of the personal-loss specifications I estimated, which trade identification against power and define exposure in three different ways. Selection into loss is non-random in every case, so I read all three as

Table 23: ITT balance: baseline-beneficiary vs non-baseline-beneficiary

Variable	Baseline exposed (Mean, SD)	Unexposed (Mean, SD)	Difference	p-value
Age (years)	31.10 (9.72)	28.65 (7.92)	2.45†	0.051
Female (%)	0.47 (0.50)	0.51 (0.50)	-0.04	0.577
Education index	0.65 (0.48)	0.71 (0.46)	-0.05	0.443
Hosting support — national (1–5)	3.40 (1.71)	2.67 (1.78)	0.73**	0.004
Hosting support — local (1–5)	3.46 (1.70)	2.54 (1.73)	0.93***	0.000
Support for camp closure (1–5)	2.25 (1.64)	3.21 (1.88)	-0.95***	0.000
Empathy index (0–1, measured W3)	0.66 (0.35)	0.62 (0.39)	0.04	0.406

Note: Wave 1 characteristics by whether the household was an aid recipient or NGO/aid employee at baseline. Columns report mean and standard deviation for the treated (exposed) and control (unexposed) groups; Difference is treated minus control. Significance: *** $p < .001$, ** $p < .01$, * $p < .05$, † $p < .10$ (Welch's t-test).

quasi-experimental evidence on the egocentric channel rather than as definitive estimates of α_i. Their value lies in triangulation: each has a different weakness, and they point the same way. The three differ in how they define which households are exposed to personal loss, not in the estimator, which is a respondent-fixed-effects difference-in-differences throughout (equations (4) through (6)).

The first is the *baseline-beneficiary* specification (Panel A), an intent-to-treat design. Treatment is a fixed Wave 1 trait: whether the household received aid or contained an NGO or aid employee before the severe cuts. Because this indicator is measured before the cuts, it cannot be a product of how a household later came to feel about refugees, which makes it the most credible of the three on identification grounds. The cost is a weaker effect estimate because not every baseline beneficiary actually lost a benefit, so the "treated" group includes households that were not affected, which pulls the estimated effect downward. That the estimate remains significant despite this strengthens rather than weakens the inference.

One parameterization note applies to Panel A. The main text's Table 3 reports the intent-to-treat specification in pooled form (equation (5)), replacing the respondent fixed effects with a baseline-beneficiary main effect so that the pre-cut level gap between beneficiary and non-beneficiary households is displayed alongside the interactions. Panel A here re-estimates the same design with respondent fixed effects (equation (4)), matching the estimator used in Panels B and C. In a fully balanced panel the two parameterizations return identical interaction estimates; they differ here only through item-level nonresponse on the outcome variables, and the camp-closure interactions are nearly identical across the two.

The second is the *cohort difference-in-differences* (Panel B), which defines exposure by realized loss. It compares households that first lost access to a benefit between Waves 2 and 3 against households that never lost access across the three waves, using the Wave 1 to Wave 2 interval as a placebo. Because loss is realized rather than assigned, this is closer to a treatment-on-treated estimate and is more exposed to the concern that attitudes drive loss classification. Its strength is the explicit falsification test: the Wave 2 interaction is a placebo that should be near zero if the two groups were on parallel trends before the Wave 3 shock, isolating the effect of the loss itself.

The third is the *switching difference-in-differences* (Panel C), which defines exposure by within-household loss events. It pools all loss events across waves and domains, regressing the outcome on an indicator that switches on once a household has lost access in a domain, with respondent and wave fixed effects. This design has the most power and, because it estimates a separate coefficient for each domain, shows whether the result is driven by one channel or holds across education, healthcare, refugee employment, and NGO employment. Its weakness is that loss is measured post treatment and that it has no accompanying placebo test.

Table 24 reports all three specifications across the two local outcomes and nationwide hosting. The three converge on the same conclusion for *Oppose Camp Closure*: households exposed to personal loss moved toward more pro-closure attitudes relative to their unexposed neighbors. The baseline-beneficiary specification shows exposed households declining at both the trough and the partial-recovery wave ($\hat{\tau}_2 = -0.711$, $\hat{\tau}_3 = -0.712$,

Table 24: Within-Kakuma hosting attitudes by household exposure to the cuts

	Oppose Camp Closure	Local Hosting	Nation Hosting
Panel A: Intent-to-treat, baseline beneficiary status			
Baseline beneficiary × Wave 2	-0.711*	0.067	-0.102
	(0.311)	(0.269)	(0.289)
Baseline beneficiary × Wave 3	-0.712*	-0.676†	-0.603†
	(0.336)	(0.352)	(0.335)
Respondent-wave obs.	612	613	614
Respondents	204	205	205
Panel B: Cohort DiD, newly lost W2–W3 vs never lost			
Newly lost W2→W3 × Wave 2 (placebo)	-0.493	-0.165	0.171
	(0.389)	(0.342)	(0.357)
Newly lost W2→W3 × Wave 3	-0.840*	-0.577	-0.562
	(0.398)	(0.466)	(0.425)
Respondent-wave obs.	375	376	377
Respondents	125	126	126
Panel C: Switching DiD by loss domain			
Any domain	-0.606*	-0.325	-0.583**
	(0.240)	(0.238)	(0.222)
Education	-0.311	-0.030	-0.207
	(0.342)	(0.290)	(0.307)
Healthcare	-0.790**	-0.024	-0.357
	(0.272)	(0.257)	(0.224)
Refugee employment	-0.970*	-0.372	-0.655
	(0.440)	(0.383)	(0.399)
NGO employment	-0.629*	-0.494†	-0.321
	(0.262)	(0.293)	(0.292)
Respondent-wave obs.	612	613	614
Respondents	204	205	205

both $p < 0.05$). The cohort difference-in-differences finds a larger Wave 3 effect among new-loss households ($\hat{\tau}_3 = -0.840$, $p < 0.05$), with a Wave 2 placebo that is insignificant, consistent with parallel pre-trends. The switching specification gives an any-domain estimate of -0.606 ($p < 0.05$), with consistent signs across all four domains and significant declines for healthcare (-0.790, $p < 0.01$), refugee employment (-0.970, $p < 0.05$), and NGO employment (-0.629, $p < 0.05$); only education is individually null. The pattern of declining support among those who personally lose out holds across all three despite their different identifying assumptions.

4 Survey experiment results and robustness

4.1 Treatment arm balance

Each respondent rated six vignette profiles whose four attributes — the level of aid, whether services are shared with hosts, whether organizations hire locally, and whether conflict is ongoing — were independently randomized. If randomization succeeded, attribute assignment should be unrelated to respondents' pre-treatment characteristics. Table 26 reports, for each attribute, the share of profiles at each level, and tests whether that assignment is predicted by respondent age, gender, or education. Each column is a regression of an attribute-level indicator on the three covariates; the reported statistic is the p-value from a joint test that all three coefficients are zero, with standard errors clustered by respondent. The results show that no attribute is jointly predicted by the covariates, consistent with successful randomization.

Table 25: Intent-to-treat estimates with inverse-probability-of-retention weights

	Oppose Camp Closure	Local Hosting	Nation Hosting
Panel A: Unweighted (main specification)			
Baseline beneficiary × Wave 2	-0.711*	0.067	-0.102
	(0.311)	(0.269)	(0.289)
Baseline beneficiary × Wave 3	-0.712*	-0.676†	-0.603†
	(0.336)	(0.352)	(0.335)
Respondent-wave obs.	612	613	614
Panel B: Inverse-probability-of-retention weights			
Baseline beneficiary × Wave 2	-0.670*	0.094	-0.062
	(0.308)	(0.268)	(0.290)
Baseline beneficiary × Wave 3	-0.666†	-0.608†	-0.560†
	(0.338)	(0.352)	(0.335)
Respondent-wave obs.	612	613	614

Note: Respondent-fixed-effects intent-to-treat specification (baseline beneficiary status interacted with wave indicators) on the balanced three-wave Kakuma panel, standard errors clustered by respondent in parentheses. Panel B weights each retained respondent by the normalized inverse of the fitted probability of appearing in all three waves, from a logit of retention on Wave 1 age, gender, education, hosting support, and beneficiary status. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 26: Vignette treatment-arm balance

Level	Share	Profiles	Joint p
Aid level (ref: 100%)			
40%	0.32	1,809	0.57
0%	0.33	1,809	0.75
Service access (ref: shared)			
Refugees only	0.51	1,809	0.14
Local hiring (ref: hire locals)			
Do not hire locals	0.50	1,809	0.55
Conflict status (ref: ongoing)			
Peace restored	0.50	1,809	0.93

Note: Each row corresponds to one non-reference attribute level. Share is the proportion of rated profiles at that level; Profiles is the number of profile ratings in the estimation sample. Joint p is the p-value from a Wald test that respondent age, gender, and education jointly do not predict assignment to that level, from a linear probability model with standard errors clustered by respondent. The estimation sample is smaller than in the main experimental analyses because respondents with missing demographics are excluded.

4.2 Comprehension check

The vignette instrument included a comprehension check, pre-specified in the pre-analysis plan, asking respondents the level of food and cash support in the last scenario they rated. Comprehension was high: 97.7 percent of respondents answered correctly, and the pass rate was stable across the three aid levels (97.8 percent at 100 percent aid, 97.0 at 40 percent, and 98.7 at 0 percent). Table 27 re-estimates the average marginal component effects excluding those who failed, alongside the full-sample estimates. The estimates are essentially unchanged. Because comprehension was nearly universal and the results do not depend on excluding the few who failed the check, I retain the full sample in the main analysis.

Table 27: Vignette AMCEs excluding comprehension-check failures

	Full sample	Passed comprehension check
Aid: 40% (vs 100%)	-0.138†	-0.148†
	(0.083)	(0.085)
Aid: 0% (vs 100%)	-0.189*	-0.188*
	(0.080)	(0.082)
Hiring: do not hire locals (vs hire)	-0.203**	-0.188**
	(0.072)	(0.073)
Access: refugees only (vs shared)	-0.184*	-0.197*
	(0.076)	(0.077)
Conflict: peace restored (vs ongoing)	-0.633***	-0.604***
	(0.081)	(0.080)
Respondent-wave obs. (profiles)	1,883	1,853
Respondents	299	294

Note: Average marginal component effects from linear models with respondent fixed effects; standard errors clustered by respondent in parentheses. The outcome is rated support for Kenya hosting refugees (1–5) under each vignette. The full-sample column reproduces the average marginal component effects shown in Figure 4 of the main paper. The second column excludes respondents who failed the comprehension check. *** $p < 0.001$,** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

4.3 Empathy index construction and robustness

The vignette heterogeneity analyses use a two-item empathy index averaging concern about refugees and their suffering with worry that refugees lack enough food because of the cuts. A third item, worded in the opposite direction (that the respondent rarely feels compassion), was fielded alongside these two and reverse-coded to the common orientation. I exclude it from the primary index because it does not cohere with the other two: Table 28 reports the pairwise correlations, which are strong and positive between the two retained items but negative between the recoded compassion item and each of them. The two retained items form an internally consistent pair (Cronbach's $\alpha = 0.72$), whereas adding the compassion item yields a negative alpha (-0.31), indicating that the three items do not scale together and that the compassion item cannot be treated as a third measure of the same construct.

Table 29 shows that the substantive conclusion does not depend on this choice: re-estimating equation (8) with the three-item index in place of the two-item index leaves the sign and pattern of the Aid interactions unchanged, with the Aid 40% interaction attenuating slightly (0.230 to 0.220) and the Aid 0% interaction essentially unchanged (0.131 to 0.141). Because neither the empathy channel nor the Aid $\times$ empathy interaction was pre-registered, these analyses are exploratory and I read them as suggestive rather than confirmatory.

Table 28: Empathy items: pairwise correlations and internal consistency

	Item 1	Item 2	Item 3
Concern for refugees' suffering (item 1)			
Worry about refugees' food (item 2)	0.56		
Reverse-coded item (item 3)	-0.40	-0.39	
Cronbach's alpha: two-item / three-item	0.72	-0.31	

Note: Pairwise Pearson correlations among the three empathy items on the vignette respondents with all three items observed (N = 301). All items are oriented so that higher values indicate more empathy; item 3 is reverse-worded and recoded. The alpha row reports Cronbach's alpha for the two-item index (items 1–2) and the three-item index.

Table 29: Empathy moderation under the two-item and three-item indexes

	× Empathy (two-item index)	× Empathy (three-item index)
Aid: 40% (vs 100%)	0.230** (0.080)	0.220** (0.079)
Aid: 0% (vs 100%)	0.131 (0.080)	0.141† (0.078)
Hiring: Do not hire locals (vs hire)	0.046 (0.074)	0.058 (0.084)
Access: Refugees only (vs shared)	-0.136† (0.077)	-0.088 (0.073)
Conflict: Peace restored (vs ongoing)	0.030 (0.084)	0.084 (0.093)

Note: Attribute-by-empathy interaction coefficients from linear models with respondent fixed effects, standard errors clustered by respondent in parentheses. Each column standardizes the stated empathy index to mean zero and standard deviation one across respondents. The two-item column reproduces the main text's specification; the three-item column adds the reverse-coded item to the index. N (two-item) = 1,883 profiles; N (three-item) = 1,883 profiles. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

4.4 Empathy marginal means

Interaction terms in equation (8) compare conditional AMCEs, which depend on the reference level of each attribute. Following Leeper, Hobolt, and Tilley (2020), subgroup comparisons in factorial designs are more transparently reported as marginal means, the average rated support at each attribute level within each subgroup. Table 30 reports marginal means for every attribute level separately for respondents below and above the sample median of the empathy index, with standard errors clustered by respondent, and the high-minus-low difference at each level.

The marginal means reproduce the interaction results. Among low-empathy respondents, support falls as aid drops from 100 percent to 40 and 0 percent, while among high-empathy respondents support is essentially flat across aid levels; the high-minus-low differences are close to zero at full aid and positive under the cut conditions.

Table 30: Vignette experiment: marginal means by empathy group

Level	Low empathy	High empathy	Difference (High − Low)
Aid level			
100%	2.640 (0.105)	2.629 (0.114)	-0.012 (0.155)
40%	2.356 (0.100)	2.572 (0.109)	0.216 (0.148)
0%	2.386 (0.103)	2.599 (0.113)	0.212 (0.153)
Local hiring			
Hire locals	2.581 (0.094)	2.686 (0.108)	0.104 (0.143)
Do not hire locals	2.349 (0.088)	2.512 (0.095)	0.162 (0.129)
Service access			
Shared with hosts	2.480 (0.095)	2.715 (0.099)	0.235† (0.137)
Refugees only	2.448 (0.095)	2.485 (0.100)	0.037 (0.138)
Conflict status			
Conflict ongoing	2.751 (0.100)	2.892 (0.101)	0.141 (0.142)
Peace restored	2.187 (0.095)	2.298 (0.098)	0.111 (0.136)

Note: Marginal means of rated hosting support (1–5) at each attribute level, by empathy group (median split on the two-item empathy index), with standard errors clustered by respondent in parentheses. The difference column reports the high-minus-low contrast among profiles at that level. Marginal means average over the randomized distribution of the other attributes and do not depend on reference-level choices. Low-empathy group: 158 respondents; high-empathy group: 143 respondents. Significance (difference column): *** p<.001, ** p<.01, * p<.05, † p<.10.

4.5 Hypothetical full-cut scenario

Table 31: Hypothetical full-cut scenario: within-respondent change and gap to general hosting support

	Support
Panel A: Over-time change in full-cut support	
Wave 2 vs Wave 1 (5-point)	-0.545***
	(0.150)
Wave 2 vs Wave 1 (support ≥ 4)	-0.136***
	(0.040)
Respondent-wave obs.	396
Panel B: Gap to general hosting support	
Full-cut vs. general (Wave 1 gap)	-0.473**
	(0.153)
Full-cut × Wave 2	-0.367†
	(0.199)
Respondent-wave obs.	813

Note:
OLS with respondent fixed effects on the balanced Wave 1–Wave 2 panel; standard errors clustered by respondent, in parentheses. The outcome is stated support for Kenya hosting refugees if international donors withdrew aid completely, on a five-point scale (higher = more supportive); the binary row codes support as a rating of four or five. Panel A gives the within-respondent Wave 1 to Wave 2 change. Panel B stacks the full-cut item against general hosting support and reports the Wave 1 gap and its widening by Wave 2, from a scenario-by-wave interaction. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, † $p < 0.10$.

In Waves 1 and 2, I asked panel respondents how strongly they would support Kenya continuing to host refugees if international donors withdrew financial support completely, a more severe cut than the one observed. Because the item is within-respondent and repeated across the two waves that bracket the trough of the contraction, I use it in three ways: the level of support gauges attitudes under a hypothetical full withdrawal at each wave; the gap between this item and general support for hosting refugees in Kenya shows how attitudes differ within respondent under current conditions versus a complete cut; and the within-respondent change from Wave 1 to Wave 2 asks whether support fell as the cuts deepened. This item is a descriptive counterpart to the full-cut arm of the vignette experiment. Standard errors are clustered by respondent throughout, and these estimates are descriptive.

The findings support the main experimental result. At baseline, support under full withdrawal was 0.473 points lower than support for hosting refugees in Kenya generally ($p < 0.01$), and this gap widened by a further 0.367 points by Wave 2 ($p < 0.10$) (Table 31). Support for hosting under full withdrawal also fell sharply across waves, from 2.64 to 2.12 on the five-point scale ($\hat{\beta} = -0.545$, $p < 0.001$), while the share expressing support fell from 42.1 to 29.4 percent ($p < 0.001$). As the real cuts deepened and the prospect of hosting without aid became more concrete, opposition to this scenario grew. Yet even under hypothetical complete withdrawal, 29 percent still supported hosting, indicating that attitudes do not rest on aid alone.

5 Additional analyses and study documentation

5.1 Enumerator fixed effects

TIFA provided interviewer identifiers for the three Kakuma panel waves, the February 2026 nationwide survey, and the May 2026 recontact. Identifiers do not exist for the 2023 baseline. Table 32 re-estimates the personal-loss specifications of Table 24 with and without enumerator fixed effects. Table 32 shows that including enumerator fixed effects leaves the estimates substantively unchanged, suggesting that the declines in hosting support among personally exposed households is not an artifact of who conducted the interview.

I do not report an analogous check for the over-time wave coefficients: 61 percent of Wave 3 interviews were conducted by enumerators who did not work the first two waves, so enumerator fixed effects are nearly collinear with the Wave 3 indicator. For those estimates, interviewer composition is instead addressed by the repeated cross-section (Section 3.3), which recovers the over-time pattern on refreshed samples with a different enumerator mix at each wave.

Table 32: Within-Kakuma personal-loss estimates with enumerator fixed effects

	Oppose Camp Closure		Local Hosting		Nation Hosting	
	Base	Enum. FE	Base	Enum. FE	Base	Enum. FE
Panel A: Intent-to-treat, baseline beneficiary status						
Beneficiary × Wave 2	-0.711*	-0.719*	0.067	0.070	-0.102	-0.177
	(0.311)	(0.317)	(0.269)	(0.271)	(0.289)	(0.288)
Beneficiary × Wave 3	-0.712*	-0.826**	-0.676†	-0.695*	-0.603†	-0.712*
	(0.336)	(0.318)	(0.352)	(0.351)	(0.335)	(0.328)
Respondent-wave obs.	612	612	613	613	614	614
Enumerators		18		18		18
Panel B: Cohort DiD, newly lost W2–W3 vs never lost						
Newly lost × Wave 2 (placebo)	-0.493	-0.501	-0.165	-0.203	0.171	0.241
	(0.389)	(0.397)	(0.342)	(0.358)	(0.357)	(0.361)
Newly lost × Wave 3	-0.840*	-0.669†	-0.577	-0.509	-0.562	-0.551
	(0.398)	(0.405)	(0.466)	(0.509)	(0.425)	(0.452)
Respondent-wave obs.	375	375	376	376	377	377
Enumerators		18		18		18
Panel C: Switching DiD, any domain						
Lost aid (post)	-0.606*	-0.554*	-0.325	-0.242	-0.583**	-0.603**
	(0.240)	(0.242)	(0.238)	(0.242)	(0.222)	(0.229)
Respondent-wave obs.	612	612	613	613	614	614
Enumerators		18		18		18

Note: Same estimators, samples, and outcomes as Table 24, restricted to interviews with a matched enumerator identifier so paired columns share a sample. Base columns absorb respondent fixed effects (and wave fixed effects in Panel C); Enum. FE columns additionally absorb the wave-specific interviewer. Enumerator fixed effects are not reported for the over-time wave coefficients, which are collinear with enumerator identity because the roster changed across waves (see text). Standard errors clustered by respondent in parentheses. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

Table 33: Nationwide 2026 host-county gap with enumerator fixed effects

	Oppose Repat. (Feb)		Local Hosting (Feb)		Nation Hosting (May)	
	Base	Enum. FE	Base	Enum. FE	Base	Enum. FE
Host county	0.009	0.008	-0.290**	-0.308**	0.137	0.107
	(0.023)	(0.024)	(0.100)	(0.103)	(0.171)	(0.173)
Observations	2,293	2,293	2,306	2,306	1,769	1,769
Enumerators		34		34		21

Note: OLS on the 2026 sample, mirroring the post-period cell of the difference-in-differences specification: unweighted, with age, gender, and education controls and HC1 standard errors in parentheses. Host county indicates Turkana or Garissa. Each pair of columns is restricted to respondents with a non-missing outcome and enumerator identifier. The first two outcomes come from the February survey and the nationwide hosting measure from the May recontact; enumerator identifiers are not available for the 2023 baseline. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

The 2023 baseline data does not include enumerator identifiers, so I cannot examine enumerator effects across the full 2023–2026 difference-in-differences. But because the interaction is the 2026 host-county gap net of the fixed 2023 gap, showing that the 2026 gap is stable to enumerator changes provides indirect evidence that the design is not driven by interviewer effects. Table 33 re-estimates the 2026 host-county comparison with and without enumerator fixed effects, using the same controls and standard errors as the difference-in-difference. The fixed effects are identified because 31 of 36 February enumerators interviewed respondents in both host and non-host counties. The *Local Hosting* gap, the post-period component that drives the difference-in-differences result, moves from −0.290 to −0.308, and the repatriation and nationwide hosting gaps are indistinguishable from zero in both columns. Together with the panel results, this indicates that neither the within-Kakuma erosion nor the nationwide gap reflects which enumerators conducted the interviews.

5.2 Minimum detectable effects (MDE)

The Kakuma panel is small, so before reading any near-zero estimate as evidence of no effect, I report the minimum detectable effect (MDE) for each headline test. For a two-sided test at the 5% level with 80% power, MDE $= (z_{.975} + z_{.80}) \times \text{SE} = 2.80 \times \text{SE}$, computed from each model's realized clustered standard error. The MDE concern is most relevant for two near-zero estimates. The aggregate over-time change in the hosting index (−0.015 SD, MDE 0.236 SD) rules out aggregate shifts larger than roughly a quarter of a standard deviation, so it is best read as the absence of a large aggregate movement rather than of any movement. The aid slope within high-empathy respondents (−0.082 points, MDE 0.366) is similarly near zero, but its MDE is large enough that a moderate reduction of high-empathy support could be not identified. The other headline effects are detected rather than null, so the MDE is not a concern.

Table 34: Minimum detectable effects for the headline tests

Test	N	Est.	Clustered SE	MDE (80%)
Over-time panel: W3 vs W1 (hosting index)	615	-0.015	0.084	0.236
ITT: baseline beneficiary × W3 (camp closure)	612	-0.712	0.336	0.942
Vignette AMCE: Aid 0% vs 100% (1–5)	1,883	-0.189	0.080	0.226
Vignette: Aid 0% × empathy (1–5)	1,883	+0.131	0.080	0.223
Vignette: Aid 0% slope, high-empathy only (1–5)	889	-0.082	0.131	0.366

Note: Each row reports the minimum detectable effect (MDE) for one headline test, computed as 2.80 times the model's clustered standard error (two-sided, 5% level, 80% power). Est. and Clustered SE are the point estimate and standard error from the corresponding results table. Units are native to each specification: standard-deviation units for the hosting index, scale points for camp-closure support (1–5), and support points for the vignette outcome (1–5).

5.3 Shirika Plan

The nationwide difference-in-differences spans a three-year window, over which attitudes could have shifted for reasons unrelated to the cuts. The most plausible such confounder is the Shirika Plan, Kenya's 2021 framework (with a 2024 implementation plan) to transition refugee camps into integrated settlements. If hosts became aware of a policy to integrate refugees locally, that awareness could move hosting support independently of the aid cuts. Tables 35 and 36 assess this by comparing hosting support between respondents aware and unaware of the Plan, within survey round, nationwide and in the Kakuma panel.

Two features of the awareness data make the Shirika Plan an unlikely driver of the results. First, nationwide awareness is too low to move the aggregate: it rose over the study period but remained limited in both rounds (1.5 percent in 2023 and 11.4 percent in 2026), so the large majority of respondents had not heard of the Plan in either year. Among those who had, national hosting support was higher than among the unaware in both rounds, and local hosting support in 2026 was slightly lower, though not significantly so. Second, in the Kakuma host area awareness was already widespread and stable across waves (66.7 percent at Wave 3), so it cannot track the within-area decline documented in the main text. At Wave 3, respondents aware of the Plan were if anything more supportive of hosting and less supportive of closing the camp than the unaware. The Shirika Plan is therefore unlikely to account for the results: nationally it is too rare to shift the aggregate, and locally it is neither new nor associated with lower support.

Table 35: Nationwide support for hosting by awareness of the Shirika Plan

Outcome	Year	Aware	Unaware	Diff.	p
National hosting	2023	3.49	3.06	+0.44	0.028
National hosting	2026	3.78	3.58	+0.20	0.060
Local hosting	2023	3.39	2.58	+0.80	0.010
Local hosting	2026	2.92	3.17	-0.25	0.110

Note: Comparisons are within survey year. Outcomes are five-point items (higher = more support). Means survey-weighted; difference tested with an unweighted Welch t-test. Awareness was 1.5% in 2023 and 11.4% in 2026.

Table 36: Support for hosting by awareness of the Shirika Plan, Kakuma (Wave 3)

Outcome	Aware	Unaware	Diff.	p
National hosting	3.59	3.04	+0.54	0.037
Local hosting	3.14	2.82	+0.33	0.225
Close camp	2.72	3.25	-0.53	0.054

Note: Outcomes are five-point items; higher values indicate more support, except close camp, where higher indicates more support for closing the camp. Unweighted means; difference tested with an unweighted Welch t-test. Awareness in the Kakuma area was 66.7 percent at Wave 3.

5.4 Reversion to the mean

Both difference-in-differences designs compare groups that were initially more supportive of refugees. In the nationwide sample, refugee-hosting counties began about one-third of a standard deviation above the rest of the country on the *Hosting Index*. In the panel, baseline-beneficiary households were also more supportive of local hosting and less supportive of camp closure. If these initial differences reflected temporary shocks or measurement error, later convergence could be mistaken for an effect of the aid cuts. Three features of the results make this unlikely.

First, the baseline differences are consistent with prior evidence rather than unusually high observations. Research from Kenya and other refugee-hosting settings finds that communities and households connected to the aid economy tend to hold more favorable attitudes toward refugees, partly because hosting brings jobs, services, business opportunities, and infrastructure (Alix-Garcia et al., 2018; Zhou, Grossman and Ge, 2023; Kadigo and Maystadt, 2023; Baseler et al., 2025; Zhou et al., 2025; Demir, Mayda and Maystadt, 2025). MacDonald and Lichtenheld (2025) similarly find higher support for refugee integration in Kenya's hosting counties in 2023. The initial gaps therefore match established patterns.

Second, reversion cannot explain the selectivity of the nationwide decline. Because the surveys use repeated cross-sections of different respondents, individual-level measurement error cannot drive the result; the concern is instead that county-level sampling error inflated the 2023 hosting-county mean. But such reversion would pull down every correlated hosting outcome measured in the same counties, in proportion to its initial gap. Instead, the decline is concentrated in local hosting, which falls by roughly two-thirds of a standard deviation, while support for hosting refugees in Kenya more broadly moves far less and support for the various integration and refugee-numbers measures does not decline at all. Because hosting counties began above the national average on all of these measures, regression to the mean should have produced declines across them. The fact that only local hosting changed suggests that the result reflects attitudes toward hosting refugees nearby rather than sampling error.

Third, the panel results follow the timing of actual losses. Beneficiary status is defined by economic ties to the aid system, not by baseline attitudes. Households that lost benefits between Waves 2 and 3 moved in parallel with other households before the loss and declined only afterward. The largest changes also occurred in the specific domains in which access was lost, while non-beneficiary households remained stable across waves. A general return toward the mean would not predict this timing or domain-specific pattern.

5.5 Formal model

This appendix derives the three empirical implications stated in Section 2.3.1 of the paper from the model in equation (1) of the paper, restated below.

Restating the model

Each host i's latent utility from refugee hosting in period t is

$$H_{it} = \underbrace{\alpha_i \theta_{it}}_{\text{egocentric}} + \underbrace{\beta_i E_i[\sigma A_t - C(A_t)]}_{\text{sociotropic}} + \underbrace{\gamma_i[E_i S(A_t) + \tfrac{1}{2} S(A_t)]}_{\text{empathy}}, \tag{9}$$

with $C(A_t) = \bar{c} - A_t$ and $S(A_t) = \bar{s} - A_t$.

Marginal response to an aid change

Substituting the linear functional forms for $C(A_t)$ and $S(A_t)$,

$$H_{it} = \alpha_i \theta_{it} + \beta_i E_i[(\sigma + 1)A_t - \bar{c}] + \gamma_i (E_i + \tfrac{1}{2})(\bar{s} - A_t). \tag{10}$$

Taking the partial derivative with respect to A_t,

$$\frac{\partial H_{it}}{\partial A_t} = \alpha_i \frac{\partial \theta_{it}}{\partial A_t} + \beta_i E_i(\sigma + 1) - \gamma_i (E_i + \tfrac{1}{2}). \tag{11}$$

A unit reduction in aid moves support in three additive ways: a discrete jump from any personal transition in θ_{it}, a negative effect through the sociotropic channel scaled by $\beta_i E_i(\sigma + 1)$, and a positive effect through the empathy channel scaled by $\gamma_i(E_i + \frac{1}{2})$. The sociotropic and empathy channels both contain E_i but enter with opposite signs, so exposure simultaneously amplifies the sociotropic loss and the empathic gain.

The cut moves aid from A_0 to a lower level A_1, so $\Delta A = A_1 - A_0 < 0$, while the weights and exposure describing each host stay fixed. The change in latent utility is

$$\Delta H_i = \alpha_i \Delta \theta_i + \beta_i E_i(\sigma + 1)\, \Delta A - \gamma_i (E_i + \tfrac{1}{2}) \Delta A. \tag{12}$$

The three implications below show what equation (12) predicts when each source of variation is moved in isolation.

Implication 1: Variation in community exposure (E_i)

In a population with community exposure E, the expected aid-cut response is

$$\mathbb{E}[\Delta H_i \mid E] = -\alpha_i\, p(E) + \beta_i E(\sigma + 1)\, \Delta A - \gamma_i (E + \tfrac{1}{2}) \Delta A, \tag{13}$$

where $p(E)$ is the probability that a host with community exposure E undergoes $\Delta\theta_i = -1$ given a cut. I take expectations over θ_i given E and treat $\alpha_i, \beta_i, \gamma_i$ as population means within E for notational simplicity.

For a low-exposure population ($E = 0$, with $p(0) \approx 0$), the expected response operates only through the baseline empathy term, which captures concern generated through media coverage and indirect ties:

$$\mathbb{E}[\Delta H_i \mid 0] = -\tfrac{1}{2} \gamma_i \Delta A. \tag{14}$$

The differential between high- and low-exposure populations is

$$\mathbb{E}[\Delta H_i \mid E] - \mathbb{E}[\Delta H_i \mid 0] = -\alpha_i\, p(E) + E[\gamma_i - \beta_i(\sigma + 1)] \Delta A. \tag{15}$$

The half-empathy term contributes to both populations and cancels in the difference, leaving the full $\gamma_i E$ amplification in the differential. Writing $|\Delta A| = -\Delta A$ to make the signs explicit under an aid cut,

$$\mathbb{E}[\Delta H_i \mid E] - \mathbb{E}[\Delta H_i \mid 0] = \underbrace{-\,\alpha_i\, p(E)}_{\text{egocentric}}\,\underbrace{-\,\beta_i E(\sigma+1)|\Delta A|}_{\text{sociotropic}} + \underbrace{\gamma_i E|\Delta A|}_{\text{empathy}}. \tag{16}$$

The first term is unambiguously negative; the second is negative; the third is positive if $\gamma_i > 0$. The differential is negative — high-exposure populations show a larger decline in hosting support relative to low-exposure ones — whenever the combined egocentric and sociotropic decrements exceed the empathy gain:

$$\alpha_i\, p(E) + \beta_i E(\sigma+1)|\Delta A| > \gamma_i E|\Delta A|. \tag{17}$$

The prediction reverses only when empathy strictly dominates sociotropic concerns and the empathy excess is large enough to outweigh the egocentric decrement from the higher rate of personal loss in the exposed population. The modal prediction is therefore backlash. This yields H1: aid cuts reduce hosting support more in high-exposure than in low-exposure populations.

For completeness, note that the sociotropic contribution to support, $\beta_i E_i\big[(\sigma+1)A_t - \bar{c}\big]$, crosses zero at $A^* = \bar{c}/(\sigma+1)$. Above A^* the community still captures net benefits from the aid economy (declining as aid falls); below A^* it bears net costs. The sociotropic channel therefore flips sign as cuts deepen past this threshold, even though its marginal slope in A_t is constant under the linear C and S functions.

Implication 2: Variation in personal loss (θ_{it})

Holding E_i fixed across two otherwise comparable hosts, one of whom loses a benefit or incurs a cost ($\theta_{it} : 0 \to -1$), differences out the sociotropic and empathy channels and isolates the egocentric channel:

$$\Delta H_i^{\theta=-1} - \Delta H_i^{\theta=0} = \alpha_i \Delta\theta_i = -\alpha_i. \tag{18}$$

Personal loss creates an additional, individual-level decrement to hosting support equal to α_i, on top of any community-level shift associated with community-level exposure. This yields H2: within an exposed community, hosts who lose aid-related benefits show a steeper decline in support than otherwise-similar neighbors.

Implication 3: Variation in empathy (γ_i)

Holding both E_i and θ_{it} fixed, the aid-cut response moves linearly in γ_i:

$$\frac{\partial(\Delta H_i)}{\partial\gamma_i} = -(E_i + \tfrac{1}{2})\Delta A = (E_i + \tfrac{1}{2})|\Delta A| > 0. \tag{19}$$

Two hosts who share the same exposure, egocentric status, and sociotropic weighting respond to an identical aid cut in opposite directions if their empathy differs in sign. The slope of this heterogeneity is $(E_i + \frac{1}{2})|\Delta A|$, so the difference in responses between high- and low-γ hosts is itself larger in high-exposure areas. This yields H3: holding community exposure and personal loss constant, aid cuts reduce hosting support less among more empathetic hosts, and where empathy is sufficiently strong, support may increase.

5.6 Deviations from the pre-analysis plan

This study was pre-registered on the Open Science Framework before the first panel wave and before the trajectory of the aid cuts was known (registered April 18, 2025), with a January 2026 amendment adding the fourth phone wave and the 2026 nationwide follow-up. This section documents where the analysis departs from the pre-analysis plan (PAP) and why.

5.6.0.1 Dropping the two-town difference-in-differences. The PAP's primary design compared Kakuma with Isiolo town, selected as a control for its similar economy and political trends, but lack of a refugee presence. During the study period, Isiolo experienced security and political issues unrelated to the aid cuts that made it unsuitable as a control case. The major crisis happened in June 2025, when a motion to impeach the Isiolo governor led to community-wide violence and insecurity, including stone-throwing, teargas, and gunfire.[1] Field reports from the research team corroborated heightened insecurity during the study period.

These events likely affected the study's outcomes, particularly perceived safety, crime victimization, and attitudes toward government. Any Kakuma–Isiolo comparison estimated over this period may therefore confound the effect of the aid cuts with the effects of the political crisis and violence. All three survey waves were completed, but I dropped the two-town design after fieldwork concluded, on the basis of the events documented above, because the identifying assumption had failed rather than because of any feature of the estimates. For transparency, this appendix reports both the contamination evidence and the pre-registered specification.

Figure 1 plots the two towns' trajectories on the security outcomes and hosting support. The share of Isiolo respondents feeling unsafe at night rose 26 percentage points over the study period, from 33 to 59 percent, despite the absence of any aid shock in Isiolo, a steeper rise than in Kakuma itself. As a consequence, the pre-registered design returns a significant positive coefficient on relative safety in Kakuma (Table 37), which taken at face value would imply that the aid cuts improved security in the hosting area. On the attitude outcomes, the pre-registered contrast on the camp-closure item is negative and significant at Wave 3 (-0.532, $p < .05$), directionally consistent with the within-Kakuma results in the main text, but given the contamination I do not interpret any of these coefficients as effects of the aid cuts.

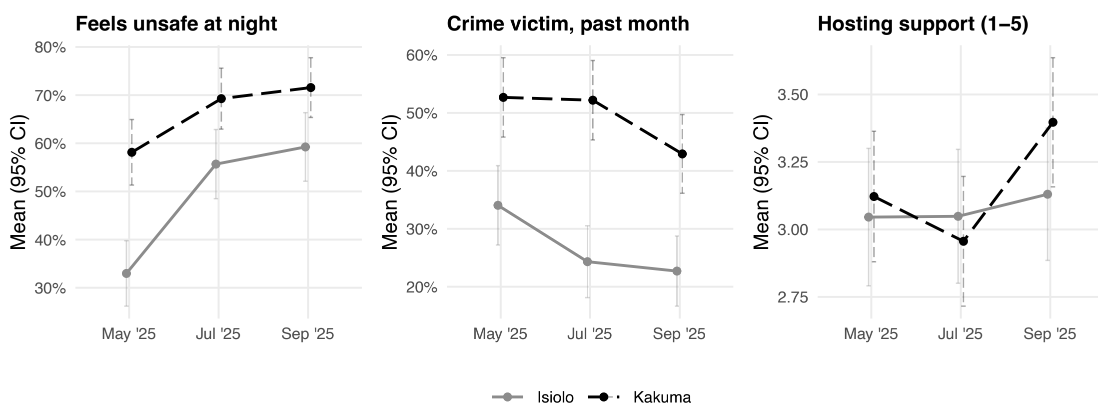

Figure 1: Two-town trajectories on security outcomes and hosting support

5.6.0.2 Promotion of the nationwide comparison to the primary H1 design. With the pre-registered control unusable, I use the two nationally representative surveys bracketing the cuts as the primary test of community-level exposure, comparing refugee-hosting counties to the rest of Kenya. Both data sources were described in the PAP (January 2026 amendment). Their combination into a difference-in-differences design, the inclusion of Garissa alongside Turkana as treated counties, the entropy-balancing post-stratification weights, and the wild cluster bootstrap inference were not pre-specified and were adopted in response to the failure of the two-town design. The Kakuma panel over-time analysis is unchanged from the PAP apart from the loss of the control town.

5.6.0.3 Reorganization of hypotheses. The PAP specified five hypotheses across three outcome families: economic and security outcomes (PAP-H1), attitudes toward refugees (PAP-H2), and attitudes toward the national government, the United States, and the United Nations (PAP-H3 through H5). The paper instead tests three hypotheses, all concerning hosting attitudes. This narrowing reflects the paper's more specific focus on attitudes toward refugee hosting. The economic and security outcomes show how the cuts

[1] "Isiolo Governor Abdi Guyo Impeached," *Daily Nation*, June 27, 2025, https://nation.africa/kenya/counties/isiolo/isolo-governor-abdi-guyo-impeached--5096528; "How Isiolo Governor Guyo Survived Impeachment on a Technicality," *Daily Nation*, July 9, 2025, https://nation.africa/kenya/news/politics/how-isiolo-governor-guyo-survived-impeachment-on-a-technicality-5111156.

Table 37: Pre-registered two-town difference-in-differences (transparency only)

Outcome	Treated x W2	SE (W2)	Treated x W3	SE (W3)	N
Hosting support	-0.182	(0.198)	0.180	(0.220)	1,158
Effects index	0.058	(0.063)	-0.117†	(0.067)	1,169
Integration index	-0.088	(0.080)	0.143	(0.092)	1,169
Oppose camp closure	-0.249	(0.217)	-0.532*	(0.230)	1,161
Full rights	0.031	(0.044)	0.098†	(0.050)	1,170
Feels unsafe at night	-0.119*	(0.058)	-0.127*	(0.062)	1,166
Crime victim, past month	0.092†	(0.055)	0.016	(0.060)	1,170

Note: Balanced two-town panel (205 Kakuma and 185 Isiolo respondents observed in all three waves). Respondent fixed effects with wave indicators (Wave 1 omitted); coefficients are the differential change in Kakuma relative to Isiolo at each post wave. Standard errors clustered by respondent. The camp-closure outcome is reversed so higher values indicate opposition to closing the camps. Reported for transparency only; the control town was contaminated during the study period (see text), so estimates should not be interpreted as effects of the aid cuts. Significance: *** p<.001, ** p<.01, * p<.05, † p<.10.

affected host households. I now treat these results as descriptive because they establish the context for the paper's main question rather than test its central argument. Section 3.2 and this appendix report all preregistered economic and security outcomes. Attitudes toward the Kenyan government, the United States, and the United Nations address the different question of how citizens assign responsibility for the cuts. I examine those outcomes in a companion working paper on attribution and accountability. This current paper therefore focuses on PAP-H2, attitudes toward refugee hosting. I divide that hypothesis into H1 and H2 to distinguish community-level exposure from individual-level exposure.

5.6.0.4 The empathy channel (H3) is exploratory. The empathy channel, the formal framework in Appendix 5.5, and the Aid × empathy vignette interaction were not pre-registered. H3 emerged inductively from the open-ended panel responses and refugee-leader interviews, in which respondents repeatedly framed their views in humanitarian terms. Because this pattern appeared prominently across the qualitative evidence, I concluded that the paper's theoretical framework would be incomplete without accounting for it. I therefore incorporated empathy as an additional mechanism and developed H3 to assess whether it was also visible in the quantitative results. These analyses should nonetheless be read as exploratory rather than confirmatory. A related limitation, discussed in the main text, is that the empathy moderator was measured post-treatment at Wave 3.

5.6.0.5 Within-treatment (H2) specifications. The baseline-beneficiary intent-to-treat design, which fixes exposure at Wave 1 aid-recipient or aid-employment status, follows the PAP's pre-specified within-treatment design. The cohort difference-in-differences and the switching difference-in-differences reported in Section 3.6.2 were added after fieldwork as robustness checks and were not pre-specified.

5.6.0.6 Attrition and follow-up survey. The PAP specified a target of 300 respondents per town in each wave, allowing for expected attrition of 15–20%. Actual attrition was higher. The balanced three-wave Kakuma panel includes 205 respondents, while the refreshed cross-sections includes at least 300 respondents per wave. Section 1.3.2 discusses attrition in detail. The fourth phone wave was implemented as specified in the January 2026 amendment, but attrition was substantial after the switch in survey mode (results are reported in Section 3.4). In the 2026 nationwide survey, a technical issue affected the *Nation Hosting* question and required a May 2026 recontact with a nationally representative sample of 1,776 respondents. Sections 1.3.1 and 1.1 provide further details.

References

Alix-Garcia, Jennifer, Sarah Walker, Anne Bartlett, Harun Onder, and Apurva Sanghi. 2018. "Do Refugee Camps Help or Hurt Hosts? The Case of Kakuma, Kenya." *Journal of Development Economics* 130: 66–83.

Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103(484): 1481–1495.

Baseler, Travis, Thomas Ginn, Robert Hakiza, Helidah Ogude-Chambert, and Olivia Woldemikael. 2025. "Can Redistribution Change Policy Views? Aid and Attitudes toward Refugees." *Journal of Political Economy* 133(9): 2874–2917.

Demir, Cansu, Anna Maria Mayda, and Jean-François Maystadt. 2025. "The Political Impact of Refugees in Africa." Working paper.

Hainmueller, Jens. 2012. "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies." *Political Analysis* 20(1): 25–46.

Kadigo, Mark Marvin, and Jean-Francois Maystadt. 2023. "How to Cope with a Refugee Population? Evidence from Uganda." *World Development* 169: 106293.

MacDonald, Mae, and Adam Lichtenheld. 2025. "Support for Refugee Integration in a Major Refugee-Hosting Country." Working Paper No. 25-01, Immigration Policy Lab, Stanford University.

Zhou, Yang-Yang, Guy Grossman, and Shuning Ge. 2023. "Inclusive Refugee-Hosting Can Improve Local Development and Prevent Public Backlash." *World Development* 166: 106203.

Zhou, Yang-Yang, Naijia Liu, Shuning Ge, and Guy Grossman. 2025. "Liberalizing Refugee Hosting Policies without Losing the Vote." OSF Preprints, Version 4.