

FEAR TRAPS

Persistent Avoidance Despite Disconfirming Evidence

Miguel Ortiz*

University of British Columbia

Eunice Atajiri

Illimi Solutions

September, 2026

Abstract

We document a fear trap that persisted in the face of both disconfirming evidence and incentives to break it, and we break it experimentally. Jos, Nigeria, has a history of deadly inter-religious riots, and Christians and Muslims there live in segregated neighborhoods and shop in segregated markets. Christian households avoid Muslim-run markets that are on average 27% cheaper, despite being aware of the price gap, because they fear that they will be attacked while they are in a Muslim area — even though a visible minority of Christians shops there regularly and returns unharmed. We randomize 500 non-goers into four weekly transfers redeemable only at a large Muslim market, or an unconditional transfer of equal value. Visits move safety beliefs sharply, with the share naming safety as a problem falling from 97% to 14%, while beliefs about prices, already accurate, do not move. The change diffuses: 44% of control households visit on their own during the intervention. Two months after incentives stop, 93% of treated and 83% of control households shop at the Muslim market, and attitudes and behavior toward Muslims improve. We interpret these findings through a model in which fear traps persist because safe returns are uninformative when observers cannot tell whether visitors are safe because the market is safe or because the visitors are an unusual type; randomization removes that ambiguity.

*Corresponding author: miguel.ortiz@ubc.ca. We are grateful to Christopher Blattman, Ernesto Dal Bo, Pascaline Dupas, Thomas Fujiwara, Seema Jayachandran, Nathan Nunn, and Jonathan Weigel for their helpful comments. Luke Chan provided excellent research assistance. We gratefully acknowledge the funding provided by the International Growth Centre and the Weiss Fund. IRB approval was granted by UC Berkeley. The pre-analysis plan of this paper is available at [\[link here\]](#)

1 Introduction

Individuals sometimes avoid actions that would make them better off because they believe those actions are dangerous when, in fact, they are not. Such mistaken fears can be persistent, because the evidence that would correct them is available only to those who take the very action the fear discourages, creating an individual-level “fear trap.” At the societal level, however, this trap should be broken. When individuals differ in their willingness to take risks, the less fearful may undertake potentially beneficial but uncertain actions, realizing the gains when they succeed and bearing the costs when they fail. Their experiences then provide fearful individuals with second-hand information, allowing them to update their beliefs and avert the trap without having to bear the risk themselves. Yet this mechanism depends critically on whether individuals learn from the experiences of entrepreneurs. If fearful individuals systematically fail to update their beliefs in response to the successes of those who take the risk, individual fear traps can persist and become a societal one.

Why do societal fear traps exist even in the face of incentives and evidence to break them, and what would it take to unravel one? This paper documents a costly fear trap that persists although the people caught in it are aware of the large economic benefits of the action they avoid and observe entrepreneurs returning from it safely. It then shows that a brief, light-touch intervention breaks the trap, setting off a self-sustaining process of adoption that spreads to people the intervention never reached.

Our setting is the city of Jos, in central Nigeria, where a history of inter-religious violence has left Christians and Muslims living in largely segregated neighborhoods and shopping in largely segregated local markets. Market segregation is economically costly due to strong group-based specialization: Muslim traders dominate wholesale and retail and operate markets with substantially lower prices—on average 27% cheaper than Christian markets—along with somewhat higher quality and greater variety. Christians are well aware of these gaps. In our surveys, 96% report knowing that Muslim markets are cheaper, and they hold accurate beliefs about the size of the price difference. In addition, despite religious tensions, they also trust and rate Muslim sellers as highly as Christian sellers. Despite this, they do not shop in Muslim markets, and have not done so for 8.6 years on average. When asked why, 97% say they feel it is unsafe: they fear being attacked while there, or an outbreak of religious violent riots starting while they are in a Muslim area (which have been common in Jos’s past). As a result, they buy the same goods at higher prices in Christian markets to avoid a place they believe is dangerous.

This fear, however, is largely unwarranted. A small number of Christian households do go to Muslim markets, and they report that going is safe, which our extensive fieldwork confirmed. What makes the persistence of fear more puzzling still is that every Christian in our sample knew that some Christians regularly visited the Muslim market. Anecdotal evidence suggests they did not see these visitors as people like themselves, but dismissed them as people who presumably had special relationships, arrangements, or knowledge that shielded them from risk. Our fieldwork, however, suggest that many were ordinary households. So, despite the economic incentive and the evidence of visitors returning safely, fear still

prevails among most Christians and prevents them from taking the action that would correct their beliefs and leave them better off.

To understand whether this fear trap can be broken, and why it holds, we ran a field experiment that incentivized Christians to visit a Muslim market and experience it firsthand. We sampled 500 Christians who were not visiting any Muslim market and randomly assigned them to treatment and control. Treatment participants received 10,000 naira (roughly \$6.50) per week for four weeks, disbursed to them only at New Market, a large, mostly wholesale Muslim market located inside a Muslim area. To address safety concerns directly, visits were made in groups, accompanied by a member of the project team, and participants were informed that the project was coordinated with the market. Control participants received the same amount as an unconditional transfer, which held income constant across arms and left them free to keep shopping where they always had. We surveyed households at baseline; after each market visit during the intervention (treatment only); at an endline about five weeks after baseline, immediately after the intervention ended; and at a follow-up two months later. We examine how this exposure affects perceptions of market characteristics such as safety and prices, their market choice, and broader attitudes toward the outgroup.

The intervention performed better than we expected. Ninety-eight percent of the treatment group visited the Muslim market, and 43% made additional unincentivized and unaccompanied visits during the intervention period—an early sign that the experience itself, and not the transfer alone, was drawing them back. Exit surveys administered after each visit show that participants spent roughly twice what they were required to spend on each visit and stayed longer each week, reaching about 50 minutes per visit by the final week. On average they report saving 26% relative to what the same basket would have cost in their regular market. They visited around 4.5 sellers per visit, and by the end 55% were buying from a seller they had bought from before. By week 4, 95% of participants reported that sellers were fair, respectful, and friendly toward them, and 67% had chatted with a seller about something other than business.

The central result at endline is that four weeks of visits sharply changed beliefs about safety. Among treated participants, the share naming safety as a downside of shopping at New Market fell from 97% to 14%. What moved was not a single belief but every component of the fear. When asked on a Likert scale how likely they thought it was that something bad would happen to them if they went to New Market, treated participants revised their estimates downward by 1.8 SD, indicating that they had learned that the day-to-day risk of a visit is essentially nil. The share reporting that they feared a violent riot breaking out while they were at the market fell from 93% to 37%, indicating that the risk of extreme events, while real, had been overweighted. And asked how likely they would be to return home safely if a riot did break out, they answers dropped by 0.74 SD. This last shift is telling. Simply surviving a rare draw would leave a person's beliefs about the worst case untouched; that participants revised what they expect to happen conditional on the feared event suggests they learned something about the people at the market rather than merely got lucky. Beliefs about prices and variety of products, by contrast, did not move,

and beliefs about quality and commute times improved only slightly — consistent with the economic benefits having been well understood all along.

Remarkably, the effects of the intervention did not stay within the treatment group. During the intervention period, 44% of control participants—who received only an unconditional transfer and no encouragement to visit—went to New Market on their own. This diffusion followed a reasonable logic: control participants who visited were more central in the social network and closer geographically to treatment group participants, consistent with being more likely to have learned that the market was safe. When asked why they had started visiting, 97% gave reasons directly tied to the project: 13% said a treatment participant had told them about their experience, and 84% said no one had told them anything, but that they saw treated participants return from the market unharmed and this made them curious enough to try it themselves.

Two months after the intervention ended, the change had not faded; it had spread and stabilized. Ninety-three percent of treated participants were still shopping at New Market, and 83% of control participants had by then started going, all of them unincentivized and unaccompanied. Across our main outcomes, treatment and control converge, and they do so in a specific way: the treatment group holds steady at its endline level while the control group rises to meet it. We read this convergence as contagion rather than fade-out. Where the treatment–control gap shrinks by follow-up, it shrinks because control households are adopting the behavior, not because treated households are reverting. Relative savings at New Market remained stable, indicating that three months after the intervention began, Christian markets had not responded with price adjustments.

This intervention also put participants into repeated contact with Muslims, but of a particular kind: market exchange is cooperative, in that both sides gain from trade, but it is also adversarial in the moment, since every price is contested in negotiation in this context. Despite not being clear whether contact of this kind would improve how people view the outgroup, we find that market interactions improved attitudes and behavior toward Muslims. After the intervention, participants rated Muslim sellers more highly on trust and fairness, from a baseline already comparable to their ratings of Christian sellers, and viewed Muslims in general as more trustworthy and less socially distant. In incentivized behavioral games with Muslim partners who were not sellers, treated participants showed more altruism, trust, and cooperation. By follow-up, these gains had grown in the treatment group, and the control group had caught up.

The question that remains is why a fear trap existed at all. Why did mistaken beliefs about the market persist for years while a subset of Christians could be seen returning safely? And why did the fear trap break in the control group once the Christians returning safely were ones our intervention had induced to visit? In the last part of the paper, we present a model that rationalizes these facts with a mechanism through which fear traps can survive even in the face of disconfirming evidence and incentives to break them.

We begin with a benchmark that formalizes the standard economic logic: if going to the market is beneficial but potentially dangerous, risk-taking entrepreneurs go and succeed, and others observe, learn, and then go themselves. We then introduce one change to the model, motivated by how non-goers described the Christians who went to New Market: as people who were presumably not at risk because they had special relationships or knowledge. In the model, there are two types of people: those who are exposed to danger if they visit and those who are shielded from it. A shielded type always visits. An exposed type visits only when her fundamentals make the risk worth bearing; these are the entrepreneurs. Now consider an exposed non-goer who is unsure whether the pool of goers consists only of shielded types or of a mix of shielded and exposed types. Observing Christians visit the market and return safely, even arbitrarily many times, will not let her conclude that the market is safe, as it would in the benchmark. Safe returns can never fully distinguish a safe market from a pool of goers who are all shielded. The existence of exposure types thus creates an ambiguity that bounds learning, so that fear cannot be eliminated even when exposed entrepreneurs generate disconfirming evidence. The fear trap survives.

The intervention breaks the fear trap not by adding to the stock of evidence but by removing its ambiguity. Its participants were non-goers selected by lottery, so when they returned safely, it was implausible to attribute their safety to them being the kind of people who were never at risk. Once payments stopped and escorts ended, remaining alternative explanations were ruled out, and the signal became stronger still. Years of watching an ambiguous set of Christians could not settle the question of safety, whereas a few weeks of watching an unambiguous set did.

Importantly, the model involves no behavioral bias, the situation it describes is fully rational: given that the market appears dangerous, it is reasonable to suspect that goers might all be shielded types. No one need be mistaken about who goes for everyone to remain mistaken about the market. In addition, the model showcases an inversion of the behavioral literature on selection neglect, which finds that people hold mistaken beliefs because they fail to account for selection in the evidence they observe (Jehiel, 2018; Enke, 2020). Here the opposite holds: people are trapped in mistaken beliefs precisely because they recognize that goers might be a selected group.

We are careful, however, about how hard we push the story the model presents. We advance the type-discounting mechanism because it is the account most consistent with what we observe in the data and learned in the field, not because we can establish it against all alternatives.

Although we study a specific setting, the mechanism we study is general and the phenomenon does not depend on the particulars of Jos. A self-sustaining fear trap of this kind can arise wherever there is an action that is valuable but risky, and two conditions hold: some agents are shielded from risk while others are exposed, and an observer cannot tell whether the people who take the action are shielded or not. These conditions plausibly hold in many settings, including the adoption of new technologies, the founding of new ventures, approaching distrusted institutions, and many forms of interaction between identity groups. In each, evidence of safety can stand in plain view of the very people who avoid the

action and yet fails to reassure them, because the choice to act can be read as a sign of the actor’s type rather than of the action’s safety. The Jos market setting is valuable precisely because it makes such a trap measurable: the forgone gains are large and quantifiable, the fear is specific and nameable, and the population of non-goers can be screened and randomized. We offer these parallels as conjectures rather than findings, but we believe the underlying structure—costly avoidance sustained by treating others’ safe experience as the mark of an unusual type—is common.

This paper contributes to three literatures. The first is work on misperceptions and why they persist. People often hold mistaken beliefs about others (Bursztyn and Yang, 2022), including about outgroup hostility (Mernyk et al., 2022; Ortiz, 2025). Explanations for why such beliefs persist include pluralistic ignorance (Bursztyn et al., 2020), motivated beliefs (Bénabou and Tirole, 2016; Zimmermann, 2020), and selection neglect (Jehiel, 2018; Enke, 2020; Esponda and Vespa, 2018). None of these explains the belief we study: the evidence against it was public, the belief was costly rather than comforting, and neglecting selection would have corrected it. We propose an explanation that requires no behavioral bias: people who suspect that the evidence they observe comes from a selected group cannot learn enough from it to correct their beliefs about the risk of a profitable action. Mistaken beliefs then persist because people account for selection, not because they neglect it.

The second is work on exploration and the adoption of new economic choices. One strand focuses on geographic exploration, showing that it tends to be limited and that incentives can expand it (Bryan et al., 2014; Davis et al., 2019; Bergman et al., 2024; Dean et al., 2026), although the effects do not always last (Pelnik, 2025). Related work shows that brief incentivized exposure can lead to lasting adoption (Bryan et al., 2014; Dupas, 2014; Carter et al., 2021; Dean et al., 2026) and that new technologies spread through social networks (Munshi, 2004; Conley and Udry, 2010; BenYishay and Mobarak, 2019; Carter et al., 2021; Beaman et al., 2021). We show that whether exploration spreads can depend on who explores: people learned from explorers selected by lottery but not from those who had chosen to go. Because the barrier was a belief rather than an aversion to the unfamiliar (Dean et al., 2026) or a standing constraint (Pelnik, 2025), others’ experience could remove it, and adoption outlasted the incentive.

The third is work on markets and relations between identity groups. Whether markets make people more cooperative (the *doux commerce* hypothesis) or erode moral behavior has long been debated (Hirschman, 1982; Falk and Szech, 2013; Bartling et al., 2015). Trade is associated with less conflict between trading partners (Martin et al., 2008; Jha, 2013; Rohner et al., 2013), and market exposure with more prosocial, universalist, and civic values (Henrich et al., 2010; Enke, 2023; Agneman and Chevrot-Bianco, 2023; Rustagi, 2024). In these studies, however, market exposure is long-standing and not randomly assigned. We randomize incentives for market interaction between two estranged groups and find support for the *doux commerce* hypothesis: within weeks, the interaction increased altruism, trust, and cooperation toward the outgroup, even though participants bargained on nearly every visit. Unlike in experiments on intergroup contact, where researchers create the contact (Paluck et al.,

2019; Scacco and Warren, 2018; Mousa, 2020; Lowe, 2021), the interaction here continued without the program because trade was profitable, and its effects extended beyond the people participants traded with.

The rest of the paper proceeds as follows. Section 2 describes the setting: the history of violence in Jos, the segregation of its markets, and New Market. Section 3 presents the experimental design and the data. Section 4 documents the fear trap at baseline: participants knew that New Market was cheaper and trusted its sellers, but stayed away because they believed it was dangerous. Section 5 reports the results, first on what participants did during the four weeks of visits, then on their beliefs, attitudes, and behavior at endline, including what drove the spillovers to the control group, and finally on where both groups stood two months after the intervention ended. Section 6 presents the model, and the final section concludes.

2 Setting

Jos is one of the main cities in Nigeria’s Middle Belt, the region where the country’s Muslim north and Christian south meet. It has around one million residents, divided close to evenly between Christians and Muslims. Since 2001 Jos has experienced recurrent outbreaks of intergroup violence, known locally as *crises*. The major episodes came in 2001, 2008, and twice in 2010, and estimates of their combined toll range from roughly 4,000 to 7,000 deaths (Krause, 2011; Higazi, 2011). The crises began without warning, touched off by ordinary local disputes (Human Rights Watch, 2001; Krause, 2011), and the fighting was carried out by ordinary residents rather than organized armed groups — neighborhood youths, market crowds, and mobs with homemade weapons (Human Rights Watch, 2001; Krause, 2011). After them, lower-level violence continued sporadically, including killings of individuals found in the wrong neighborhood and attacks at roadblocks.

The crises scarred the city and reorganized it, setting in motion a process of religious segregation. Mixed neighborhoods sorted along religious lines through displacement and relocation, and the city’s neighborhoods are now largely identifiable as Christian or Muslim (Krause, 2011; Madueke, 2018). Figure ?? maps them, along with the city’s principal markets. Christians and Muslims rarely enter the other group’s areas, believing that being identified there could be dangerous. They fear that violent members of the other group might attack them, or that a new crisis might break out while they are somewhere they cannot easily leave.

Among the things that segregated were the markets where households buy their groceries, which had been integrated before the crises (Obateru and Adama, 2011; Lalo and Sadiq, 2011). This separation is particularly costly because the two groups specialize in different economic activities. Muslim traders built the long-distance networks that supply the region (Krause, 2011), so goods enter Jos through the Muslim markets, which operate largely as wholesale markets, and reach the Christian markets downstream, where they are sold at retail, at prices that include the margins added along the way. Muslim

markets are therefore substantially cheaper than Christian markets. Over three months we priced the eleven products Christian households buy most often and found that prices at Muslim markets are on average 27% lower.¹²

Christian households know about this price difference. Muslim markets are also not much farther from them than the Christian markets they use, and they trust Muslim sellers, who have a long trading tradition, no less than they trust Christian sellers (Section [X]). They nonetheless stay away, because they are afraid of the areas the Muslim markets sit in, and so pay higher prices to feel safe (Lalo and Sadiq, 2011). However, that fear is currently not warranted: visiting a Muslim market today is safe. A small number of Christian households have shopped at Muslim markets regularly for years and report that they are safe to go to — a claim we verified in several ways.³ In addition, most Christians know that these few go and return safely, but they appear to discount that experience as unrepresentative of what they themselves would face, saying that such people probably have special relationships and know how to move in a Muslim area. So despite a standing incentive to visit and save money, and despite the evidence of other Christians going and returning safely, most remain afraid and keep paying higher prices.

One of the main Muslim markets in Jos is New Market, among the largest in the city. It was created by Muslim traders who had held stalls in what was then the city’s main mixed market; when the violence made it impossible for them to remain there, they relocated to an area that was becoming Muslim and set up a new market. It sells both wholesale and retail, and it is one of the main places where retailers from across the city buy their stock. Christian areas surround it, none of them far away. This is the market where we run our intervention to correct mistaken beliefs.

3 Experimental Design and Data

This section presents the design of the experiment. We incentivized a randomly selected group of Christian households to visit a Muslim market over four weeks, and measured how their beliefs about the market, their market choices, and their attitudes and behavior toward Muslims changed as a result.

Sample and randomization

We recruited participants door to door in three Christian neighborhoods of Jos: Angwan Rukuba, Apata, and Tina Junction. These are the Christian neighborhoods closest to New Market, and so the ones for which shopping there is a realistic alternative to the markets residents currently use. Within each

¹A gap of this size is in line with Aker et al. (2014), who find that crossing an ethnic boundary within a West African country is about as costly as crossing an international border, and raises prices by roughly 25%.

²Appendix XX includes a methodological note on how the price difference between markets was calculated.

³We interviewed Christian households who shop there, spoke with Muslim market managers and with the police stations serving these markets, had our enumerators visit the markets themselves, and ran a pilot. All confirmed that visiting is safe.

household we surveyed the person responsible for buying the groceries, and allowed only one participant per household. The resulting sample is almost entirely women.

The eligibility criterion was not currently shopping at New Market. Around nine in ten of the households we approached were eligible. In line with the pre-analysis plan, we recruited 500 participants and assigned half to treatment and half to control, stratifying by neighborhood. At recruitment, participants were told that only some of them would be selected to visit New Market and that the selection would be random. They were also told that those not selected would still be paid for being part of the project and answering the surveys.

Treatment and control

Treatment participants were invited to visit New Market once a week for four weeks, and received 10,000 naira on each visit — about three days' wages, or 6.50 US dollars. They were told the schedule and the amount at the outset.

Participants traveled to New Market accompanied by an enumerator and in groups of eight. Each group met its enumerator in its own neighborhood on the same day and at the same time each week and traveled to the market together. The same enumerator accompanied the same group for all four weeks. Participants were also told that the project had been coordinated with the management of New Market. Payment was made in cash on arrival at the market, so a participant who did not travel did not receive it.

At the market, participants were required to spend at least 4,000 naira and to stay at least 30 minutes. They then returned to the point where the group had separated, where the enumerator checked their purchases and administered a short exit survey. After the survey they could wait and return home with the group or leave on their own.

Control participants received the same amount on the same weekly schedule, with no condition attached to visiting any market; to collect it, they had to meet their assigned enumerator once a week and answer a short survey.

Surveys

All surveys were administered in person at participants' homes. We conducted a baseline before the intervention, an endline about five weeks after baseline and immediately after the intervention ended, and a follow-up two months after endline.

Treatment participants also completed a five-minute exit survey after each of their four visits, administered at the market. The exit surveys recorded various aspects of the visit, including how much the participant had spent, how long she had stayed, and how many sellers she had visited. Control participants answered a short survey each time they collected their weekly payment, which asked which markets they had visited that week and how much they had spent.

Outcomes

We group outcomes into five families and describe each in more detail in Section 5, where the corresponding results appear.

Attendance. The number of visits a participant made to New Market. Rather than rely on self-reports, we gave participants at endline a punch card that a project partner stamped each time they visited New Market over the two months to follow-up. Each stamp increased their follow-up survey payment by an amount set below the cost of transport, so that the card would not by itself induce visits.

Beliefs about the Muslim market. The price difference between their regular market and New Market, both for a general consumer basket and for specific products; the perceived risk of everyday threats and of tail events; the quality and variety of products relative to their regular market; and the composition of the market’s customers.

Perceptions of Muslim sellers. Whether they are fair, honest, free of bias, respectful, and deliver on their promises. We also analyze a pre-registered index of these items.

Perceptions of non-seller Muslims. Whether a Muslim could be trusted as a neighbor, could become a good friend, would help in a moment of need, would be loyal and reliable, and would want the best for Christians. We also analyze a pre-registered index of these items.

Behavior toward non-seller Muslims. Four incentivized behavioral games: a dictator game, measuring *altruism*; a participant’s belief about her counterpart’s dictator decision, measuring *trust*; a two-player public-goods game, measuring *cooperation*; and a conditional public-goods game, measuring *reciprocity*. The games were one-shot, with real stakes of 2,000 naira and real matching. We never described the counterpart as Muslim. Instead we presented a list of names, all of them Muslim names, and told participants that their counterpart was one of the people on the list but not which one, to preserve anonymity. In this setting, names reveal religion.

Balance, attrition, and estimation

Attrition was 3 percent. A long array of baseline variables is balanced across arms; Appendix Table [X] reports the balance tests.

We estimate

$$Y_i = \alpha + \beta T_i + X_i' \gamma + \varepsilon_i, \quad (1)$$

where T_i indicates assignment to treatment and X_i is a vector of controls containing the baseline value of the outcome, age, and neighborhood fixed effects. The coefficient β is the intention-to-treat effect. Standard errors are heteroskedasticity-robust.

Because a substantial share of the control group also visited New Market, the intention-to-treat effect in equation (1) understates the effect of visiting the market. We therefore also report two-stage

least squares estimates in which visiting New Market is instrumented with assignment to treatment:

$$V_i = \delta + \pi T_i + X_i' \mu + \nu_i, \quad (2)$$

$$Y_i = \alpha + \beta V_i + X_i' \gamma + \varepsilon_i, \quad (3)$$

where V_i indicates whether participant i visited New Market and T_i is the excluded instrument. The coefficient β in equation (3) is the effect of visiting on Y among the participants whose visit was induced by their assignment. Appendix Table [X] reports the first stage.

The study received IRB approval from the University of California, Berkeley [LINK], and Appendix [X] contains a discussion of the ethical considerations raised by the design. The pre-analysis plan is registered with the AEA registry [LINK], and Appendix [X] lists the deviations from it.

4 Motivating Facts: A Costly Fear Trap at Baseline

This section describes our participants' situation before the intervention. Table 1 reports baseline means and standard deviations, and column (3) reports tests of balance between the treatment and control groups.

Participants are 40 years old on average, and 95% are women, a consequence of our surveying the person responsible for buying the household's groceries. They shop 2.3 times a week at a regular market 17 minutes from home. On average, their last visit to a Muslim market was 8.6 years before the baseline survey, and 11% had never visited one.

Participants *know* that the main benefit of New Market is that it is considerably cheaper. Asked to name its benefits, 96% name lower prices, and they estimate that goods cost 28% less there than at their regular market — a belief that turned out to be accurate. Fewer name variety (62%) or quality (39%), and the gaps they expect on these dimensions are small: 10% for variety and 5% for quality.

Why do they not go to New Market, then? Asked to name the downsides of New Market, 97% say it is unsafe. The other reasons we offered them get almost no traction: 3% say the products they need are missing, 3% that they cannot trust the sellers, and 2% that they do not want to give Muslims their business. Panel D shows the same about trust: participants rate Muslims as business partners as highly as they rate Christians, 3.74 against 3.72 on a scale from 1 to 5, which likely reflects the reputation Muslim traders have built in the region. Twenty percent say New Market is far, but they place it only five minutes farther from their homes than their regular market, and even that turns out to be an overstatement of about three minutes.

What are they afraid of, exactly? 93% say a crisis — one of the spontaneous outbreaks of intergroup violence described in Section 2. However, they do not fear only the extreme, unlikely event: 49% name an ordinary attack against them and 21% harassment. Only 3% name a terrorist attack, which indicates that what they have in mind is not an armed group entering the city but ordinary residents turning

Table 1: Situation at Baseline

	Mean (1)	SD (2)	Balance (p-v) (3)
<i>Panel A: Sample and shopping behavior</i>			
Age	40.08	10.31	0.55
Female (%)	0.95	0.21	1.00
Shopping trips per week	2.29	1.27	0.79
Travel time to regular market (min)	17.43	9.39	0.47
Travel time to New Market (min)	22.64	9.64	0.93
Years since last visit to a Muslim market	8.64	6.28	0.68
Never visited a Muslim market (%)	0.11	0.31	0.17
<i>Panel B: Beliefs about New Market</i>			
Benefits of New Market (% who selected each)			
Cheaper	0.96	0.20	0.82
Variety	0.62	0.48	0.71
Quality	0.39	0.49	0.93
Other	0.05	0.22	0.23
Perceived price gap (%diff in cost)	-0.28	0.13	0.87
Perceived quality gap (%diff in score)	0.05	0.47	0.66
Perceived variety gap (%diff in score)	0.10	0.56	0.27
Downsides of New Market (% who selected each)			
Unsafe	0.97	0.18	0.61
Far	0.20	0.40	0.32
Products missing	0.03	0.17	0.59
Can't trust sellers	0.03	0.16	0.78
Don't want to patronize them	0.02	0.15	1.00
Believed % of customers that are Christian	29.18	15.12	0.84
<i>Panel C: Beliefs about Safety at New Market</i>			
What do you fear could happen? (% who selected each)			
Violent crisis	0.93	0.26	0.73
Ordinary attack	0.49	0.50	0.48
Harassed	0.21	0.41	0.08
Terrorist attack	0.03	0.17	0.79
Likely something bad will happen to me (1-9)	6.34	1.68	0.69
Likely I would be at risk (1-9)	6.22	1.59	0.69
If crisis happens, likely I would not return home (1-9)	6.60	1.92	0.58
If all safe, how uncomfortable feels to visit (1-10)	6.44	2.02	0.81
<i>Panel D: Attitudes and behavior towards Muslims</i>			
Perception of Muslims for business rel. (avg. 1-5)	3.74	0.63	0.63
Perception of Christians for business rel. (avg. 1-5)	3.72	0.60	0.03
Perception of Muslims for personal rel. (avg. 1-5)	2.51	0.83	0.15
Altruism (dictator game, out of 2000)	552.82	317.64	0.35
Trust (beliefs on match's dictator, out of 2000)	627.66	344.66	0.27
Cooperation (public good game, out of 2000)	843.20	459.17	0.63
Reciprocity (conditional public good game, index)	0.94	0.24	0.37

Notes: This table reports summary statistics from the baseline survey administered to the sample of 500 households (250 treatment, 250 control) prior to random assignment. Columns (1) and (2) report the mean and standard deviation. Column (3) reports the p-value from a test of equality of means across treatment and control.

violent, as the crises in Jos have unfolded. On scales from 1 to 9, participants rate the likelihood that something bad would happen to them at New Market at 6.3, and the likelihood that they would be at risk there at 6.2. Asked to suppose that a crisis broke out while they were at New Market, they rate the likelihood that they would not return home safely at 6.6. Asked how uncomfortable they would feel visiting even if they went and returned safely, they answer 6.4 on a scale from 1 to 10.

None of this is because they believe that no Christian ever goes to New Market. On average, they think that 29% of its customers are Christian. However, our fieldwork suggests that they see those Christians as a different type of person: traders and others who have relationships and ways of moving in a Muslim area that protect them, and that a household like theirs does not have.

Despite these fears, our fieldwork established — and the experiment results confirm — that going to New Market was safe for ordinary Christian households. Why these unwarranted fears persisted for years, despite a standing economic incentive to overcome them and the evidence of other Christians going and returning safely, is the question we take up in Section 6.

Beyond the market, Panel D records how participants see Muslims outside of trade and how they behave toward them. They rate Muslims lower for personal relationships, at 2.51 on the same scale from 1 to 5. In the dictator game they give a Muslim counterpart 553 of 2,000 naira on average; they expect a Muslim counterpart to give them 628 in the same game; and in the public-goods game they contribute 843. Reciprocity is already high at baseline, as participants contribute on average 94% of their Muslim match contribution in a conditional public good game.

5 Results

This section presents the results of the intervention in three parts. The first describes what participants did during the four weeks of the intervention itself, drawing on the exit surveys administered after each visit to New Market. The second reports the effects of the intervention at endline, measured immediately after the last week of visits. The third reports the effects of the intervention at follow up, two months after the incentives had stopped.

5.1 Behavior During the Intervention

Table 2 describes what happened on each of the four visits, as recorded in the exit surveys that enumerators administered at the market immediately after participants finished shopping.

Attendance was near universal and held up throughout. Of the 250 participants assigned to treatment, 246 made the first visit and 239 the fourth.

Participants consistently did more than the visit required of them. They had to spend at least 4,000 naira and stay at least 30 minutes; in the first week they spent twice that minimum, just over 8,000 naira of the 10,000 they received, and stayed 39 minutes. Time at the market then grew: by the fourth week they were staying 50 minutes, 11 minutes longer than in the first ($p < 0.01$).

Table 2: Behavior During the Intervention

	Week 1	Week 2	Week 3	Week 4	$\Delta W4-W1$
	(1)	(2)	(3)	(4)	(5)
Number of visitors (out of 250 treated)	246	245	239	239	
Time spent (min)	39.13 (8.33)	47.95 (12.38)	50.30 (12.49)	50.49 (12.81)	11.37*** (0.93)
Money spent (naira)	8107 (1829)	8110 (1690)	8098 (1827)	8334 (1887)	227 (142)
Expenses reduced by (% of money spent)	0.27 (0.12)	0.25 (0.11)	0.26 (0.12)	0.26 (0.11)	-0.01 (0.01)
Safety score (1–10)	6.83 (1.98)	7.94 (1.59)	8.29 (1.47)	8.63 (1.28)	1.80*** (0.12)
Number of sellers visited	4.52 (1.57)	4.50 (1.68)	4.31 (1.76)	4.47 (1.77)	-0.05 (0.13)
Repeated seller (%)	0.07 (0.26)	0.33 (0.48)	0.47 (0.50)	0.55 (0.52)	0.47*** (0.03)
Haggled with seller (%)	0.91 (0.28)	0.94 (0.23)	0.92 (0.27)	0.91 (0.29)	-0.01 (0.02)
Sellers showed positive behavior (index, 0-1)	0.80 (0.24)	0.86 (0.22)	0.93 (0.19)	0.96 (0.12)	0.16*** (0.02)
Chatted with sellers about other things (%)	0.61 (0.49)	0.63 (0.48)	0.63 (0.48)	0.67 (0.47)	0.07* (0.03)
Chatted with Muslims other than sellers (%)	0.52 (0.50)	0.63 (0.48)	0.54 (0.50)	0.64 (0.48)	0.12*** (0.04)

Notes: This table reports mean weekly responses to the exit survey administered to treatment households after each visit to the New Market during the four-week intervention. Standard deviations are reported in parentheses in Columns (1)–(4); Column (5) reports the difference between the Week 4 and Week 1 means, with its standard error in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$).

The economics of the trip, by contrast, were settled from the start. Spending held between 8,100 and 8,300 naira across the four weeks, and participants reported that what they bought cost about 26% less than the same purchase would have cost at their regular market, a saving that lines up with the price gap we measured directly. Saving that amount didn't require any learning, as the percentage saved was consistent throughout. The same holds for how they shopped: they visited between four and five sellers on each trip and haggled virtually always.

What did change was how safe they felt and how they engaged with the market and the people in it. Asked at the end of each visit to rate how safe they had felt, participants gave 6.8 out of 10 in the first week — already a high score given the fears they had described at baseline — and 8.6 in the fourth, a further gain of 1.8 points ($p < 0.01$). Their dealings with sellers moved from anonymous to repeated: by the fourth week, 55% bought from a seller they had bought from before. How they felt sellers treated them was also high from the start and improved further, with the index of whether sellers were fair, honest, respectful, helpful, and friendly rising from 0.80 to 0.96 ($p < 0.01$). And contact spread beyond

the transaction. By the fourth week, 67% were talking with a seller about something other than business, and 64% were talking with a Muslim at the market who was not a seller.

These patterns are descriptive: they come from the treatment group alone and have no control counterpart. What they establish is that the intervention delivered the exposure it was designed to deliver. Participants went, they transacted on ordinary terms, they came to no harm, and over four weeks they grew more comfortable in a place they had avoided for years.

5.2 Effects of Visits at Endline

This subsection reports the effects of the intervention as measured at endline, administered immediately after the last week of visits. We look first at participants' beliefs about the market, then more specifically at their beliefs about safety there, and finally at how they see Muslims more broadly and how they behave toward them.

Visiting quickly outran the incentives. Ninety-eight percent of treatment participants visited New Market, and 43% made further visits beyond the four we incentivized, without a project enumerator accompanying them. Visits also spread to the control group: during the same period, 44% of control participants went to New Market on their own, despite not being incentivized to do so. Section 5.2.3 examines what drove these spillovers.

5.2.1 Beliefs about New Market

Table 3 reports how the intervention changed what participants believe about New Market. Panel A gives intention-to-treat estimates, the effect of being assigned to treatment. However, these estimates understate the effect of visiting, because assignment did not completely determine who went, as many in the control group did as well. Assignment raised the probability of visiting by 56 percentage points rather than by 100, and Appendix Table [X] reports this first stage. Panel B therefore instruments visiting with assignment, so its coefficients give the effect of visiting for those whose visits the assignment induced, and they run about 1.8 times the corresponding ITT estimates. Every regression table below reports both panels, and we read Panel B as our best estimate of the magnitude of the effect of visiting.

After four weeks of visits, what participants believe about the economic benefits of New Market barely moved. The null in Column 1 confirms that participants had accurate beliefs about the price gap between markets: four weeks of visits did not change their estimates, and the standard error of 0.01 rules out anything but the smallest revision. Their view of variety did not move either (Column 3). Only their assessment of quality improved, and it did so modestly: visiting raised the perceived quality advantage 12 percentage points (Column 2), an estimate significant at the 10% level. This suggests that whatever kept Christians away from New Market, it was not a mistaken belief about the economic benefits of going.

Table 3: Effects on Beliefs about New Market

	Benefits			Downsides		
	Price Gap (1)	Rel.Quality (2)	Rel.Variety (3)	Unsafe (4)	TranspTime (5)	%C.Buyers (6)
<i>Panel A: ITT</i>						
Treated	0.00 (0.01)	0.07* (0.04)	0.03 (0.04)	-0.43*** (0.04)	-1.79*** (0.62)	1.72 (1.30)
<i>Panel B: 2SLS</i>						
1(Visits > 0)	0.00 (0.02)	0.12* (0.06)	0.05 (0.07)	-0.77*** (0.07)	-3.22*** (1.10)	3.06 (2.32)
Mean DepVar	-0.28	0.05	0.10	0.97	22.64	29.18
Controls & FE	Y	Y	Y	Y	Y	Y
Observations	480	485	485	485	480	485

Notes: This table reports treatment effects on perceptions of New Market. Column (1) is the perceived % difference in price relative to respondents' usual market; Column (2) is the perceived % difference in quality; Column (3) is the perceived % difference in variety; Column (4) is an indicator for viewing the New Market as unsafe; Column (5) is the reported travel time to the New Market (minutes); Column (6) is the believed % of the New Market's customers who are Christian. Panel A reports ITT estimates of random assignment to treatment; Panel B reports 2SLS estimates instrumenting 1(Visits > 0) with treatment assignment. Standard errors are clustered at the individual level. Controls & FE includes controls for age, gender, education, and the outcome variable at baseline, and neighborhood fixed effects. Mean DepVar reports the baseline sample mean of each outcome; sample sizes vary slightly by column due to item non-response.

The one belief about the market that moved sharply was safety. At baseline, 97% of participants named safety as a downside of shopping at New Market; visiting cut that by 77 percentage points (Column 4). Among those the intervention induced to go, calling New Market unsafe went from near-universal to near-absent. This is the central change the intervention produced, and we take it apart in the next set of results.

Participants also came back with a more accurate sense of how far away the market is. They put the trip at 22.6 minutes at baseline, and visiting shortened that estimate by 3.2 minutes (Column 5), a revision that is small in absolute terms. Column 6 shows a different pattern: beliefs about who shops at New Market moved, but not differentially. Visiting shifted the estimated Christian share of the market's customers by 3 percentage points relative to control, which we cannot distinguish from zero. Across the full sample, however, the estimate rose from 29% at baseline to 37% at endline, probably because both groups could see the influx of Christian shoppers that the project itself had created.

Table 4 examines which parts of that fear moved. The table groups the questions by the kind of event they ask about: everyday events, the things that might happen to a shopper on a normal trip, and extreme events, the crises that have periodically broken out in the city. Within each group we asked whether participants were afraid the event would happen and, on a scale from 1 to 9, how likely they judged it to be. Column (7) holds the outcome fixed and asks how uncomfortable they would feel going even if nothing went wrong. All seven estimates show a statistically significant decrease in fear. As above, we quote the 2SLS estimates.

Table 4: Effects on Beliefs about Safety at New Market

	Everyday events				Extreme events		
	Afraid it happens (%)		Likelihood (1-9 scale)		Afraid it happens	Likelihood (1-9 scale)	
	Ordinary Attack (1)	Harassed (2)	Something Bad (3)	Safety at Risk (4)	Violent Crisis (5)	If crisis, no return safe (6)	If all safe, uncomfort (7)
<i>Panel A: ITT</i>							
Treated	-0.06** (0.03)	-0.06*** (0.02)	-1.68*** (0.14)	-1.64*** (0.15)	-0.18*** (0.04)	-0.79*** (0.18)	-1.49*** (0.16)
<i>Panel B: 2SLS</i>							
1(Visits > 0)	-0.11** (0.05)	-0.11*** (0.03)	-2.99*** (0.26)	-2.92*** (0.27)	-0.32*** (0.08)	-1.42*** (0.32)	-2.66*** (0.30)
Mean DepVar	0.49	0.21	6.34	6.22	0.93	6.60	6.44
Controls & FE	Y	Y	Y	Y	Y	Y	Y
Observations	485	485	485	485	485	485	485

Notes: This table reports treatment effects on perceived safety at New Market, with the columns grouped by the kind of event each question is about. Column (1) is an indicator for fearing an ordinary attack; Column (2) is an indicator for fearing being harassed; Column (3) is the likelihood (Likert scale, 1–9) that something bad would happen; Column (4) is the likelihood (1–9) that the respondent’s safety would be at risk; Column (5) is an indicator for fearing a violent crisis could happen while visiting; Column (6) is the likelihood (1–9) of not returning home safely in the event of a crisis; Column (7) is how uncomfortable (1–10) the respondent would feel visiting even if everything were safe. Panel A reports ITT estimates of random assignment to treatment; Panel B reports 2SLS estimates instrumenting 1(Visits > 0) with treatment assignment. Standard errors are clustered at the individual level. Controls & FE includes controls for age, gender, education, and the outcome variable at baseline, and neighborhood fixed effects. Mean DepVar reports the baseline sample mean of each outcome.

Beliefs about everyday risk fell sharply. The share afraid of an ordinary attack fell by 11 percentage points from a baseline of 49%, and the share afraid of harassment by 11 points from a baseline of 21% (Columns 1 and 2). The graded measures moved further still. Asked how likely it was that something bad would happen to them at New Market, participants answered 6.3 out of 9 at baseline, and visiting moved them down by 3.0 points, about 1.8 baseline standard deviations (Column 3). Their assessment of whether their own safety would be at risk moved almost identically, by 2.9 points from a baseline of 6.2 (Column 4). These are the largest effects in the table, and they are the ones four weeks of ordinary shopping trips are best placed to produce. Participants went, nothing happened, and they concluded that on a normal day the market is not dangerous.

Beliefs about extreme events also moved a substantial amount, although by less. At baseline, 93% of participants said they feared a violent crisis breaking out while they were at the market, and visiting cut that by 32 percentage points (Column 5). The likelihood they attached to not returning home safely if a crisis did break out fell by 1.4 points from a baseline of 6.6, or 0.74 baseline standard deviations (Column 6). That both moved less than the everyday measures is what we should expect, since four visits tell a participant far more about ordinary risk than about a rare event. But Column (6) is conditional on

the crisis happening, and four safe trips carry no direct information about it. That it moved at all means participants learned something about the people around them: they came away expecting those people to be more likely to help them than to harm them in an extreme situation.

The last column asks participants how uncomfortable they would feel while at the market, even if they went and returned safely. At baseline they answered 6.4 out of 10, and visiting cut that by 2.7 points, a drop of 1.3 baseline standard deviations. This drop likely reflects both greater ease with Muslims and less anxiety that something dangerous might happen.

One caveat applies to all of these magnitudes. Because 44% of the control group visited New Market during the intervention period, each estimate is a difference between two groups that both moved, and the levels fell by more than the coefficients alone suggest.

5.2.2 Beliefs about and behavior towards Muslims

The intervention put participants into repeated contact with Muslims, but of a particular kind: they went to buy things, and the Muslims they met were the people selling to them. Whether contact of this sort improves how a person sees an outgroup is not obvious a priori. Market exchange is cooperative, in that both sides gain from the trade, but it is also adversarial in the moment, since every naira of the price is contested in negotiation. We ask three questions in turn. Did participants change their view of the sellers they dealt with? Did anything carry over to Muslims they never met? And did any of it show up in behavior rather than only in what people say?

We begin with the sellers. Table 5 reports effects on five statements about Muslim sellers, each rated on a scale from 1 to 5: whether they are honest about the quality of their goods, set fair prices, keep their business promises, treat everyone the same in business dealings regardless of religion, and are respectful even to customers they dislike. Column (6) is a pre-registered standardized index of the five. Every item rises, and the index rises by 0.30 standard deviations, though the individual items are estimated less precisely than the index.

The baselines are as informative as the effects. Participants already rated Muslim sellers highly before they ever visited: 3.79 on honesty about quality, 3.84 on fair prices, 4.05 on keeping promises, and 4.11 on respectfulness, all well above the midpoint of the scale. As we noted at baseline, they rated Muslims as business partners no lower than Christians. The intervention therefore improved an assessment that was already favorable.

The exception is Column (4), and it is the revealing one. Asked whether Muslim sellers treat everyone the same in business dealings, regardless of religion, participants answered 2.91 at baseline, the only item below the midpoint and more than a full point below the others. This was the one commercial belief on which they harbored real doubt, and it is about religion rather than about competence or honesty. Visiting raised it by 0.31 points, moving participants toward the belief that they themselves would be dealt with on the same terms as a Muslim customer. Appendix Table [X] runs the same estimates

Table 5: Effects on Beliefs about Muslim Sellers in Business Relationships

	Honest ab. Quality (1)	Fair Prices (2)	Keep Promises (3)	Not Biased by Relig. (4)	Respectful (5)	Index (6)
<i>Panel A: ITT</i>						
Treated	0.19** (0.09)	0.15* (0.09)	0.12** (0.05)	0.17* (0.10)	0.18*** (0.06)	0.17*** (0.05)
<i>Panel B: 2SLS</i>						
1 (Visits > 0)	0.34** (0.15)	0.27* (0.16)	0.21** (0.09)	0.31* (0.17)	0.32*** (0.10)	0.30*** (0.08)
Mean DepVar	3.79	3.84	4.05	2.91	4.11	0.00
Controls & FE	Y	Y	Y	Y	Y	Y
Observations	485	485	485	485	485	485

Notes: This table reports treatment effects on perceptions of Muslim sellers in business relationships. Column (1) is agreement (Likert scale, 1–5) that Muslim sellers are honest about the quality of their products; Column (2) is agreement that they set fair prices; Column (3) is agreement that they keep their business promises; Column (4) is agreement that, in business, they treat everyone the same, regardless of religion; Column (5) is agreement that they are respectful, even if they don't like you; Column (6) is a standardized index of Columns (1)–(5) (mean zero by construction). Panel A reports ITT estimates of random assignment to treatment; Panel B reports 2SLS estimates instrumenting 1 (Visits > 0) with treatment assignment. Standard errors are clustered at the individual level. Controls & FE includes controls for age, gender, education, and the outcome variable at baseline, and neighborhood fixed effects. Mean DepVar reports the baseline sample mean of each outcome.

for Christian sellers, where almost nothing moves except that participants now judge them less fair on price.

It is important to remember that these gains did not come from a frictionless encounter. Participants haggled virtually every time, which is to say that most of their contact with these sellers was spent bargaining against them over price. They nonetheless came away rating them higher on every dimension we asked about. One limit is worth stating: these are beliefs about the particular Muslims participants transacted with. Whether the change extends to Muslims they did not meet is the question we take up next.

The second question is whether these effects carried beyond sellers in business relationships to Muslims in general in personal ones. Table 6 asks about Muslims we explicitly described to respondents as not being sellers: whether participants would trust one as a neighbor, could become good friends with one, would be helped by one in a moment of need, believe that Muslims want the best for Christians, would find them loyal whatever the circumstances, and could rely on them in a moment of crisis. Column (7) is a pre-registered standardized index of the six. The index rises by 0.27 standard deviations, close in size to the effect we found for sellers.

These beliefs start from a much lower base, however. Where participants rated Muslim sellers between 3.79 and 4.11 on the commercial items, they rate Muslims in general at 2.46 as neighbors, 2.45 on loyalty, 2.23 on wanting the best for Christians, and 1.62 on being someone to rely on in a crisis.

Table 6: Effects on Beliefs about Non-Seller Muslims in Personal Relationships

	Trust to Be Neighbor (1)	Could Be Friends (2)	Help When I Need It (3)	Wants Best for Christ. (4)	Loyal (5)	Reliable (6)	Index (7)
<i>Panel A: ITT</i>							
Treated	0.18* (0.09)	0.08 (0.09)	0.06 (0.08)	0.10 (0.08)	0.15* (0.09)	0.26*** (0.09)	0.15** (0.06)
<i>Panel B: 2SLS</i>							
$\mathbb{1}(\text{Visits} > 0)$	0.33** (0.17)	0.14 (0.15)	0.12 (0.14)	0.18 (0.15)	0.28* (0.15)	0.47*** (0.16)	0.27** (0.11)
Mean DepVar	2.46	3.15	3.13	2.23	2.45	1.62	0.00
Controls & FE	Y	Y	Y	Y	Y	Y	Y
Observations	485	485	485	485	485	485	485

Notes: This table reports treatment effects on perceptions of non-seller Muslims in personal relationships. Column (1) is agreement (Likert scale, 1–5) that the respondent would trust Muslims to be their neighbors; Column (2) is agreement that they could become good friends; Column (3) is agreement that they would help the respondent in a moment of need; Column (4) is agreement that they want the best for Christians; Column (5) is agreement that they would be loyal despite the circumstances; Column (6) is agreement that the respondent could rely on them in a moment of crisis; Column (7) is a standardized index of Columns (1)–(6) (mean zero by construction). Panel A reports ITT estimates of random assignment to treatment; Panel B reports 2SLS estimates instrumenting $\mathbb{1}(\text{Visits} > 0)$ with treatment assignment. Standard errors are clustered at the individual level. Controls & FE includes controls for age, gender, education, and the outcome variable at baseline, and neighborhood fixed effects. Mean DepVar reports the baseline sample mean of each outcome.

Participants seemed to draw a line between the two domains: Muslims were people you could do business with, but not people you could count on. The intervention improved a favorable assessment in business relationships and an unfavorable one in personal ones.

All six items move in the same direction, and three of them are statistically significant. Trust in a Muslim as a neighbor rises by 0.33 points (Column 1), loyalty by 0.28 (Column 5), and reliance in a moment of crisis by 0.47 (Column 6), the largest effect in the table and the item that began lowest. The neighbor result is worth noting on its own: in the crises, when neighborhoods were still mixed, neighbors turned on each other. The remaining three items — whether Muslims could become good friends (Column 2), whether they would help in a moment of need (Column 3), and whether they want the best for Christians (Column 4) — move in the same direction but are not distinguishable from zero.

The third question is whether any of this shows up in behavior rather than in what people say. At endline, participants played four one-shot incentivized games against a Muslim counterpart, with real stakes of 2,000 naira and real matching. We never described that counterpart as Muslim. We showed participants a list of names, all of them Muslim names (which in Jos is an unmistakable signal of religion), and told them their counterpart was one of the people on it. Table 7 reports the effects. Column (1) is the amount given in a dictator game, which we take as altruism; Column (2) is what the participant expected her counterpart to give her in the same game, which we take as trust, and which was itself incentivized, since a participant earned an additional payment if she guessed correctly; Column (3) is

Table 7: Effects on Behavior Towards Non-Seller Muslims

	Altruism (1)	Trust (2)	Cooperation (3)	Reciprocity (4)
<i>Panel A: ITT</i>				
Treated	83.32*** (24.46)	61.73** (26.38)	109.22*** (36.24)	0.01 (0.02)
<i>Panel B: 2SLS</i>				
$\mathbb{1}(\text{Visits} > 0)$	149.28*** (43.20)	109.93** (45.95)	194.92*** (63.85)	0.03 (0.03)
Mean DepVar	597.62	599.60	860.40	0.96
Controls & FE	Y	Y	Y	Y
Observations	485	484	485	485

Notes: This table reports treatment effects on incentivized behavioral games with non-seller Muslims. Column (1) is the amount given to a Muslim match in a dictator game, out of 2,000 Naira; Column (2) is the amount the respondent believed their Muslim match would give them in the same game, out of 2,000; Column (3) is the amount contributed to a public goods game with a Muslim match, out of 2,000; Column (4) is a measure of reciprocity from a conditional public goods game: respondents stated their contribution for each of five hypothetical contributions by their Muslim match, and Column (4) is the average, across the five cases, of the respondent's contribution as a share of their match's contribution. Panel A reports ITT estimates of random assignment to treatment; Panel B reports 2SLS estimates instrumenting $\mathbb{1}(\text{Visits} > 0)$ with treatment assignment. Standard errors are clustered at the individual level. Controls & FE includes controls for age, gender, education, and the outcome variable at baseline, and neighborhood fixed effects. Mean DepVar reports the baseline sample mean of each outcome.

the amount contributed in a two-player public-goods game, which we take as cooperation; and Column (4) is a measure of reciprocity from a conditional public-goods game.

Three of the four move, and the magnitudes are considerable. Visiting raised giving in the dictator game by 149 naira against a baseline of 598 (Column 1), what participants expected to receive by 110 against a baseline of 600 (Column 2), and public-goods contributions by 195 against a baseline of 860 (Column 3). Measured against baseline these are increases of a fifth to a quarter; measured against the 2,000-naira endowment they run between five and ten percentage points. Unlike a survey answer, each of them is money the participant does not keep.

Column (2) is worth noting separately, because it records a belief rather than an action. That it rises means participants came away expecting more generous treatment from Muslims, not only behaving more generously toward them. Reciprocity, in Column (4), is the one measure on which we find no effect, but the baseline suggests this reflects a ceiling rather than the absence of an effect. At 0.96, participants were already matching almost all of their counterpart's contribution, so the measure had nowhere left to go, and the null is uninformative rather than evidence that nothing changed.

5.2.3 What drives spillovers?

Forty-four percent of the control group visited New Market during the intervention period. We did not anticipate this, because our fieldwork had shown how hard it was to persuade anyone to visit a Muslim market at all. In this section we investigate what distinguished the control households who went from

Table 8: What Drives Spillovers?

<i>Outcome variable: Visits to New Market</i>							
Indep.Variable:	Beliefs at baseline				Network and proximity		
	Price Gap	Fear Index	Business Rel. Index	Personal Rel. Index	Degree Centrality	Eigenvector Centrality	# Treated In 25m
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	0.08 (0.58)	-0.05 (0.11)	0.09 (0.11)	-0.11 (0.10)	0.28*** (0.10)	2.02** (0.98)	0.30** (0.13)
Controls & FE	Y	Y	Y	Y	Y	Y	Y
Observations	241	241	241	241	241	241	241

Notes: This table reports correlates of visits to the New Market by control households, to explore potential channels of spillovers. The dependent variable is the same in every column; each column is a separate regression on the independent variable named in its header. Columns (1) to (4) are beliefs measured at baseline: Column (1) is the perceived price gap, as in Table 3; Column (2), the Fear Index, is a standardized index of the three likelihood measures in Table 4 (Columns 3, 4, and 6); Column (3) is the same Business Relationship Index as Column (6) of Table 5; Column (4) is the same Personal Relationship Index as Column (7) of Table 6. Columns (5) and (6) are the respondent's degree and eigenvector centrality in the social network of participants; Column (7) is the number of treated households located within a 25-meter radius of the control household, controlling for the total number of participant households in that radius. Standard errors are clustered at the individual level. Controls & FE includes controls for age, gender, and education, and neighborhood fixed effects.

those who did not, in order to understand what led them to go. We read the answers as suggestive. Table 8 restricts the sample to the control group and regresses the number of visits a household made on a series of baseline characteristics, one at a time. Columns (1) to (4) are the beliefs we have tracked throughout, measured before the intervention began. Columns (5) and (6) place the household in a social network we constructed among participants, using two standard measures of centrality. Column (7) counts the treated households living within 25 meters of the respondent, controlling for the total number of participant households in that radius.

Baseline beliefs do little to predict who went. The perceived price gap is essentially zero (Column 1). The fear index is negative and the business relationship index positive, both in the direction one would expect if the least fearful and the best disposed went first, but neither is distinguishable from zero (Columns 2 and 3), and the personal relationship index carries the opposite sign to the one we would have predicted (Column 4). With 241 households these estimates are not precise enough to rule out modest relationships, and we would not want to claim that predispositions were irrelevant. What we can say is that nothing in what control households believed beforehand separates the visitors from the rest.

On the other hand, a household's position in the social network does predict whether it visited. Both measures of centrality are positive and statistically significant, degree centrality at the 1% level (Column 5) and eigenvector centrality at the 5% level (Column 6), and the two are constructed differently. Column (7) adds a third measure of exposure to information, one that does not depend on anything participants told us: the number of treated households living within 25 meters. It is positive and significant as well. This matters because the network measures are built from participants' own accounts of whom they

know, and could in principle pick up how forthcoming a respondent was rather than how well connected she is. Physical proximity is not open to that objection, and it points the same way. Whatever moved control households to New Market traveled through the network and reached those best placed to hear about it first. Being disposed to go seems to have mattered less than being placed where news of the market would arrive.

One might object that these visits had nothing to do with the project. Christian households could have been returning to Muslim markets on their own, for reasons unrelated to anything we did, and our control group would then be part of a change that was already under way rather than a group responding to the treated households around them. The baseline data make this hard to believe. Control households had not visited a Muslim market for 8.6 years on average. Behavior that had not changed in almost nine years is unlikely to change on its own in 44% of the control group, in the same four weeks that we were paying their neighbors to shop at New Market.

Beyond this, we also asked the control households who had started going why they had done so. Ninety-seven percent gave an answer connected to the project: 13% said a treatment participant had told them about her experience, and 84% said no one had told them anything, but that they had seen treated participants come back unharmed. Self-reported reasons have limited worth on their own, but this one appears to track something real. Control households who said a treatment participant had told them were four times more likely to have a treatment participant living close by than those who said they had only observed others returning safely. The channel people named lines up with the channel we can observe.


5.3 Situation at Follow-up (Two Months After)

The follow-up survey took place two months after the intervention was over, and by then the control group had become a group of New Market shoppers in its own right. This changes what the comparison can tell us, as a control group that has adopted the behavior is no longer a useful counterfactual. Therefore, everything in this subsection is descriptive. The following graphs show raw means of different variables at baseline, endline (the week after the intervention finished), and at follow up two months after the intervention concluded.

Visits over these two months were verified rather than self-reported. At endline we gave each participant a punch card, and a project partner at New Market stamped it on every visit she made. Each stamp raised the payment she received at the follow-up survey, but by less than the cost of transport to the market, so that collecting stamps was never worth a trip on its own.

Panel A of Figure 1 plots the share of each group that had visited New Market. As we have mentioned before, at endline essentially every treated participant had gone, against 44% of the control group. During the two months that followed the intervention, the share of the treatment group that visited New

Figure 1: Visits to New Market by Follow-up



Notes: This figure plots visits to the New Market by treatment status across three survey rounds: baseline (prior to random assignment), endline (at the end of the four-week intervention), and a later follow-up survey two months after the intervention ended. Panel A plots the percentage of respondents who had visited the New Market at least once; Panel B plots the average number of visits per week. Vertical bars are 95% confidence intervals.

Market, now without incentives or a project companion, remained high at 93%. The share of the control group, instead, kept rising, reaching 83% by follow-up.

Panel B plots the average number of visits per week and shows the same movement. By endline the treatment group was visiting roughly once a week. By follow-up, treated participants were visiting more often, 1.4 times a week. The control group was at 0.15 visits a week at endline, since most were not going at all, and by follow-up its frequency had caught up with the treatment group, at 1.25 visits a week.


The important thing to note is that in both panels the gap between the groups closes because the control group rises, not because the treated group falls away. This pattern, the persistence of visits in the treatment group and the almost complete spread to the control group, suggests that a trap of mistaken fears had indeed been holding these communities back from real economic gains, and that our simple intervention broke it by giving first-hand experience to some and then letting information flow.

Figure 2 follows perceived risk over the same three rounds. Panel A plots the share who name safety as a downside of shopping at New Market; Panels B and C plot the rated likelihood, on a scale from 1 to 9, that something bad would happen and that the respondent's own safety would be at risk; Panel D plots the rated likelihood of not returning home safely if a crisis broke out.

For the treatment group, little further learning seems to happen after endline. In Panel A the share calling New Market unsafe falls from 97% to 14% during the four weeks of visits, and then only to 11% over the two months that follow. In Panel B the rating falls from 6.4 to 3.2 and then sits at 3.3; in Panel C it falls from 6.2 to 3.3 and stays at 3.3. Treated households kept going, and went more often than before, but two further months at the market did not move their assessment of the risk much further.

Panel D behaves somewhat differently. The belief that one would not return home safely if a crisis broke out moves least of the four, although still by 0.74 standard deviations. Among treated households

Figure 2: Beliefs about Safety at New Market




Notes: This figure plots beliefs about safety at the New Market by treatment status across three survey rounds: baseline (prior to random assignment), endline (at the end of the four-week intervention), and a later follow-up survey two months after the intervention ended. Panel A plots the percentage of respondents who reported the New Market as unsafe. Panels B–D plot mean responses (Likert scale, 1–9) to: how likely something bad would happen to the respondent (B); how likely the respondent’s safety would be at risk (C); and, in the event of a crisis, how likely the respondent would not return home safely (D). Vertical bars are 95% confidence intervals.

it falls from 6.7 to 5.1 by endline and stands at 5.6 at follow-up, a difference that is not statistically significant, so the belief is best read as having held roughly steady. If some of it did revert, it is worth remembering that this is the belief on which visits carry the least information, since a small number of uneventful trips carries less information about the likelihood of extreme events.

In terms of fear, the control group repeats the catching-up pattern. Its beliefs had already moved considerably by endline, driven by the 44% who had visited by then, but they remained well above those of the treatment group, because most of the control group still had not gone. Then, as time went by and more of them took up going, their fear kept falling until it matched that of the treatment group. The share calling the market unsafe falls from 96% to 57% and then to 15%, against 11% among the treated. On the likelihood questions the ratings go from 6.3 to 4.9 to 3.6 in Panel B, and from 6.2 to 5.0 to 3.6 in Panel C, in both cases ending within a third of a point of the treatment group.

Figure 3 follows the four incentivized games over the same three rounds. Panels A and B plot the amount given to a Muslim counterpart in the dictator game and the amount the participant expected

Figure 3: Behavior Towards Non-Seller Muslims



Notes: This figure plots incentivized behavioral game outcomes towards non-seller Muslims by treatment status across three survey rounds: baseline (prior to random assignment), endline (at the end of the four-week intervention), and a later follow-up survey two months after the intervention ended. Panel A plots the amount given to a Muslim match in a dictator game, out of 2,000 Naira; Panel B plots the amount the respondent believed their Muslim match would give them in the same game, out of 2,000; Panel C plots the amount contributed to a public goods game with a Muslim match, out of 2,000; Panel D plots a measure of reciprocity from a conditional public goods game: the average, across five hypothetical contributions by the match, of the respondent's contribution as a share of their match's contribution. Vertical bars are 95% confidence intervals.

to receive from him. Panels C and D plot contributions in the public-goods game and the reciprocity measure.

Altruism and trust toward Muslims keep improving after the intervention ends. Among treated households, giving rises from 607 naira at baseline to 697 at endline and then to 715 at follow-up, and expected giving rises from 609 to 674 and then to 718. This is a different pattern from the one in Figure 2, where beliefs about risk settled once the four weeks of visits were over; how participants treated Muslims in these games went on changing as they kept going to the market. The control group shows the same catching-up as before. It moves little by endline, from 588 to 609 in Panel A and from 590 to 612 in Panel B, and then rises sharply to 705 and 700 at follow-up, within about ten naira of the treated.

Panel C is harder to read from the raw means. Treated households start below the control group at baseline, at 830 naira against 891, rise well above it by endline, at 977 against 891, and then fall back to 903 while the control group drifts down to 872. Measured against both the control group and their own

starting point, treated households are about 90 naira higher at follow-up than at baseline, so the direction of the change survives even though the endline peak does not. Panel D shows no separation between the arms at any point. Both begin at about 0.96, close to the maximum reciprocity, since participants were already returning nearly everything their counterpart had contributed, and both drift down together to 0.89 and 0.87 by follow-up.

6 A Model of Fear Traps in Segregated Markets

At this point in the paper, an important question lingers: Why did mistaken beliefs about the dangers of the market persisted for so long? Standard economic reasoning suggests that if an action brings large economic benefits but is believed to be risky, a mistaken belief about it should not last. People differ in how afraid they are and in how much they stand to gain, so some will find the gain worth the risk and take it. These are the entrepreneurs. If the risk is exaggerated, entrepreneurs succeed and profit, others observe them, learn, and copy them, and the mistaken belief does not last (Bolton and Harris, 1999; Foster and Rosenzweig, 1995; Bandiera and Rasul, 2006; Conley and Udry, 2010). In Jos, the gains were clear, and entrepreneurs existed and were observed: a minority of Christians had been going to New Market for years, and every participant in our sample knew of them. So why did learning not take place? Adding to the puzzle is what happened once the intervention began. Our treatment participants were, in a sense, a new set of entrepreneurs, and this time, within a few weeks, the control group observed them, updated their beliefs, and started going themselves. Why was the evidence from these entrepreneurs different?

The model in this section proposes an answer, which we believe is a novel mechanism by which mistaken beliefs can survive even in the face of incentives to abandon them and evidence that contradicts them. We begin with a benchmark that formalizes the standard logic: entrepreneurs appear, and others observe, update, and adopt—mistaken beliefs do not survive. We then introduce one change to the model, motivated by the way non-goers described the Christians who went to New Market, as people who presumably had special relationships, arrangements, or knowledge: there are two types of Christians, those who would be exposed to danger if they went to the market and those who are shielded from it, and an observer cannot tell them apart. The model shows how the existence of types bounds what observers can learn, allowing mistaken beliefs to survive indefinitely even though entrepreneurs exist and keep producing disconfirming evidence. The section concludes by showing how the intervention removed the ambiguity created by unobserved types, and with it the bound on learning that had kept mistaken beliefs in place.

6.1 Environment

Time is discrete, $t = 0, 1, 2, \dots$, indexing weeks. A unit continuum of Christian households, indexed by $i \in [0, 1]$, chooses each week which market to visit for groceries, $d_{it} \in \{C, M\}$, where C is the Christian market and M is the Muslim market.

Payoffs. The flow utility of visiting M relative to C is

$$\Delta u_{it} = v_i + m_{it} - \hat{p}_{it} L, \quad v_i \equiv \pi_i^{\text{sav}} + q_i - \tau_i - a_i, \quad (4)$$

where π_i^{sav} is the weekly expenditure saving from M 's lower prices, q_i the quality and variety differential, τ_i the differential travel cost, a_i the animus cost of patronizing outgroup sellers, $m_{it} \in \{0, m\}$ a program payment (zero outside the intervention), $\hat{p}_{it}$ household i 's subjective probability of coming to harm during a visit to M in week t , and $L > 0$ the composite loss from such harm. We call v_i household i 's *fundamentals*.

Risk. The uncertain object is $\rho \in \{0, \bar{\rho}\}$, the per-visit probability that a Christian visiting M comes to harm. Beliefs are $\mu_{it} \equiv \mathbb{P}_i(\rho = \bar{\rho})$, so subjective risk is $\hat{p}_{it} = \mu_{it} \bar{\rho}$. The true state throughout is $\rho = 0$: the danger, real in the past, has ended. Households, however, do not know this.

Heterogeneity and priors. Initial beliefs $\mu_{i0} \in (0, 1)$ are anchored on realized past violence. Fundamentals v_i are likewise heterogeneous, and the joint distribution of (μ_{i0}, v_i) has full support.

Network and information. Christian households are all connected. Each week, household i observes, for each j , the current market choice d_{jt} and, if $d_{jt} = M$, the outcome $o_{jt} \in \{\text{safe}, \text{harm}\}$.

Decision rule. A household visits M in week t iff $\mu_{it} \bar{\rho} L \leq v_i + m_{it}$, that is, iff its fearful beliefs lie below the cutoff

$$\mu_{it} \leq \bar{\mu}_i(m_{it}) \equiv \frac{v_i + m_{it}}{\bar{\rho} L}, \quad \bar{\mu}_i \equiv \bar{\mu}_i(0) \quad (5)$$

Equation (5) shows multiple things. A larger price gap raises v_i and makes a household more likely to overcome their fear of visiting. However, the possible harm being greater (bigger L) or the possible harm being more likely to happen if the market is in fact dangerous (greater $\bar{\rho}$) make people less likely to overcome their fear of visiting.

6.2 Benchmark: Learning when there are no types

This benchmark illustrates the logic by which the market should fix mistaken beliefs and make impossible a fear trap. Showcasing it will allow us to better understand afterwards why it fails.

Suppose—as we will not, from Section 6.3 onward—that all households face the same risk ρ and that this is common knowledge. Updating is Bayesian. Write $\mathcal{O}(x) \equiv x/(1-x)$ for odds; with two states, Bayes’ rule takes a multiplicative form, so posterior odds equal prior odds times the likelihood ratio of the signal.

A safe visit is possible in both states but more likely in the safe one: it occurs with probability $1 - \bar{\rho}$ if $\rho = \bar{\rho}$ and with probability 1 if $\rho = 0$. Since all households are alike, this is true of one’s own visit and of an observed visit equally, and outcomes are independent across visits and visitors conditional on the state. Hence every safe visit a household observes or takes multiplies its odds by $\text{LR} = 1 - \bar{\rho}$, so a record of k safe visits, whoever made them, multiplies the odds by $(1 - \bar{\rho})^k$, which converges to zero. Second-hand evidence is exactly as informative as first-hand evidence, and its informativeness is *unbounded*: enough of it settles the question.

Proposition 1 (No trap without types). *In the benchmark environment, for any distribution of priors:*

- (i) **The economic incentive generates entrepreneurs.** *By full support and (5), a positive mass of households—those with low enough initial fear (μ_{i0}) or high enough fundamentals (v_i)—visits M at $t = 0$.*
- (ii) **Entrepreneurs learn and persist.** *Their subjective odds of M being dangerous evolve as $\mathcal{O}(\mu_{it}) = \mathcal{O}(\mu_{i0})(1 - \bar{\rho})^t \rightarrow 0$, so they continue to visit and truthfully report that no harm occurs.*
- (iii) **Fearful households observe and learn.** *Any household observing entrepreneurs has odds falling by the factor $(1 - \bar{\rho})$ every safe return, so its belief crosses its cutoff in finite time, when it decides to visit M ; each adopter is a new source of information for those still reluctant.*
- (iv) **The mistaken beliefs are corrected.** *On a connected network, $\mu_{it} \rightarrow 0$ for every household and the visiting set converges to $\{i : v_i \geq 0\}$. The only households that do not visit are those for whom the Muslim market is not worth visiting on fundamentals. No intervention is required and no belief-based avoidance survives.*

Proposition 1 is the natural economic reading of the situation: markets should correct mistaken beliefs. This requires only that the distribution of fear and of stakes be non-degenerate, so that *somebody’s* threshold is crossed, and that what happens to that person be publicly visible. The optimistic conclusion rests entirely on the fact that a safe return by *anyone* is a safe return by *someone like me*. This is what we relax next.

6.3 Types and the Composition of the Goer Pool

We now suppose that risk is not common to all, but there are two types of households, $g_i \in \{s, e\}$. Type s households are *shielded* from risk (because they know the sellers, know when and where to go, and can get out) such that for them going to M is never dangerous. Type e households are *exposed* to risk, and face the uncertainty $\rho \in \{0, \bar{\rho}\}$ so they could get harmed if $\rho = \bar{\rho}$. Subjective risk is now

$\hat{p}_{it} = \mu_{it}\bar{\rho} \cdot \mathbf{1}\{g_i = e\}$, and the cutoff rule (5) governs exposed households while shielded households visit iff $v_i \geq 0$.

The key assumption we make is that types are not observable, so households must form a view about *who goes*. Households consider two hypotheses about the pool of Christians observed visiting M :

$$\begin{aligned} H_S : & \text{goers are all shielded;} \\ H_R : & \text{a share } 1 - \lambda \text{ of goers exposed.} \end{aligned}$$

Let $\pi_i \equiv \mathbb{P}_i(H_S)$ denote household i 's prior on the hypothesis H_S .

6.4 Learning Under Type Uncertainty

Own experience is unaffected, because a household knows its own type: for an exposed household, the likelihood ratio of visiting M and returning safely is $\text{LR}^{\text{own}} = 1 - \bar{\rho}$, so K own safe visits multiply its odds by $(1 - \bar{\rho})^K \rightarrow 0$, exactly as in the benchmark. What changes is the value of watching others.

Consider a household observing n distinct goers, each returning safely k times. The evidence must now be evaluated over four joint states, (ρ, pool) . A safe record has probability one in three of them: if the market is safe, nobody is ever harmed regardless of the pool; and if the goers are all shielded, nobody is ever harmed regardless of how safe the market is. Only under *both* $\rho = \bar{\rho}$ and H_R is a long safe record unlikely, with probability

$$\Lambda_n(k) \equiv [(1 - \lambda)(1 - \bar{\rho})^k + \lambda]^n \rightarrow \lambda^n \quad \text{as } k \rightarrow \infty. \quad (6)$$

Applying Bayes' rule, observing people return safely now updates priors on the safety of the market in the following way:

$$\mathcal{O}(\mu_i) = \mathcal{O}(\mu_{i0}) \cdot [\pi_i + (1 - \pi_i)\Lambda_n(k)] \geq \mathcal{O}(\mu_{i0}) \cdot \pi_i \quad (7)$$

Equation (7) shows the crux of the model: because there are types, learning from safe returns is bounded. In the benchmark case, observing enough safe returns would lead one to conclude the market is safe, because $\mathcal{O}(\mu_{i0})(1 - \bar{\rho})^k$ went to zero. In this case, observing safe returns for years will not lead you to update your beliefs about danger below $\mathcal{O}(\mu_{i0})\pi_i$, because it can always be the case that goers are returning safe because all of them are shielded, not because the market is safe. The higher π_i is, the less informative safe returns are.

6.5 Equilibrium: A Fear Trap Despite Entrepreneurs

Proposition 2 (Belief trap with voluntary entrepreneurship). *If $\pi_i > 0$ for a positive mass of households, a stationary belief trap appears in which:*

- (i) **Exposed entrepreneurs exist and are correct.** A positive mass of exposed households visits M , learns $\rho = 0$ from own experience, persists, and their safe returns are observed by others.
- (ii) **Non-goers' learning is bounded.** Every non-goer household observes an arbitrary number of safe returns and has beliefs $\mu_{it} \downarrow \mu_i^\infty$, where $\mathcal{O}(\mu_i^\infty) \geq \mathcal{O}(\mu_{i0}) \pi_i$.
- (iii) **A positive mass is permanently trapped by mistaken beliefs.** Any exposed household satisfying the trap condition

$$\mathcal{O}(\mu_{i0}) \pi_i > \mathcal{O}(\bar{\mu}_i) \quad (8)$$

never visits M despite observing safe returns, because it cannot learn enough for its fears to cross its threshold. This situation destroys $\int \mathbf{1}\{i \text{ trapped}\} v_i di$ of surplus per week, in perpetuity.

Comparing Propositions 1 and 2 locates the trap precisely. Economic incentives still exist and lead to the appearance of exposed entrepreneurs. However, the existence of types bounds non-goers' learning from safe returns, because observers can no longer take a safe return by someone else as a safe return by someone like themselves. That single wedge converts a self-correcting situation into a permanent one, where non-goers never overcome their fears despite the evidence.

Condition (8) says that fear traps are greater and more likely to appear in society if initial fears are high, goers are likely of a different class, and the feared event is catastrophic or likely to happen.

Remark 1 (π_i is self-confirming). π_i is not a misperception but an accurate reading of the equilibrium: The trap keeps exposed households away from M , so the goer pool is indeed heavily selected toward the shielded. No one need be mistaken about who goes for everyone to remain mistaken about the market.

Remark 2 (What can and cannot move π). Outcomes cannot lower π : safe returns actually raise it (see appendix). Only evidence about *identity* lowers π —personal knowledge that particular goers are exposed households; or, a perceived goer population too large to consist only of connected people.

Remark 3 (Selection neglect debiases). The model showcases an inversion of the behavioral literature on selection neglect (Jehiel, 2018; Enke, 2020): Here, households are trapped in mistaken beliefs because they are aware that goers might be a selected class. Having selection neglect, instead, would allow households to escape the trap, as they would treat safe returns as fully informative.

6.6 Randomized Encouragement as an Information Technology

The model shows that what would break the trap is lowering π . And Remark 2 notes that information about the identity of goers (not just of their safe returns) is able to lower π . In this sense, the intervention was well suited to break the fear trap. Participants knew that the project was selecting only non-goers to participate, that those incentivized to go would be picked at random, and knew some of the participants as their neighbors. In this sense, the intervention made the hypothesis H_S , where all goers were shielded, very unlikely, and therefore π very low or zero. Safe returns now were fully informative and

watching enough safe returns would convince them that the market is safe. The intervention worked not because it added to the stock of evidence of safe returns, but because it removed the ambiguity—years of ambiguous examples could not do what a few weeks of unambiguous ones did.

7. Conclusion

We have documented a costly mistake that survived for years in plain sight of the evidence that would have corrected it. Christian households in Jos knew that Muslim markets were cheaper, rated Muslim traders as highly as they rated their own, and lived barely farther from New Market than from the market they already used. They stayed away all the same, for an average of 8.6 years, because they believed that going was dangerous. It was not. Four accompanied visits were enough to change their views: among treated participants the share naming safety as a downside of shopping at New Market fell from 97% to 14%, while their beliefs about prices, which had been accurate all along, did not move. The change then spread and held. Forty-four percent of the control group visited during the intervention and 83% had gone within two months of its end, by which point treated households were going more often, unpaid, than they had while being paid to go. Their assessments of Muslim sellers, of Muslims in general, and their behavior toward Muslims in incentivized games all improved alongside.

The harder question is not why participants learned, but why they had not learned already. A visible minority of Christians had been shopping at New Market for years and returning unharmed, and everyone knew it. Our reading is that this evidence failed because its source was ambiguous. A safe return can be attributed to the market being safe or to the person returning being the sort who was never at risk, and while the second reading remained available, no quantity of safe returns could settle the matter. Randomization removed the ambiguity. The people who went were known to be ordinary non-goers picked by lottery, so what happened to them could not be explained away by who they were. We advance this account because it is the one most consistent with what we observe and with what we learned in the field, not because we can establish it against every alternative. If it is right, the same structure may recur wherever an action is worth taking but believed dangerous, some people are shielded from the danger and others are not, and observers cannot tell them apart. What breaks a trap of that kind is not more evidence, but evidence that cannot be explained away.

References

- Agneman, G. and Chevrot-Bianco, E. (2023). Market participation and moral decision-making: Experimental evidence from Greenland. *Economic Journal*, 133(650):537–581.
- Aker, J. C., Klein, M. W., O’Connell, S. A., and Yang, M. (2014). Borders, ethnicity and trade. *Journal of Development Economics*, 107:1–16.

- Bandiera, O. and Rasul, I. (2006). Social networks and technology adoption in northern Mozambique. *Economic Journal*, 116(514):869–902.
- Bartling, B., Weber, R. A., and Yao, L. (2015). Do markets erode social responsibility? *Quarterly Journal of Economics*, 130(1):219–266.
- Beaman, L., BenYishay, A., Magruder, J., and Mobarak, A. M. (2021). Can network theory-based targeting increase technology adoption? *American Economic Review*, 111(6):1918–1943.
- Bénabou, R. and Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–164.
- BenYishay, A. and Mobarak, A. M. (2019). Social learning and incentives for experimentation and communication. *Review of Economic Studies*, 86(3):976–1009.
- Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F., and Palmer, C. (2024). Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice. *American Economic Review*, 114(5):1281–1337.
- Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.
- Bryan, G., Chowdhury, S., and Mobarak, A. M. (2014). Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh. *Econometrica*, 82(5):1671–1748.
- Bursztyn, L., González, A. L., and Yanagizawa-Drott, D. (2020). Misperceived social norms: Women working outside the home in Saudi Arabia. *American Economic Review*, 110(10):2997–3029.
- Bursztyn, L. and Yang, D. Y. (2022). Misperceptions about others. *Annual Review of Economics*, 14:425–452.
- Carter, M., Laajaj, R., and Yang, D. (2021). Subsidies and the African green revolution: Direct effects and social network spillovers of randomized input subsidies in Mozambique. *American Economic Journal: Applied Economics*, 13(2):206–229.
- Conley, T. G. and Udry, C. R. (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review*, 100(1):35–69.
- Davis, D. R., Dingel, J. I., Monras, J., and Morales, E. (2019). How segregated is urban consumption? *Journal of Political Economy*, 127(4):1684–1738.
- Dean, J. T., Kreindler, G., and Mbonu, O. (2026). The effect of exposure: Evidence from spatial choices in Nairobi. NBER Working Paper 34799, National Bureau of Economic Research.

- Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica*, 82(1):197–228.
- Enke, B. (2020). What you see is all there is. *Quarterly Journal of Economics*, 135(3):1363–1398.
- Enke, B. (2023). Market exposure and human morality. *Nature Human Behaviour*, 7(1):134–141.
- Esponda, I. and Vespa, E. (2018). Endogenous sample selection: A laboratory study. *Quantitative Economics*, 9(1):183–216.
- Falk, A. and Szech, N. (2013). Morals and markets. *Science*, 340(6133):707–711.
- Foster, A. D. and Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103(6):1176–1209.
- Henrich, J., Ensminger, J., McElreath, R., et al. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science*, 327(5972):1480–1484.
- Higazi, A. (2011). The jos crisis: A recurrent nigerian tragedy. Discussion Paper 2, Friedrich-Ebert-Stiftung, Abuja.
- Hirschman, A. O. (1982). Rival interpretations of market society: Civilizing, destructive, or feeble? *Journal of Economic Literature*, 20(4):1463–1484.
- Human Rights Watch (2001). Jos: A city torn apart. Technical report, Human Rights Watch, New York.
- Jehiel, P. (2018). Investment strategy and selection bias: An equilibrium perspective on overoptimism. *American Economic Review*, 108(6):1582–1597.
- Jha, S. (2013). Trade, institutions, and ethnic tolerance: Evidence from South Asia. *American Political Science Review*, 107(4):806–832.
- Krause, J. (2011). A deadly cycle: Ethno-religious conflict in jos, plateau state, nigeria. Technical report, Geneva Declaration Secretariat, Geneva.
- Lalo, M. and Sadiq, L. (2011). Jos: A tale of ‘two cities’. *Daily Trust*.
- Lowe, M. (2021). Types of contact: A field experiment on collaborative and adversarial caste integration. *American Economic Review*, 111(6):1807–1844.
- Madueke, K. L. (2018). Routing ethnic violence in a divided city: Walking in the footsteps of armed mobs in jos, nigeria. *Journal of Modern African Studies*, 56(3):443–470.

- Martin, P., Mayer, T., and Thoenig, M. (2008). Make trade not war? *Review of Economic Studies*, 75(3):865–900.
- Mernyk, J. S., Pink, S. L., Druckman, J. N., and Willer, R. (2022). Correcting inaccurate metaperceptions reduces Americans’ support for partisan violence. *Proceedings of the National Academy of Sciences*, 119(16):e2116851119.
- Mousa, S. (2020). Building social cohesion between Christians and Muslims through soccer in post-ISIS Iraq. *Science*, 369(6505):866–870.
- Munshi, K. (2004). Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution. *Journal of Development Economics*, 73(1):185–213.
- Obateru, T. and Adama, E. (2011). Despite crisis, christians and muslims still do business together in jos. *Vanguard*.
- Ortiz, M. (2025). Hate, fear and intergroup conflict: Experimental evidence from Nigeria. Working Paper.
- Paluck, E. L., Green, S. A., and Green, D. P. (2019). The contact hypothesis re-evaluated. *Behavioural Public Policy*, 3(2):129–158.
- Pelnik, C. (2025). Moving to profitability? Alleviating constraints on microentrepreneur location. Working Paper.
- Rohner, D., Thoenig, M., and Zilibotti, F. (2013). War signals: A theory of trade, trust, and conflict. *Review of Economic Studies*, 80(3):1114–1147.
- Rustagi, D. (2024). Market exposure, civic values, and rules. Warwick Economics Research Paper 1501, University of Warwick.
- Scacco, A. and Warren, S. S. (2018). Can social contact reduce prejudice and discrimination? Evidence from a field experiment in Nigeria. *American Political Science Review*, 112(3):654–677.
- Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, 110(2):337–361.