Preaching Against the Choir: Religious Leaders and Reproductive Norms in Kenya*

Eduard van der Merwe†

This version: August 31, 2026

Abstract

Religion shapes beliefs and behaviors globally, yet whether religious authority can shift entrenched social norms remains largely unmeasured. The answer depends on the strength of conformity, which observed behavior entangles with authority and private conviction, posing an empirical challenge. I extend a discrete-choice model of social interactions to include a religious leader and scripture alongside peer conformity, and estimate it via a vignette survey of reproductive-age women in Meru County, Kenya. In third-person scenarios, I independently randomize the community norm, the leader's stance, and the scriptural frame, and identify the weight of each. I find that elicited secret use exceeds stated intention, which reveals high latent demand for family planning, but peer norms suppress use. Standard self-reports underestimate this peer influence, ranking it near the bottom even though it is the dominant channel. A religious authority crowds out the peer norm rather than reinforcing it. Therefore, religious leaders are most effective when preaching against an entrenched social norm. Difference-in-differences estimates from two faith-based programs corroborate the mechanism where faith and trust are densest. By quantifying how and when authorities shift norms, I demonstrate that religion can act as a measurable lever for development policy.

Keywords: religion; social norms; religious leaders; family planning; Kenya

JEL Classification: O12, Z12, Z13, D91, J13, C93

*I gratefully acknowledge financial support from the Social Sciences Research Institute, the Wagoner Foreign Study Scholarship Program, and the Department of Economics at Rice University. This study was approved by the Rice University Institutional Review Board (protocol IRB-FY2026-386) and conducted in partnership with Sodzo Kenya.

†Department of Economics, Rice University, Houston, Texas, United States. e.vdmerwe@rice.edu.

I am grateful to Rossella Calvi, Flávio Cunha, Elaine Howard Ecklund, Isabelle Perrigne, Willa Friedman, Kerby Goff, and Jeff Dominitz for their invaluable mentorship and insightful advice throughout this project. I also thank participants at ASREC 2026 and the 2026 IRES Graduate Workshop for their helpful comments and vibrant discussion. I thank Mona Bormet for her expert insights on family planning in Kenya. Finally, I extend my sincere thanks to Sodzo International and their Kenya team for their indispensable partnership and logistical support during the fieldwork.

1 Introduction

Across much of the developing world, religion shapes belief and behavior, and religious institutions command an authority that few others rival. Whether this authority can move entrenched social norms, and so serve as an instrument of development policy, remains largely unmeasured. Nowhere is the question sharper than in Africa's reproductive norms. Religious institutions speak directly to childbearing, and they do so where the norm around family planning (FP)[1] is both strong and contested. What a religious leader does with that authority is theoretically ambiguous. A leader may hold an entrenched norm in place, or be the one figure capable of breaking it.

For five decades, FP has been a cornerstone of development strategy, with well-documented returns for women's health, education, and economic participation (Bailey, 2013; Goldin & Katz, 2002). Despite this investment and renewed commitments through the Sustainable Development Goals (United Nations, 2019), the effectiveness of standard FP interventions has become questionable, particularly in sub-Saharan Africa (SSA).[2] In 2022, roughly 164 million women of reproductive age had an unmet need for FP, and across much of SSA, fewer than half of those who wish to avoid pregnancy use a modern method (United Nations, 2022). The conventional argument for these persistent gaps is a supply-side one, namely that affordable contraception is both hard to access and of low quality (Blackstone et al., 2017; Bongaarts & Bruce, 1995; Tessema et al., 2016). Two recent studies find otherwise. Dupas et al. (2025) traces low uptake to high desired family size, and Bau et al. (2024) to misconceptions about contraception and infertility.[3] These findings shift the focus to traditional belief systems, culture, and social norms across SSA, a terrain religious leaders already occupy.

Whether such norms can be moved, and by whom, is the question I address through the lens of religious authority. Religious leaders are among the most trusted figures in SSA, and where state institutions fail they have the standing and the reach to move a community's norm. I ask whether such leaders change what women privately want, or instead coordinate behavior around preferences women already hold. Coordination of that kind would resolve a collective-action problem that can keep a community at low levels of FP use.

Answering that question requires separating three channels that observed behavior blends together: (1) an individual's private conviction, (2) her conformity to community norms, and (3) her deference to religious authority. Empirically isolating these forces is challenging. They are entangled, and direct questions about them are distorted by social desirability because the very

[1] Christian Connections for International Health (CCIH) defines FP as "enabling couples to determine the number and timing of pregnancies, including the voluntary use of methods for preventing pregnancy—not including abortion—harmonious with their values and religious beliefs."

[2] Appendix Figure A1 visualizes this for unmet need, contraceptive use, and fertility.

[3] Bongaarts and Casterline (2013), Caldwell and Caldwell (1987), and Timæus and Moultrie (2020) also explore the phenomenon of high desired family size.

factors at issue make women reluctant to admit that peers or leaders sway their choices. Because these mechanisms are so seldom told apart, separating them in this paper does more than merely explain FP use. It offers a general way to measure how trusted authorities shape the social norms that hold communities in place.

To measure these forces, I pair a structural model with an original survey experiment. I extend Brock and Durlauf's (2001) discrete choice model with social interactions in which a woman's decision responds to her private conviction, the community norm she perceives, and her religious leader's position. I estimate this model using a vignette-based survey of reproductive-age women fielded in Meru County, Kenya. Because direct questions about contraception invite social-desirability bias, the survey presents short third-person scenarios about a fictional woman, "Kendi," and independently randomizes the community norm she faces, her leader's stated position, and the scriptural frame she holds. This randomization identifies the weight a woman places on conformity, scripture, and leader authority from experimental variation. The estimates answer the question posed above, and show how strongly conformity amplifies a leader's influence. The estimated parameters then support counterfactual policy simulations of how far a leader's advocacy can shift community norms.

I find a wedge between what women privately want and what they publicly do. In the same scenarios, women judge that Kendi has about a 40% chance of using FP when her choice is public, but 70% say she would use it secretly. Latent demand is high, and the constraint that holds use down is social. This social constraint is also the force women are least willing to name. In a first-person pebble allocation game over the drivers of her own choice, women give the peer norm only about 14% of their pebbles, less than they give their own conviction or scripture. The vignettes reveal that same norm to be the force that moves them most. Therefore, standard self-reports invert the true ordering, a social-desirability gap produced by the very forces I set out to measure.

How religious authority works is the paper's central result. Conformity to the peer norm carries more weight than a woman's deference to her religious leader, and scripture carries less still, with each weight precisely estimated. The more telling result is how a leader enters this system. When a religious leader takes a stance, or scripture is invoked, the weight women place on the peer norm falls, so authority displaces the norm rather than reinforcing it. A leader accomplishes the most by preaching against an entrenched norm and the least by preaching to an aligned one. A leader's blessing also moves behavior more than a prohibition, an endorsement premium consistent with a supportive message licensing women to act on demand that the community norm suppresses.

With latent demand already high, a religious leader's work is to resolve a collective-action problem. The estimated weights do not vary systematically with a woman's demographics, though the strength of conformity itself falls with household income and differs across denominations. This lever is real but bounded. Conformity is strong, its social multiplier is close to 10 where no author-

ity speaks, and about 3.3 averaged over the cues. However, it stays below the tipping threshold throughout, so a leader bends the FP norm at the margin without moving the community to a new equilibrium.

To validate these experimental findings in real-world behavior, I exploit the rollout of two faith-based FP initiatives in Kenya. The CHAK-Packard and Christian Advocacy for Family Planning in Africa (CAFPA) programs trained religious leaders as 'champions' to reduce FP stigma. Using a difference-in-differences design built on health facility locations and Demographic and Health Survey (DHS) data, I find heterogeneous effects concentrated where the faith-based delivery network is dense and trust in religious leaders is high. While these aggregate results cannot by themselves separate the distinct behavioral channels, their pattern serves as the mechanism's signature at a national scale. It appears exactly where the structural model predicts a leader's influence should be greatest.

The policy implications follow from the estimated weights. Because leaders resolve a coordination problem, engaging them is most effective where conformity is strong, trust in them is high, and a faith-based network can deliver on a shifted norm. The structural estimates locate the first two conditions, and the programmatic evidence confirms the third. The counterfactual simulations then quantify how far leader advocacy can move a community and reveal the theoretical thresholds at which conformity would be strong enough to tip a community to a new equilibrium. This turns a diffuse call to "involve religious leaders" into a targeted, testable design. Because the framework is not specific to FP, the same logic extends to other norm-governed behaviors, like female genital cutting and intimate-partner violence, where the binding constraint is social rather than material.

I contribute to four strands of literature. First, I advance the measurement of social interactions. Identifying the strength of conformity is difficult because it is confounded with correlated preferences and shared environments, a version of Manski's (1993) reflection problem. Building on the discrete-choice framework of Brock and Durlauf (2001), I use experimentally randomized vignettes to separate conformity, religious authority, and private conviction. I then validate these estimates against a first-person elicitation, and use the divergence between the two to measure social-desirability bias. Second, I contribute to the study of leaders as agents of social and cultural change. A theoretical literature casts prominent actors as coordinators who move communities between social equilibria (Acemoglu & Jackson, 2015), and leaders are known to shape economic outcomes and cultural traits more broadly (Jones & Olken, 2005; Prummer & Siedlarek, 2017). Closest to my setting, Bassi and Rasul (2017) show that religious authority shifts fertility-related behavior. I add experimental evidence on the mechanism, quantifying how far a leader's influence operates by coordinating behavior rather than by changing private preferences.

Third, I contribute to the literature on the persistence and change of social norms. Suboptimal

norms can persist as self-enforcing equilibria even when they are privately unpopular (Ashraf et al., 2020; Bursztyn et al., 2020), and recent work shows they can be shifted by correcting misperceptions or by well-placed interventions (Field et al., 2021; Gulesci et al., 2025). I identify a distinct lever, the trusted religious leader, and the conditions under which engaging one shifts a community's norm. Finally, I contribute to the economics of religion by treating religion as a contemporary policy instrument rather than a historical determinant (Becker et al., 2021). Where much of this literature studies religion's long-run economic legacies, I show, in the spirit of Bryan et al. (2020), that engaging religious authority can shift present-day behavior, reframing faith and its leaders as a lever for development policy.

2 The Kenyan Context

This section establishes the religious and institutional landscape of Kenya to contextualize the empirical analysis. Kenya reflects a broader pattern across sub-Saharan Africa (SSA), where entrenched social norms coexist with trusted religious authorities, making the setting relevant for the wider region. First, I document religion as a dominant social force, showing that trust in religious leaders exceeds that of any other formal or informal institution. I then provide evidence of the central role faith-based organizations (FBOs) play in the national healthcare infrastructure, particularly in maternal care. Turning to individual behavior, I outline current reproductive patterns, identifying a substantial gender gap in fertility preferences and a substantial latent demand among women for birth spacing. Finally, I demonstrate that the primary barriers to family planning (FP) uptake stem from social norms and informational misconceptions rather than doctrinal opposition. Together, these stylized facts suggest that religious leaders possess the social capital necessary to act as intermediaries, potentially aligning community norms with underlying individual preferences.

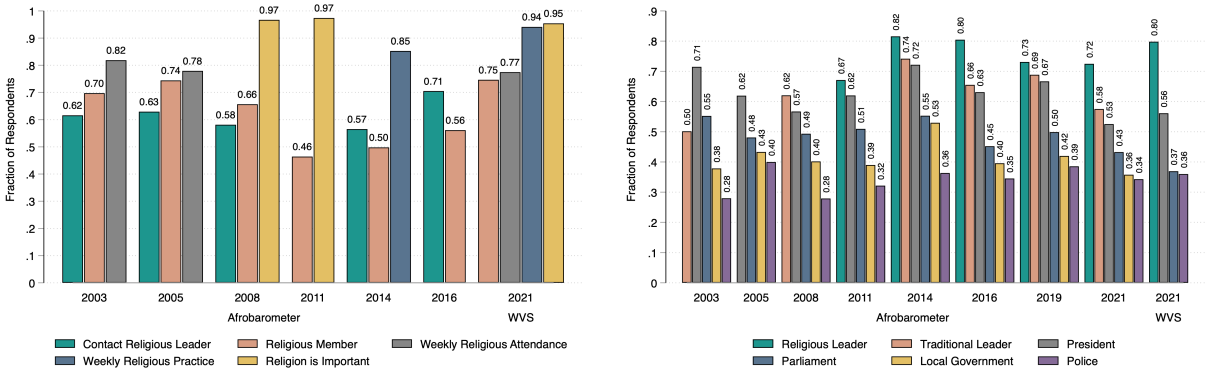
Religious Landscape Globally, more than three-quarters of the population is affiliated with a religious faith (Pew Research Center, 2025). This affiliation is even more pronounced in Kenya, where religion plays a central role in both public and private life. According to the 2019 Kenya Population and Housing Census, 85% of Kenyans identify as Christian and 11% as Muslim, while only 1.5% report no religious affiliation (KNBS, 2019).[4] Beyond mere identification, Figure 1a illustrates the high intensity of religiosity in Kenya. Consistently, over 95% of Kenyans report that religion is "somewhat" or "very" important in their lives. As of 2021, a large majority of Kenyans are active members of religious groups, attend weekly services, and engage in regular private religious practices such as prayer. While self-reported measures of devotion are occasionally scrutinized for social desirability bias, these figures collectively underscore the significance of religion

[4]Disaggregated, the population is 20% Catholic, 34% Protestant, 20% Evangelical, 7% African Independent Churches, and 1% Traditionalist.

in the Kenyan socio-cultural fabric.

Figure 1b illustrates that religious leaders consistently garner the highest levels of institutional trust in Kenya. In 2021, 80% of Kenyans reported trusting religious leaders at least "somewhat," with trust levels exceeding the continental average. This high degree of trust reflects the deep integration of religious leaders within local communities and cultural belief systems. These leaders therefore act as influential moral authorities, which positions them as key intermediaries in community decision-making. The physical infrastructure of religious institutions further bolsters this influence. Across SSA, FBOs provide between 20% and 50% of all healthcare services (Bormet et al., 2021).

Figure 1: Religiosity and Trust



(a) Indicators of Religiosity (b) Measures of Trust

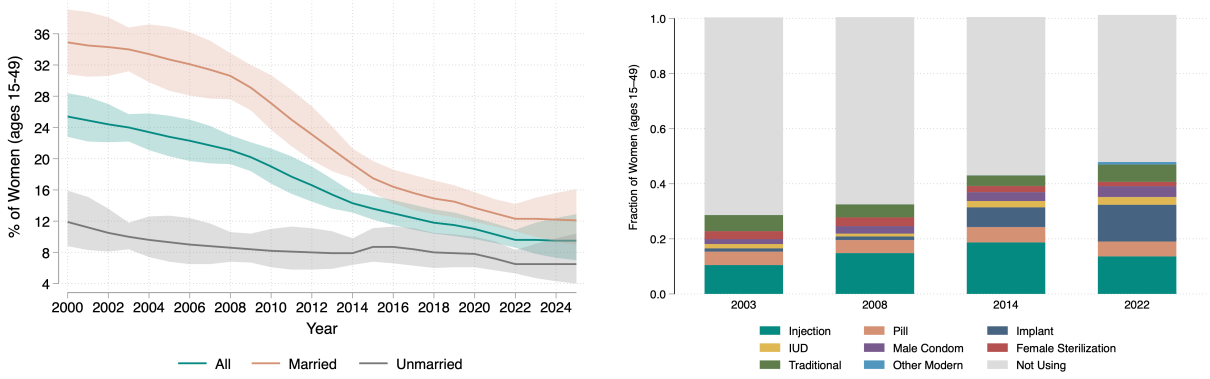
NOTE: Panel (a) reports the weighted fraction of respondents across five dimensions of religiosity: (1) active membership or leadership in a religious group, (2) weekly or more attendance at religious services, (3) weekly or more frequency of prayer, (4) those who consider religion "somewhat" or "very" important, and (5) those who have ever contacted a religious leader. Panel (b) reports trust measures which represent the proportion of respondents reporting "somewhat" or "a lot" of trust in religious, traditional, presidential, and local authorities. Data are from the Afrobarometer and World Values Survey.

Faith-Based Social Services In Kenya, data from the Kenya Master Health Facility Registry (KMHFR) show that of the 16,249 total health facilities, 1,026 are faith-based. Many of the 7,929 private facilities operate as not-for-profit or through NGOs frequently affiliated with religious institutions. These facilities often serve as the primary (or sole) source of care for remote populations. While the 2018 Kenya Household Health Expenditure and Utilization Survey (KHHEUS) indicates that 3.3% of rural and 5.2% of urban households seek outpatient care at FBOs, their systemic importance is most visible in specialized care. Critically, FBOs account for 11% of Level 4 facilities (sub-county referral hospitals) and 8% of Level 3 facilities (health centers and maternity homes). Given that these levels represent the primary points for maternity services and clinical referrals, religious institutions are a structural cornerstone of FP delivery.

Table A1 summarizes the KMHFR data on FP provision. Government facilities are the most likely to offer FP services across all levels of care, reflecting their mandate to ensure widespread service availability. Overall, FBOs have lower rates of FP provision. However, facilities associated with the Christian Health Association of Kenya (CHAK) and registered with the Kenyan health ministry are roughly 20 percentage points more likely to provide FP services than the national FBO average. Within CHAK, Seventh-day Adventist, African Independent Churches, and Methodist churches are the leading Protestant denominations in FP supply, with 83.3%, 66.0%, and 56.3% of their associated facilities offering these services, respectively.[5]

Reproductive Outcomes Figure 2a shows the evolution of reproductive health needs in Kenya over the past quarter-century. Following a period of marginal decline between 2000 and 2008, the unmet need for contraceptive methods fell sharply through 2015. While this downward trend has persisted, unmet need remains above 10% for the general population. Despite these gains, broader maternal health indicators reflect persistent challenges, according to the World Development Indicators. The adult female mortality rate, the probability of dying between ages 15 and 60, stands at 30%, and maternal mortality has largely stagnated relative to global benchmarks. Teenage childbearing remains prevalent, with nearly 15% of women aged 15–19 having already begun childbearing. However, Kenya has seen a steady decline in crude birth rates and infant mortality since 2014.

Figure 2: Women's Reproductive Measures



(a) Unmet Need: Modern Methods

(b) Method Mix

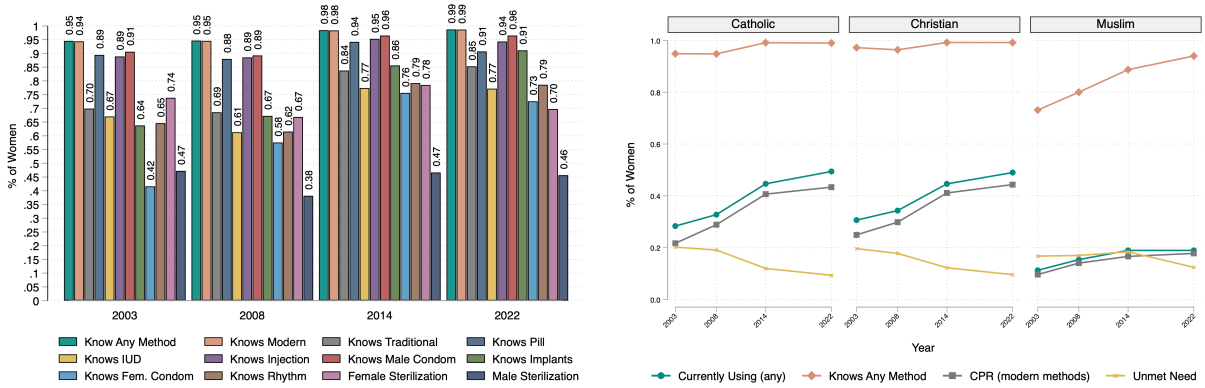
NOTE: Panel (a) presents estimates of the percentage of Kenyan women (ages 15–49) with an unmet need for modern contraceptive methods, stratified by marital status. The shaded regions denote the uncertainty intervals surrounding the estimates. Data are sourced from UNDESA (2024). Panel (b) illustrates the changing distribution of contraceptive method use among the same demographic across four survey waves (2003, 2008, 2014, and 2022). Data are derived from the Kenya Demographic and Health Surveys.

Figure 2b illustrates the shifting landscape of contraceptive method mix among Kenyan women

[5]These figures reflect KMHFR facility records as of September 2025. Only facilities that explicitly list FP as an active service are counted toward the FP provision rate.

between 2003 and 2022. While injectable contraceptives have historically been the most prevalent modern method, the data reveal a marked shift by 2022, where the uptake of implants has grown to reach parity with injections. For women, particularly those in rural contexts, the rise of long-acting reversible contraceptives (LARCs) like implants offers practical advantages in terms of privacy, long-term effectiveness, and reduced frequency of facility visits. Despite the expansion of these modern options, the use of traditional methods has remained remarkably stable over the two-decade period. Figure 3a highlights that knowledge of at least one contraceptive method is near-universal among all women. As shown in Figure 3b, when analyzing by religious affiliation, the differences in usage, knowledge, and unmet need between Catholics and other Christians are negligible. In contrast, while Muslim women report a lower contraceptive prevalence rate (CPR), they experience similar rates of unmet need compared to Christian groups.

Figure 3: Contraceptive Knowledge and Religious Differences



(a) Method Knowledge

(b) Religious Differences

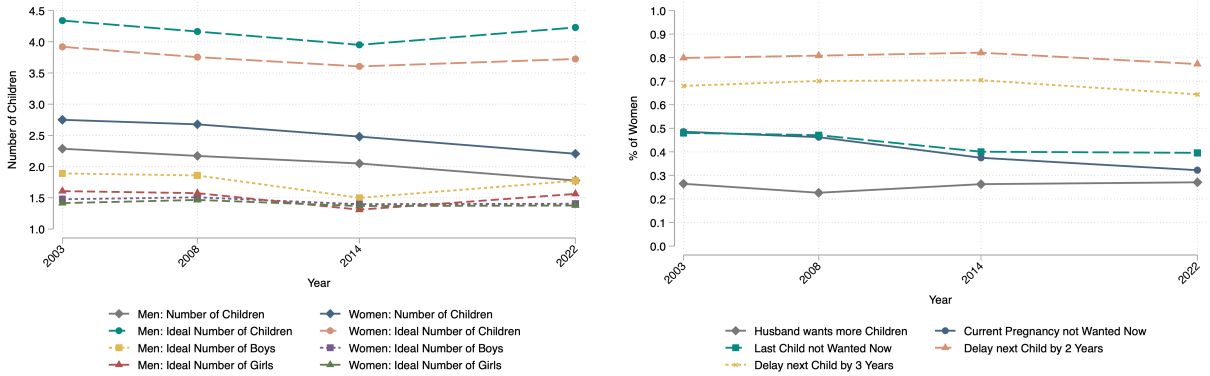
NOTE: Panel (a) presents the proportion of Kenyan women reporting knowledge of various contraceptive methods across four survey waves (2003, 2008, 2014, and 2022), detailing awareness of broad categories as well as specific modern and traditional methods. Panel (b) illustrates trends in key reproductive health indicators, including knowledge of any method, current use of any method, modern contraceptive prevalence rate (CPR), and unmet need, stratified by the respondents' religious affiliation (Catholic, Other Christian, and Muslim) over the same period. Data are sourced from the Kenya Demographic and Health Surveys.

Fertility Preferences A premise of this study is that effective development policy must move beyond "top-down" impositions and instead leverage existing religious structures to align policy goals with local preferences. Figure 4 displays the ideal number of children alongside total children ever born for both men and women. Across the survey periods, men's ideal family size consistently exceeds women's by approximately 0.5 children. Men are also more likely to express a preference for sons and, in this sample, report fewer total children ever born than women. While the desired family size in Kenya remains high, as noted by Dupas et al. (2025), a substantial portion of women express a desire to space or limit future births.

Specifically, 80% of women report a desire to delay their next pregnancy by at least two years

(and 65% by at least three years). A further 40% of women stated that their last child was either mistimed or unintended, while 30% of currently pregnant women report that the pregnancy was not desired at this point in time. These preferences often clash within the household, as 30% of women report that their husbands desire more children than they do. These discrepancies indicate that a major barrier to FP uptake is not a lack of latent demand for birth spacing, but rather intra-household disagreements and the social costs associated with deviating from established community norms. In this context, the high trust placed in religious leaders suggests they are uniquely positioned to act as intermediaries, bridging the gap between prevailing community norms and underlying individual preferences.

Figure 4: Fertility Preferences



(a) Ideal Number of Children

(b) Fertility Outcomes

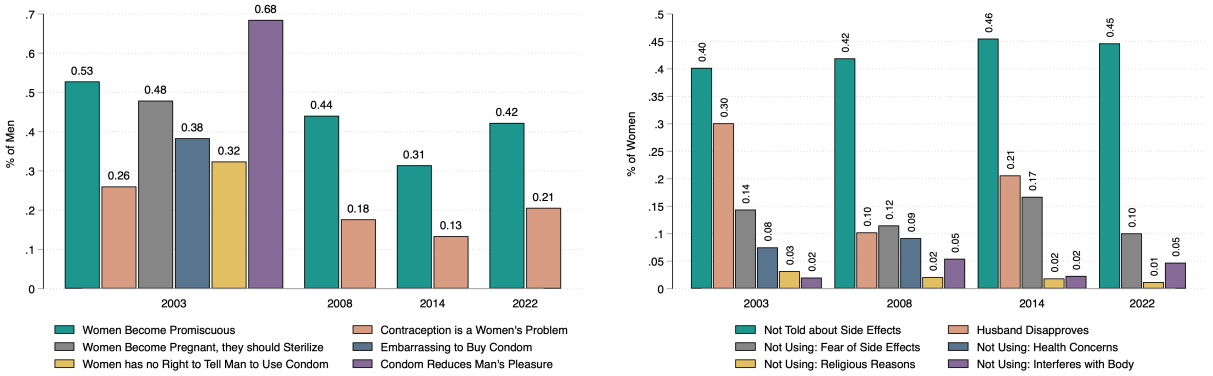
NOTE: Panel (a) compares the actual number of children to the reported ideal family size, disaggregated by gender, alongside respondents' specific preferences for the ideal number of boys and girls. Panel (b) illustrates women's fertility spacing preferences and pregnancy intentions over the same period. It tracks the desire to delay subsequent births, the retrospective reporting of recent or current pregnancies as mistimed or unwanted, and the proportion of women indicating that their husbands desire more children than they do. Data for all indicators are drawn from the 2003, 2008, 2014, and 2022 waves of the Kenya Demographic and Health Surveys.

Norms, Stigmas and Misconceptions Lastly, persistent social stigmas and misconceptions regarding contraceptive use represent a significant barrier to FP uptake. Figure 5a details male attitudes toward FP, which frequently associate contraceptive use with infidelity and social transgression. In both 2008 and 2022, over 40% of men surveyed stated that women who use contraceptives become "promiscuous." Furthermore, a gendered perception of responsibility persists, with approximately one-fifth of men viewing contraception as an exclusively female concern. First-round data from 2003 suggest even more rigid norms, with nearly half of men believing that women should undergo sterilization because they are the ones who become pregnant, and 40% reporting that purchasing condoms was socially embarrassing.

These stigmas are particularly acute for unmarried and young women, for whom the "social cost" of seeking FP services can be prohibitive. While knowledge of contraceptives is nearly univer-

sal, with over 98% of women aware of modern methods, Figure 5b illustrates that informational frictions remain. Fear of side effects is a primary driver of non-use, with nearly half of women reporting they never received counseling on potential side effects. Male opposition remains a hurdle, with one-fifth of women reporting husband disapproval in 2014. Only 2% of women cite direct religious prohibition as their reason for non-use. This suggests that the intersection of religion and FP is not primarily doctrinal, but rather embedded within broader social stigmas and intra-household dynamics. The severity of these gender norms extends well beyond contraception, with female genital cutting still affecting 15% of women ages 15–49, and 41% of ever-partnered women reporting violence from a partner (KNBS & ICF, 2023), both of which reflect the deeply entrenched community pressures that characterize this setting.

Figure 5: Norms and Misconceptions around FP



(a) Norms Around Contraception Use (b) Reasons why Women do not use Contraception

NOTE: Panel (a) illustrates the prevalence of various attitudes, norms, and misconceptions regarding family planning among Kenyan men, tracking beliefs such as contraception encouraging promiscuity or being exclusively a woman's responsibility. Panel (b) details the primary self-reported reasons women cite for not using contraception (including fear of side effects, spousal disapproval, and religious objections) alongside the proportion of women who report not being informed about potential side effects by family planning providers. Data are drawn from the 2003, 2008, 2014, and 2022 waves of the Kenya Demographic and Health Surveys. Note that several indicators in Panel (a) were only recorded during the 2003 survey wave.

3 A Model of Authority, Conformity, and Social Norms

To formalize how a trusted authority shapes an entrenched social norm, I extend the discrete-choice model of social interactions developed by Brock and Durlauf (2001), and use it to separate peer influence, authority, and scripture from an individual's private conviction. Although I study family planning (FP) in a context where the authority is a religious leader, the framework generalizes to any setting in which a trusted authority confronts a suboptimal norm sustained by coordination failure. Female genital cutting (FGC) and intimate-partner violence are two such settings. In much of sub-Saharan Africa (SSA), few other institutions combine that trust with that reach. In

the model, an individual's utility is shaped not only by the private costs and benefits of acting, but by a triangulated system of social influence. Horizontal influence is peer conformity, the pull of community behavior and stigma, and it carries weight α. Vertical influence is the authority's endorsement, which lowers the social cost of adoption, and it carries weight γ. Internalized influence is the pull of scriptural or doctrinal teaching, and it carries weight δ. Casting leader advocacy as a shift within this system allows me to examine the extent to which a trusted authority can coordinate community behavior and, when conformity is sufficiently strong, transition the community to a new equilibrium.

3.1 Model Setup

Each individual i in a finite population of I agents chooses a binary action $y_i \in \{1, -1\}$. The choice responds to four forces: (1) the individual's own preference, (2) a normative frame, (3) the community norm, and (4) the stance of a trusted authority. I apply the model to FP, where $y_i = +1$ denotes the use of an FP method and $y_i = -1$ denotes non-use. The trusted authority is the religious leader whose stance is R_i, and the normative frame is the scripture S_i that the individual internalizes. Abstracting from within-household bargaining, the utility of individual i choosing y_i is

$$U_i(y_i) = \underbrace{y_i\,(\boldsymbol{x}_i\beta + \delta S_i)}_{\text{Individual Preference}} + \underbrace{\alpha y_i E_i[y_{-i}]}_{\text{Social Norm}} + \underbrace{\gamma y_i R_i}_{\text{Religious Leader}} + \varepsilon_i(y_i), \tag{1}$$

where $\boldsymbol{x}_i$ is a vector of individual characteristics (e.g., age, parity, marital status) and β captures the baseline intrinsic preference for FP in a norm- and religious-neutral environment. That is, the community is evenly split ($E_i[y_{-i}] = 0$), the individual feels neither morally obligated nor morally prohibited ($S_i = 0$), and the religious leader is neither in favor of nor against the use of FP methods ($R_i = 0$). A positive β implies that the individual derives higher utility from FP use, while a negative β implies the individual derives higher utility from non-use. These personal preferences may reflect beliefs about health complications, desire for more children, or economic factors.

The parameter δ captures the weight the individual places on scriptural teaching, and S_i is the scriptural frame the individual holds, normalized to lie in $[-1, 1]$. Here, S_i close to 1 corresponds to a reading that supports spacing and contraception, while S_i close to -1 corresponds to a pronatalist reading that opposes it. Section 4 gives the specific passages used to construct these frames. The parameter α represents the strength of conformity to social norms. If the individual is very concerned about their status in the community or fears sanctions or ostracism, the parameter will be large. The term $E_i[y_{-i}]$ denotes the individual's subjective belief about the average behavior of others in the community. It is close to 1 when the individual believes most women use an FP method, and close to -1 when they believe few do. A value of $E_i[y_{-i}] = 0$ therefore represents the point of maximum social ambiguity (a perceived 50/50 split in the village).

The parameter γ measures the individual's deference to the religious leader. A high value of γ implies that the individual weighs the leader's position heavily. The leader's stance on FP, denoted R_i, is normalized to lie in $[-1, 1]$, where R_i close to 1 signals support, R_i close to -1 signals opposition, and $R_i = 0$ indicates a neutral or silent position. Lastly, $\varepsilon_i(y_i)$ is an idiosyncratic preference shock, assumed to be i.i.d. Type I Extreme Value, which yields the standard logistic choice probabilities. It captures individual-specific factors that affect the desirability of the action. In observed data, these forces are entangled. A community's norm $E_i[y_{-i}]$, its leader's stance R_i, and the scriptural frame S_i move together and respond to one another, so I cannot recover their separate weights from observational correlation alone. To break this entanglement, I independently randomize the three in a survey experiment, which identifies α, γ, and δ from exogenous variation, not from equilibrium behavior. I describe the design and estimation in Section 4.

3.2 Equilibrium Analysis

The magnitude of α shapes how a community responds to its religious leader, and it opens two routes. If α is modest, the community has a single equilibrium. A leader's endorsement or a shift in scripture moves behavior, and the reactions of neighbors amplify it through a social multiplier, but the community still settles close to where it began. In this case, the leader bends the norm. If α is large enough, multiple equilibria coexist, and a sufficient push can carry the community from widespread non-use to widespread use in a self-reinforcing cascade. In this case, the leader tips the norm, and otherwise similar villages can settle at very different rates of use. Which route is available depends on the size of α, which I estimate in Section 4 and characterize formally below.

Following Brock and Durlauf (2001), I assume individuals have rational expectations, so their subjective expectation of others' behavior, $E_i[y_{-i}]$, equals the actual average behavior in the community, $\bar{y} \in [-1, 1]$. Given the i.i.d. Type I Extreme Value distribution of the error terms ε_i, the probability that individual i chooses to use FP ($y_i = 1$) is

$$\Pr(y_i = 1) = \frac{\exp\left(\boldsymbol{x}_i\beta + \delta S_i + \gamma R_i + \alpha\bar{y}\right)}{\exp\left(\boldsymbol{x}_i\beta + \delta S_i + \gamma R_i + \alpha\bar{y}\right) + \exp\left(-\left(\boldsymbol{x}_i\beta + \delta S_i + \gamma R_i + \alpha\bar{y}\right)\right)}. \tag{2}$$

Using the identity $\tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$, this probability can be compactly expressed as

$$\Pr(y_i = 1) = \frac{1 + \tanh\left(\boldsymbol{x}_i\beta + \delta S_i + \gamma R_i + \alpha\bar{y}\right)}{2} \tag{3}$$

The Self-Consistency Condition Under rational expectations in a large population, and since $y_i \in \{-1, 1\}$, the average choice $\bar{y}$ is defined by the expected value $\mathbb{E}[y_i] = 2\Pr(y_i = 1) - 1$. This

yields the fundamental self-consistency equation:

$$\bar{y} = \int \tanh\left(\boldsymbol{x}\beta + \delta S + \gamma R + \alpha\bar{y}\right) dF(\boldsymbol{x}, S, R). \tag{4}$$

Replacing the population average with a representative member at the mean characteristics gives the deterministic utility $h = \bar{\boldsymbol{x}}\beta + \delta\bar{S} + \gamma\bar{R}$, which aggregates private preference, scripture, and the leader's stance. This is exact when the community is homogeneous and a mean-field approximation otherwise, since the average of $\tanh(\cdot)$ need not equal $\tanh(\cdot)$ at the mean. The equilibrium is then a fixed point of the mapping

$$\bar{y} = \tanh\left(h + \alpha\bar{y}\right). \tag{5}$$

Since $\tanh(\cdot)$ is continuous and maps $[-1, 1]$ into itself, Brouwer's Fixed Point Theorem guarantees at least one equilibrium. How many exist, and whether the leader bends or tips the norm, depends on the strength of conformity α relative to the deterministic utility h.

Equilibrium Multiplicity A central feature of this framework is that the strength of social conformity (α) dictates whether a community faces a single equilibrium or multiple self-fulfilling ones. Proposition 1 adapts the framework of Brock and Durlauf (2001) to the present setting. Proofs for this and all subsequent propositions are provided in Appendix C.

Proposition 1. *Let $h = \bar{\boldsymbol{x}}\beta + \delta\bar{S} + \gamma\bar{R}$.*

1. *If $\alpha \leq 1$, there exists a unique equilibrium.*
2. *If $\alpha > 1$, there exists a critical threshold $h^* > 0$ such that:*
   - *If $|h| > h^*$, the equilibrium is unique.*
   - *If $|h| < h^*$, there are three equilibria, two stable steady states (a "low-usage trap" and a "high-usage norm") and one unstable middle equilibrium.*
   - *If $|h| = h^*$, there are two equilibria with one being a double root.* ■

This multiplicity explains why two villages with nearly identical demographic characteristics ($\boldsymbol{x}$) and religious leanings (S) can exhibit radically different uptake rates. A village where $\bar{y}$ is near -1 is stuck in a norm trap, so even if individuals possess a latent desire to use FP ($\beta > 0$), the fear of social sanction (α) keeps the community anchored at a low-usage steady state.

The Social Multiplier A shift in the deterministic utility h, whether from the leader's stance R or the scriptural frame S, moves the community norm by more than its direct effect, because neighbors respond to one another. Figure 6 illustrates this dynamic. In the sub-critical regime ($\alpha < 1$), this multiplier is finite. The leader smoothly bends the norm. The multiplier diverges only as the community approaches a tipping point.

Proposition 2. *Any deterministic shift in a community's utility is amplified by the social multiplier M, which depends on the strength of conformity α. Holding demographics $\bar{\boldsymbol{x}}$ fixed, the marginal effects of the*

mean leader stance $\bar{R}$ and the mean scriptural frame $\bar{S}$ on the aggregate adoption rate $\bar{y}$ are

$$\frac{\partial \bar{y}}{\partial \bar{R}} = \gamma \left(1 - \bar{y}^2\right) M(\bar{y}; \alpha), \qquad \frac{\partial \bar{y}}{\partial \bar{S}} = \delta \left(1 - \bar{y}^2\right) M(\bar{y}; \alpha),$$

where the social multiplier M is defined as:

$$M(\bar{y}; \alpha) = \frac{1}{1 - \alpha(1 - \bar{y}^2)}. \tag{6}$$

The multiplier M satisfies $M \geq 1$ for all $\alpha \in [0, 1)$, and $M \to \infty$ as the system approaches a tipping point where $\alpha(1 - \bar{y}^2) \to 1$. ■

In both regimes, the leader acts as a coordination mechanism. They change behavior by shifting the social and moral costs (h) instead of altering the underlying intrinsic preference (β). This change is marginal when conformity is weak, and sharp when it is strong. Which regime a community occupies is an empirical question that I resolve by estimating α, γ, β, and δ in Section 4. The estimated model then supports policy-relevant counterfactuals and Figure 6 illustrates both routes.

3.3 Authority Salience and the Endorsement Premium

Salience and Crowding-Out The baseline holds the conformity weight α constant. Yet an authority cue may do more than shift utility directly through $\gamma \bar{R}$ and $\delta \bar{S}$. When a leader speaks, their word weakens the community's grip on its peer norm and pulls behavior toward their own message and away from what others do. This displacement does not depend on what they say. Whether they support the behavior or oppose it, speaking makes the peer norm less salient.

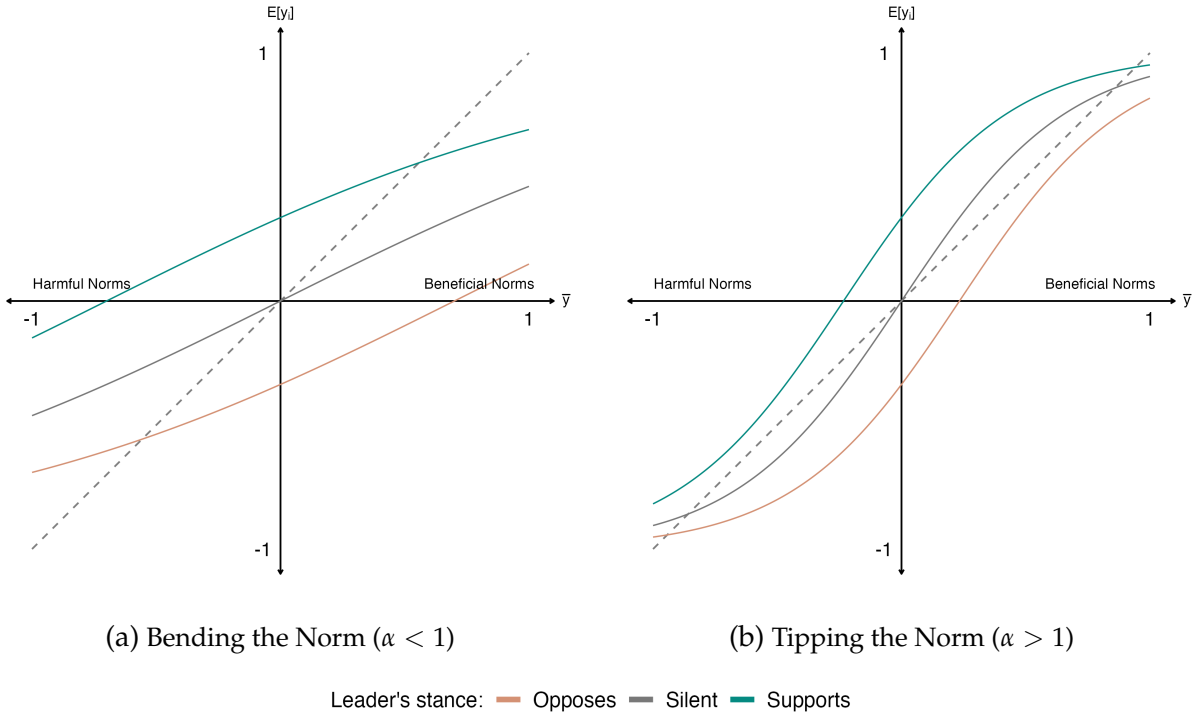
Because this displacement pulls behavior off the norm, its value turns on whether the leader is fighting the norm or echoing it. When they oppose an entrenched norm (preaching against the choir), pulling behavior off the norm moves it toward their message, so their direct push and the loosened norm reinforce one another. When they echo the prevailing norm (preaching to the choir), they are redundant, and loosening the norm undercuts the very behavior they favor. A leader's leverage is therefore greatest against a strong, misaligned norm and weakest when reinforcing an aligned one.

I capture this by letting the conformity weight respond to the cues an individual faces. The extended utility is

$$U_i(y_i) = \underbrace{y_i \left(\boldsymbol{x}_i \beta + \delta S_i\right)}_{\text{Individual Preference}} + \underbrace{\alpha(R_i, S_i)\, y_i E_i[y_{-i}]}_{\text{Social Norm}} + \underbrace{\gamma y_i R_i}_{\text{Religious Leader}} + \varepsilon_i(y_i), \tag{7}$$

identical to the baseline except that the conformity weight now falls with the salience of the au-

Figure 6: Bending versus Tipping Equilibrium Outcomes



(a) Bending the Norm ($\alpha < 1$)

(b) Tipping the Norm ($\alpha > 1$)

NOTE: Each panel plots the actual average behavior $g(\bar{y}) = \tanh(h + \alpha\bar{y})$ against the expected average behavior $\bar{y}$, so equilibria lie on the 45-degree line where the two coincide. The three curves show an opposing, silent, and supportive leader ($h = -0.35, 0, +0.35$), and the scriptural frame shifts h in the same way. In Panel (a), conformity is sub-critical ($\alpha < 1$): the curve crosses the line once, and the equilibrium slides smoothly with h. In Panel (b), conformity is super-critical ($\alpha > 1$): the curve can cross three times, and once h moves past the threshold h^*, a stable crossing vanishes and behavior jumps to the surviving one.

thority cues,

$$\alpha(R_i, S_i) = \alpha_0 - \alpha_R|R_i| - \alpha_S|S_i|, \tag{8}$$

where α_0 is conformity in an authority vacuum. The absolute values encode salience and are blind to direction, since a cue displaces the norm whether it points for or against. Each cue enters the model twice: through a direct effect on utility (γR_i and δS_i) and an unsigned salience effect on the conformity weight ($\alpha_R|R_i|$ and $\alpha_S|S_i|$). The two are separately identified because one is linear in the cue and the other in its absolute value. Scripture enters symmetrically, which makes this crowding-out a general property of salient authority and not a peculiarity of the leader.

Aggregating to a representative community member as in Section 3.2, the conformity weight is $\alpha(\bar{R}, \bar{S}) = \alpha_0 - \alpha_R|\bar{R}| - \alpha_S|\bar{S}|$, and the equilibrium norm solves $\bar{y} = \tanh(h + \alpha(\bar{R}, \bar{S})\,\bar{y})$, with the weight itself responding to the cues. Two consequences follow, both stemming from the same displacement mechanism. First, since $\alpha(\bar{R}, \bar{S}) \leq \alpha_0$, a speaking authority lowers the social multiplier, so a leader's influence is amplified less than the constant-α model implies. Second, this same

loosening allows a leader to override an entrenched norm, since pulling behavior off a strong, misaligned norm is what breaks a norm trap. The tipping condition becomes cue-dependent, $\alpha(\bar{R}, \bar{S}) > 1$. Because a salient authority strictly reduces this weight ($\alpha(\bar{R}, \bar{S}) \leq \alpha_0$), their intervention moves a community further from the tipping threshold, never toward it.

Corollary 1. *For a supportive leader ($\bar{R} > 0$) and scriptural frame ($\bar{S} > 0$), the marginal effects on the equilibrium norm are*

$$\frac{\partial \bar{y}}{\partial \bar{R}} = (\gamma - \alpha_R \bar{y}) \left(1 - \bar{y}^2\right) M, \qquad \frac{\partial \bar{y}}{\partial \bar{S}} = (\delta - \alpha_S \bar{y}) \left(1 - \bar{y}^2\right) M,$$

where $M = 1/\left(1 - \alpha(\bar{R}, \bar{S})(1 - \bar{y}^2)\right)$, with the opposing case analogous. ■

Both factors formalize this mechanism. The effective push, $\gamma - \alpha_R \bar{y}$ for the leader and $\delta - \alpha_S \bar{y}$ for scripture, generates the "against-the-choir" asymmetry. It is larger against a misaligned norm (where $\bar{y}$ opposes the cue) and smaller when reinforcing an aligned one. The shared multiplier M, evaluated at the reduced weight $\alpha(\bar{R}, \bar{S})$, represents the substitution effect, as a salient cue amplifies less than an authority vacuum would. The gap between the leader's effect in this model and its constant-α benchmark is what I quantify in the leader counterfactual.

Endorsement versus Prohibition Thus far, the model restricts the direct authority effect to the symmetric term γR_i, which forces a supportive and an opposing message to move behavior by equal and opposite amounts. I relax this restriction. Allowing a blessing and a prohibition to operate separately, the direct leader effect becomes

$$\gamma^+ \mathbb{1}[R_i = 1] + \gamma^- \mathbb{1}[R_i = -1], \tag{9}$$

where γ^+ is the response to a supportive leader and γ^- to an opposing one. Scripture enters the preference parameter similarly through δ^+ and δ^-. The symmetric baseline is simply the restriction $\gamma^+ = -\gamma^-$. Whether a blessing and a prohibition carry equal weight therefore becomes an empirical question, no longer an assumption maintained in advance.

This endorsement premium, the gap between a blessing and a prohibition, is distinct from the against-the-choir asymmetry formalized in Corollary 1. That asymmetry concerns how a message's effect varies with the peer norm $\bar{y}$. Conversely, the endorsement premium concerns the sign of the message at a fixed norm. The two mechanisms are separately identified. The premium is identified by the main effect, whereas crowding-out is identified by its interaction with the peer norm. Because the endorsement premium term enters the deterministic utility h additively, it leaves the equilibrium and tipping analysis unchanged. If a supportive message primarily licenses behavior a woman already favors, while an opposing one merely restates a constraint she already feels, a blessing should outweigh a prohibition ($|\gamma^+| > |\gamma^-|$). I test this empirically in Section 5.2.

3.4 Extension to Respondent Heterogeneity

The baseline model holds the structural weights constant across respondents. However, respondents differ in whose guidance they prioritize. Allowing these weights to vary with a respondent's own characteristics addresses a between-respondent question of who is most responsive to peers, to the leader, and to scripture. This is distinct from the within-story salience mechanism of Section 3.3, where a given respondent's weight shifts dynamically with the specific cues she faces. Provided that these individual traits are cross-sectionally demeaned ($\mathbb{E}[Z_i] = \mathbf{0}$), these variations average out at the community level, and the aggregate tipping analysis of Section 3.2 is unaffected.

Let Z_i denote respondent i's vector of demographic and attitudinal characteristics. This is distinct from the vignette persona Kendi's profile $\boldsymbol{x}$, which determines the private preference β as established in Section 3.1. Theory dictates which respondent trait to pair with which structural weight. I allow a respondent's self-reported sanction-sensitivity to enter the norm weight α, her trust in leaders to enter the leader weight γ, and her scriptural orientation to enter the scripture weight δ. Parameterizing each weight as a linear function of Z_i yields

$$\beta_i = \beta + Z_i'\beta_Z, \quad \delta_i = \delta + Z_i'\delta_Z, \quad \alpha_i = \alpha + Z_i'\alpha_Z, \quad \gamma_i = \gamma + Z_i'\gamma_Z,$$

where β_Z, δ_Z, α_Z, and γ_Z are vectors of interaction coefficients to be estimated. Substituting these respondent-specific weights into the baseline utility of Section 3.1 leaves its structure and closed-form logit choice probabilities unchanged, with each weight now carrying an interaction with Z_i. Because Z_i is mean-centered, the model collapses to the common-weight baseline for the average respondent. The vectors of interaction coefficients therefore measure only heterogeneous departures from the population mean.

Estimation proceeds via maximum likelihood on the stacked vignette data, with interactions between the randomized vignette factors and Z_i. Standard errors are clustered at the respondent level. This extension leaves the core identification strategy unchanged while enabling direct hypothesis tests of heterogeneity, for example whether the leader-influence parameter γ_i is larger for highly religious respondents, or whether conformity to peer norms α_i varies with age. All average parameters (β, δ, α, and γ) and their associated heterogeneity terms remain cleanly identified by the independent randomization of the vignette factors.

4 Survey Design and Structural Identification

Section 3 showed that the community norm, the leader's stance, and the scriptural frame are entangled in observed behavior, so their separate weights cannot be read from correlation, and that asking about them directly invites social-desirability bias. This section describes the survey I de-

signed to overcome both obstacles and recover the structural parameters. The instrument is an original vignette survey that independently randomizes the three forces within each respondent, so that the weights α, γ, and δ are each identified from exogenous variation rather than from behavior. Because the scenarios concern a fictional woman, "Kendi," respondents report what someone else would do, the psychological distance that a direct question about their own contraceptive use cannot provide. The belief they report maps directly onto the choice probability of the model, so the estimated coefficients are the model's structural weights. I first describe the survey setting, sample, and composition, then the experimental design and the balance of its randomized factors, then the outcome measures, and finally the mapping from vignette responses to the model and the random-effects estimator.

4.1 Setting, Sample, and Descriptive Statistics

The primary data are collected through an original survey of $N = 635$ Christian women of reproductive age in Meru County, Kenya. Meru serves as an ideal setting, as it is a primarily agrarian society where traditional customs remain strong. It possesses a deep-seated religious institutional history stemming from the 19^{th} century Methodist Missions. This "dual-authority" environment, where traditional and religious structures coexist, allows the observation of how religious leaders can successfully act as coordination mechanisms to shift social norms. Meru was also among the three counties targeted by the CAFPA program in Phase 1, which grounds the structural estimation in the same institutional landscape as the quasi-experimental evidence in Section 6.

The women were recruited through Sodzo Kenya's table-banking savings groups, known locally as *Kuja Pamoja kwa Jamii* (Come Together for Community), in villages where Sodzo was operating or had recently begun to operate. The groups meet in the mornings, and enumerators surveyed the members who were Christian, of reproductive age, and willing to participate. Because Sodzo's presence differs across villages, the groups range from long-established to only a meeting or two old, so the sample is neither uniformly selected nor uniformly empowered. Trained Sodzo enumerators administered the survey in Kimeru. The interview was tablet-assisted for the questionnaire and data capture, while the vignette stimuli were delivered on physical printed cards and the allocation task used physical pebbles and boxes, which keep the tasks concrete for respondents with limited numeracy. I restrict the sample to Christian women to preserve doctrinal coherence in the scripture frames used in the vignette instrument.

Table B1 describes the sample. The respondents are on average 35 years old, and 81% are married or in a union. Formal schooling is limited, with 69% having completed at most primary education and one in four reaching secondary or beyond. The sample spans Kenya's major Christian traditions. Pentecostal churches and mainline Protestants, who in Meru are largely Methodist, Anglican, and Presbyterian, are each about a fifth of the sample, African Independent Churches a fur-

ther 15%, and Catholics and Adventists smaller minorities. The remaining third belong to a dense landscape of small independent congregations, split between charismatic and gospel churches and African-initiated Zionist and apostolic churches, which women identify by name rather than by a broad label. Religiosity is near-universal, as 95% attend services weekly and 90% report trusting religious leaders.

Family planning (FP) is familiar in this sample. Almost every woman knows a modern method, 95% have ever used one, and 84% currently use a modern method, even as the ideal family size remains high at roughly four children. Use is genuinely high across this area, where modern methods are provided free and are widely available, so the level reflects the setting rather than the recruitment frame. This makes it a high-use setting, unlike much of the lower-use sub-Saharan Africa (SSA) the paper concerns, and a demanding place to detect a normative constraint. The high use is a strength for what follows, for two reasons. First, the structural weights are identified from the randomized variation within each respondent's own vignettes, so the level of use in the sample bears on the descriptive quantities but not on the parameters I estimate. Second, these experienced users still judge a hypothetical peer's use to fall sharply under an adverse community norm and to rise once her choice is hidden, so the vignette recovers a normative constraint that their own behavior conceals. The design surfaces a force that revealed behavior in this setting would miss.

4.2 Experimental Design

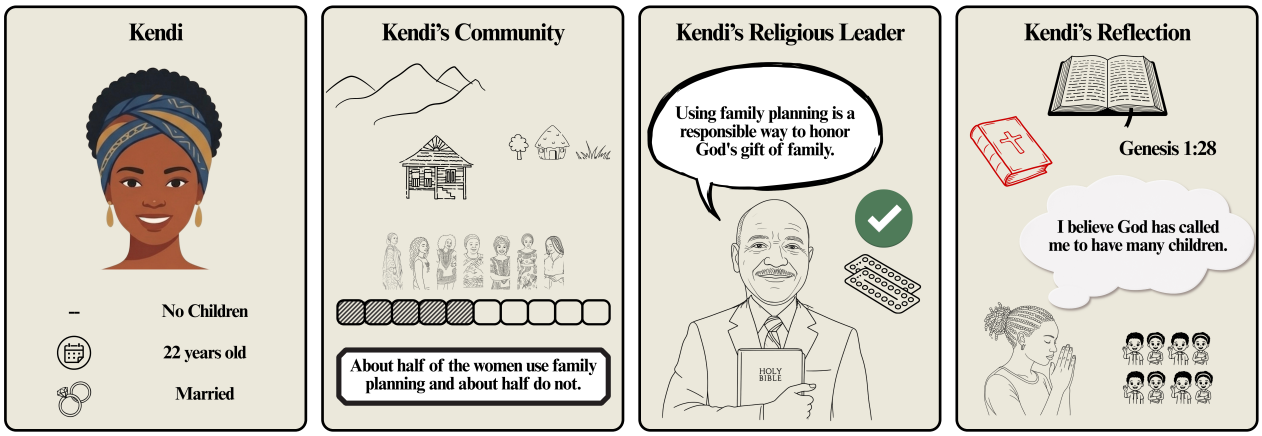
The survey implements a structured factorial vignette design with three independent treatment factors, each randomized across the eight stories presented to each respondent. These factors correspond directly to the three sources of influence in the utility model of Section 3.1. The levels are normalized so that the neutral level (0) corresponds to an ambiguous community norm, a silent religious leader, and no scripture exposure, respectively. The community-norm level is shown as a card marking how many of 10 women in the village use FP, 2, 5, or 8 of 10, as shaded blocks paired with the phrase "almost no," "about half," or "almost all." Pilot and focus-group testing confirmed that respondents read the card as the intended share. Mapping the depicted share p to the mean action $E[y_{-i}] = 2p - 1$ gives the coded values -0.6, 0, and $+0.6$. The norm regressor therefore lies on the same $[-1, 1]$ scale as the average behavior it represents, and α carries the multiplicity threshold of one.

Each story is laid out as four physical cards in front of the respondent, one for Kendi's profile and one for each treatment factor, with the story read aloud from the tablet. The profile card is fixed for the respondent, while the tablet randomizes each treatment factor's level, including its sign, and the order of the stories, with assignment balanced across respondents. The eight stories per respondent follow a structured design that isolates first-order parameters through single-factor variation, identifies pairwise interactions through two-factor variation, and provides a misspecifi-

cation test via the fully saturated omnibus story. The three treatment cards each carry a concrete message. The community-norm card reports how common FP is among the women in Kendi's village. The leader card conveys her religious leader's public position. He has taught the congregation that FP is a responsible way to honor God's gift of family, has not spoken about it, or has taught that it shows a lack of trust in God's plan.

The scripture card sets the passage Kendi has been reading. Its neutral level presents no passage. Its pronatal level gives Genesis 1:28 and the call to be fruitful and multiply, which she takes as a call to have many children. Its stewardship level gives 1 Timothy 5:8, on providing for one's household, which she takes as a call to care and plan for each child. Because the leader and scripture cues are specific, recognizable messages, the estimated weights γ and δ can be read as responses to real religious content rather than to abstract labels. Pilot and focus-group testing also confirmed that the two nonzero levels of each card are balanced around its neutral level. The low and high norm shares, the supportive and opposing leader messages, and the pronatal and stewardship scriptural frames each carry equal weight. Any asymmetry the estimates later reveal between a supportive and an opposing cue is therefore a property of how women respond (Section 5.2).

Figure 7: A Complete Vignette Example



(a) Kendi's profile (b) Community norm (c) Religious leader (d) Scripture frame

NOTE: The four cards placed in front of a respondent for a single story are Kendi's profile, which is fixed across her eight stories, and the three treatment cards, whose levels the tablet randomizes independently. The examples shown are a married and childless profile, a village where about half of women use family planning, a religious leader who supports it, and a pronatal scripture frame. The complete set of cards, the five profiles and the three levels of each treatment factor, appears in Appendix D.

Before the vignettes begin, each respondent is randomly assigned one of five *Kendi lifecycle profiles*, and that profile remains fixed across all eight stories. The profiles are A (age 16, unmarried, no children), B (age 22, unmarried, no children), C (age 22, married, no children), D (age 22, married, one child), and E (age 28, married, one child). The profiles form a ladder in which each

Table 1: Vignette Treatment Factors

Factor	Symbol	Levels	Parameter Identified
Community Norm	$E_i[y_{-i}]$	$-0.6,\ 0,\ +0.6$	α_0
Leader Advocacy	R_i	$-1,\ 0,\ +1$	$\gamma^+,\ \gamma^-$
Scripture Frame	S_i	$-1,\ 0,\ +1$	$\delta^+,\ \delta^-$

adjacent pair differs in a single attribute (age for A→B, marriage for B→C, parity for C→D, and age again for D→E), so that the four adjacent contrasts are pre-specified single-attribute effects. This between-respondent variation in Kendi's observable characteristics provides an experimental source of identifying variation for the baseline preference parameter β, and tests the model's prediction that these lifecycle attributes shift the baseline utility index.

Table 2: Survey Design Matrix

Story	**Name**	S	E	R	**Identification Purpose**
1	Baseline	0	0	0	Anchor: all factors neutral
2	Scripture	±1	0	0	Main effects $\delta^+,\ \delta^-$
3	Norm	0	±0.6	0	Main effect α_0
4	Leader	0	0	±1	Main effects $\gamma^+,\ \gamma^-$
5	Scripture × Norm	±1	±0.6	0	Scripture salience α_S
6	Norm × Leader	0	±0.6	±1	Leader salience α_R
7	Scripture × Leader	±1	0	±1	Leader × scripture check
8	Omnibus	±1	±0.6	±1	Over-identification test

NOTE: $\pm$ indicates that the tablet randomly assigns the valence for each non-zero factor, balanced across respondents. Story order is fully randomized within respondent. S: scripture frame; E: community norm ($E[y_{-i}]$); R: leader advocacy.

Because the three factors and the profile are assigned at random, they should be uncorrelated with one another and with the women who receive them, and they are. Table B2 reports that the community norm, the leader's stance, and the scripture frame are mutually uncorrelated across the 5,080 stories and uncorrelated with story position and respondent characteristics. Table B3 reports that the randomly assigned Kendi profile is balanced across every respondent characteristic, so no profile is systematically shown to a different kind of woman. This orthogonality is the design validity on which the structural identification rests.

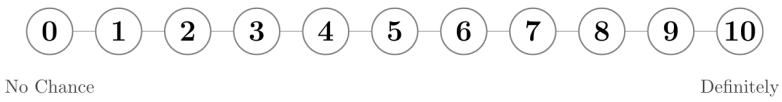
4.3 Outcome Measures

Direct questions about contraceptive use or intentions are susceptible to social desirability bias in this context, precisely because community norms, scriptural conviction, and religious leader in-

fluence are the objects of interest. A respondent who faces strong social pressure against FP will underreport her true preference when asked directly. Individual preferences and conformity pressures are then indistinguishable. To circumvent this, the survey employs a third-person vignette design. Respondents are presented with short stories about a fictional woman named "Kendi" and asked to report what they believe women in Kendi's situation would do, not what they themselves would do. This framing reduces social desirability bias by creating psychological distance between the respondent and the sensitive behavior.

Two questions are asked after each vignette. The primary outcome is a *likelihood belief* elicitation: "On a scale of 0 to 10, where 0 is no chance and 10 is definitely, what is the chance that Kendi will go to the clinic to receive a family planning method?" Responses are recorded on a 0–10 integer chance scale (0 = no chance, 10 = definitely) shown in Figure 8. Rather than reporting a number verbally, the respondent points to the position corresponding to her estimate. This instrument is cognitively natural for respondents with limited numeracy, and directly elicits her belief about the individual choice probability $\Pr(y_i = 1)$ for Kendi. Eliciting the chance for a single agent matches the estimand directly, since the model's left-hand side is the probability that one individual facing the posted norm chooses FP. Asking about a single hypothetical also keeps the posted norm E_k exogenous rather than inviting an equilibrium reading in which the women's choices are mutually determined.

Figure 8: The Zero-to-Ten Chance Scale



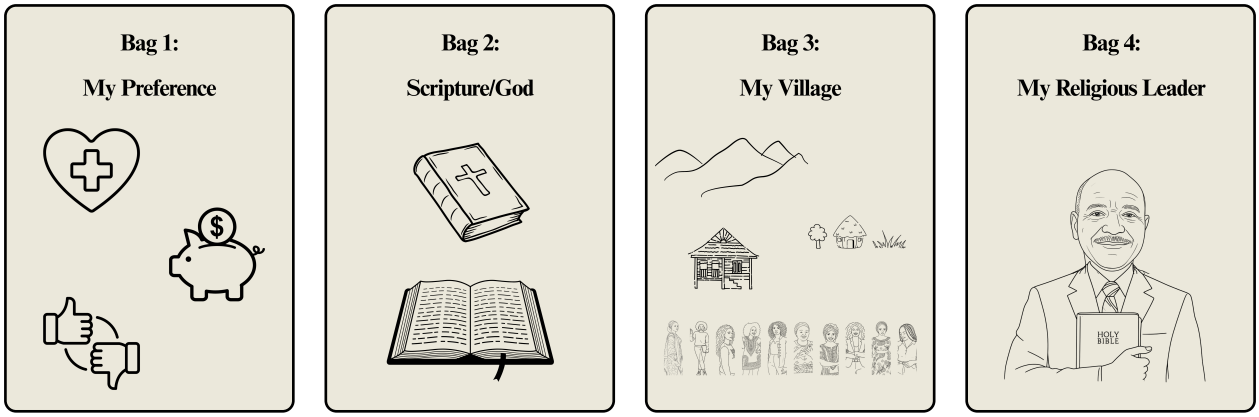
NOTE: The scale respondents used to report the primary outcome. Pointing to a position from 0 ("no chance") to 10 ("definitely"), the respondent gave her belief about the chance that Kendi would go to the clinic for family planning, rather than reporting a number verbally.

The secondary outcome is the *normative agreement* question measured on a five-point Likert scale asking how strongly the respondent agrees that "Kendi should use a family planning method." This captures the normative dimension separately from the descriptive belief above. The gap between the two responses is informative. If a respondent believes that most women in Kendi's situation would use FP, but that Kendi herself should not, her restraint reflects an internalized moral constraint rather than a descriptive expectation about peers.

A key supplementary instrument is the *secret injection question*. After selected vignettes, respondents are asked whether Kendi would secretly obtain a family planning injection at a clinic far from her village, where it leaves no visible sign and no one, not even her husband, could find out (Appendix D). By removing all observable social consequences, this question switches off both the conformity pressure ($\alpha \cdot E[y_{-i}]$) and the leader signal (γR_i). Only the internal channels remain,

so the question isolates the respondent's belief about latent private demand for FP. The question is administered once per respondent, on one randomly selected story, and provides a between-respondent benchmark on β that complements its identification from the profile randomization.

Figure 9: The Four Boxes of the Pebble-Allocation Game



NOTE: The four picture cards used in the first-person pebble-allocation game. Each respondent divided 12 pebbles across the four boxes, her own preference, scripture, the community, and her religious leader, to show how much each force weighs on her own family planning decision. Section 5.3 compares these stated weights against the weights the vignettes reveal.

A final instrument turns the question inward. Where the vignette measures the forces on Kendi indirectly, the pebble-allocation game asks each respondent directly what drives her own FP decision. She is given 12 pebbles and four labeled boxes, shown in Figure 9, standing for her own preference, scripture, the community, and her religious leader, and she distributes the pebbles to reflect how much each force weighs on her choice. Because this elicits in the first person the same drivers the vignette recovers indirectly, the gap between the two becomes a direct measure of social-desirability bias, which I exploit in Section 5.3. As with the other tasks, the physical pebbles and picture cards keep the exercise concrete for respondents with limited numeracy.

4.4 Identification and Estimation

The structural model yields the following choice probability for individual i:

$$\Pr(y_i = 1) = \frac{1 + \tanh(\boldsymbol{x}_i\beta + \delta S_i + \gamma R_i + \alpha E_i)}{2}. \tag{10}$$

For respondent i and story k, the likelihood belief response $m_{ik} \in \{0, 1, \ldots, 10\}$ is interpreted as a stated belief about this probability: $\hat{p}_{ik} = m_{ik}/10$. Applying the inverse hyperbolic tangent to both sides yields the linearized estimating equation:

$$\tanh^{-1}(2\hat{p}_{ik} - 1) = \boldsymbol{x}_i\beta + \delta S_k + \alpha E_k + \gamma R_k + \nu_i + \eta_{ik}, \tag{11}$$

where ν_i is a respondent-level random effect and η_{ik} is vignette-level noise. This transformation is exact under the logistic error assumption of the model, so the estimated coefficients map directly to the structural utility weights, not to reduced-form averages. Responses at the boundary ($\hat{p}_{ik} = 0$ or 1) are recoded to 0.05 and 0.95 respectively before applying the transformation.

The estimated model relaxes two restrictions in (11), following the extensions of Section 3.3. It lets the leader and scripture effects differ by direction, replacing the symmetric γR_k and δS_k with separate loadings for support and opposition, and it lets the conformity weight fall with cue salience, so the norm term becomes $\alpha(R_k, S_k)\, E_k$:

$$\begin{aligned}\tanh^{-1}(2\hat{p}_{ik} - 1) = \underbrace{\boldsymbol{x}_i\beta}_{\text{Between}} + \underbrace{\gamma^+ \mathbb{1}[R_k{=}1] + \gamma^- \mathbb{1}[R_k{=}{-}1]}_{\text{Leader (asymmetric)}} + \\ \underbrace{\delta^+ \mathbb{1}[S_k{=}1] + \delta^- \mathbb{1}[S_k{=}{-}1]}_{\text{Scripture (asymmetric)}} + \underbrace{\alpha_0 E_k - \alpha_R |R_k| E_k - \alpha_S |S_k| E_k}_{\text{Norm with crowding-out}} + \\ \underbrace{\nu_i + \eta_{ik}}_{\text{Composite error}} \quad ,\end{aligned} \tag{12}$$

where $\nu_i \sim N(0, \sigma_\nu^2)$ captures stable unobserved heterogeneity in FP disposition across respondents and η_{ik} is within-person vignette-level noise. The structural parameters

$$\beta, \gamma^+, \gamma^-, \delta^+, \delta^-, \alpha_0, \alpha_R, \alpha_S$$

and variance components $(\sigma_\nu^2, \sigma_\eta^2)$ are estimated jointly by maximum likelihood, integrating out ν_i, with enumerator fixed effects and standard errors clustered at the respondent level. This random effects MLE framework exploits both within- and between-respondent variation, and uses each source for what it identifies. As a boundary-native robustness check I re-estimate the same specification as a fractional logit on the untransformed $\hat{p}_{ik}$ scale. A specification ladder shows that the model captures the systematic variation, with any residual departure from a fully saturated cell model diffuse and without a dominant omitted term.

Within-respondent variation, generated by the randomized assignment of (S_k, E_k, R_k) across stories, causally identifies every within-respondent parameter. Stories 1–4 isolate the first-order effects through single-factor variation. Story 4 against the Story 1 anchor identifies the leader loadings γ^+ and γ^-, and Story 2 the scripture loadings δ^+ and δ^-. The endorsement premium is therefore a within-respondent test of the symmetry restrictions $\gamma^+ = -\gamma^-$ and $\delta^+ = -\delta^-$. Story 3 identifies the vacuum conformity weight α_0. Stories 5 and 6 then recover the crowding parameters through two-factor variation, Story 5 (scripture $\times$ norm) identifying α_S and Story 6 (norm $\times$ leader) identifying α_R, each isolated by the orthogonality of the design. Story 7 (scripture $\times$ leader) tests for a leader-scripture interaction that the parsimonious model omits. Story 8 (omnibus) provides a Sargan-style over-identification test, in which the predicted index from the additive speci-

fication should equal the observed response. Any systematic deviation indicates misspecification of the linearity embedded in the $\tanh^{-1}$ transformation.

Because Kendi is the decision-maker in each vignette, $\boldsymbol{x}_i$ collects *Kendi's* profile attributes (age, marital status, and parity), which are randomly assigned across respondents. β is therefore identified experimentally, from between-respondent variation in the randomized profile, and the four adjacent ladder contrasts recover the single-attribute effects. Two of these bundle their attribute with a co-moving feature. The marriage step (B→C) also introduces a husband and a joint-decision frame, and the parity step (C→D) also shifts the story to a birth-spacing frame. Since this parameter is identified between rather than within respondents, β is estimated less precisely than the within-respondent parameters. The heterogeneity extension of Section 3.4 interacts the vignette factors with *respondent* characteristics Z_i (age, education, parity, religiosity, and marital status), which are identified from the same within-respondent variation and add no new estimation complications.

5 Results

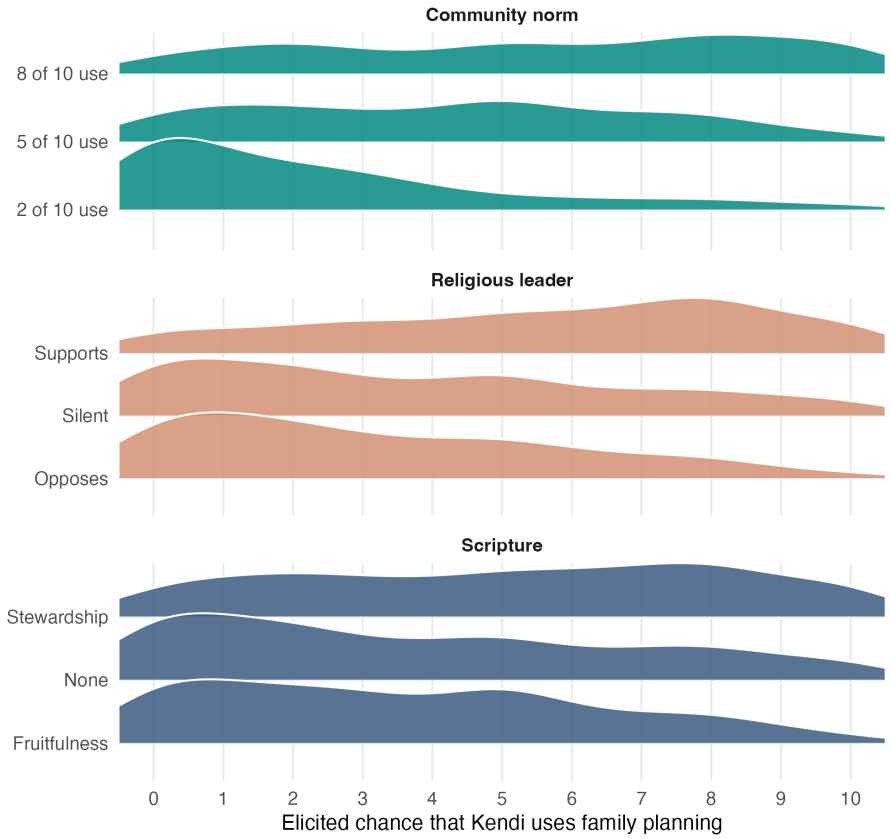
This section estimates the model on the vignette data and reads its parameters as a description of how one family planning (FP) decision responds to three social forces. I first recover the conformity, leader, and scripture weights from the randomized factors and place the community in the sub-critical regime, where private conviction still governs behavior and no single force can tip the decision on its own (Section 5.1). I then relax the two restrictions the design can lift, symmetric authority and a fixed conformity weight. A religious leader's endorsement moves a woman far more than an equal prohibition, and its power grows precisely where the community norm is most set against her (Section 5.2). A supplementary question about secret provision then shows that this conviction is real but understated, and that women act on a private demand they will not state aloud. Religious leaders read that demand accurately, while community and scriptural reports distort it (Section 5.3). The three weights are common to women across age, education, wealth, and parity, and to Kendi across her own life stages, so the mechanism is broadly shared, even as the strength of conformity itself varies with income and denomination (Section 5.4). A final subsection shows the estimates are invariant to the estimator, the functional form, and the mode of inference (Section 5.5).

5.1 Structural Estimates

The randomized factors move stated behavior in the raw data, before any model is imposed. Figure 10 plots the full distribution of the elicited chance across the levels of each factor. In every panel the distribution slides toward higher chances as the community norm rises, the religious leader turns supportive, or the scriptural frame turns toward provision. The structural model

attaches weights to exactly these patterns.

Figure 10: Raw Belief Distributions by Factor Level



NOTE: Distribution of the elicited chance that Kendi uses family planning, on the 0 to 10 scale, across the three levels of each randomized factor, pooled over the 5,080 vignette responses. Within every panel the distribution shifts toward higher chances as the community norm rises, the religious leader turns supportive, or the scriptural frame turns from fruitfulness (Genesis 1:28) to stewardship (1 Timothy 5:8).

Table 3 reports the structural weights from the simplified model, in which the conformity weight is constant and the leader and scripture effects are symmetric. The three social weights are all precisely estimated and clearly ordered. Conformity to the peer norm carries the most weight, $\alpha \approx 0.70$, ahead of deference to the religious leader, $\gamma \approx 0.31$, with the scriptural frame smallest, $\delta \approx 0.22$. The peer norm is the dominant channel moving a respondent's belief, over twice the weight of the leader and three times that of scripture. This is the ordering that a first-person elicitation will later invert (Section 5.3).

Because the conformity weight lies below one, the community has a single equilibrium (Proposition 1), so within this sample a leader bends the norm without tipping it. Conformity is nonetheless strong enough to matter. Evaluated at the point of maximum ambiguity, $\bar{y} = 0$, the social multiplier is $M = 1/(1-\alpha) \approx 3.3$, so a shift in the deterministic utility is amplified about threefold as neighbors respond to one another. The sample sits inside the sub-critical regime, so a trusted

Table 3: Structural Estimates: Simplified Model

Panel A. Social influence weights	
Conformity to peer norm (α)	0.70***
	(0.02)
Religious leader (γ)	0.31***
	(0.01)
Scripture (δ)	0.22***
	(0.01)
Panel B. Private preference (β): lifecycle ladder contrasts	
Age (16 to 22)	0.08
	(0.05)
Marriage	0.01
	(0.04)
First child (parity 0 to 1)	0.51***
	(0.05)
Age (22 to 28)	-0.09*
	(0.05)
Social multiplier M ($\bar{y} = 0$)	3.29
Respondent SD	0.32
Residual SD	0.56
Respondents	635
Stories	5,080

NOTE: Random-effects maximum-likelihood estimates of the simplified additive model (constant conformity weight, symmetric leader and scripture effects), with a respondent random intercept, enumerator fixed effects, and cluster-robust standard errors (at the respondent) in parentheses. The outcome is $\tanh^{-1}(2\hat{p} - 1)$. Panel B reports adjacent ladder contrasts, each isolating one lifecycle attribute relative to the previous profile (base A: 16, unmarried, 0 children). The social multiplier is $M = 1/(1 - \alpha)$, evaluated at $\bar{y} = 0$. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

authority moves behavior at the margin without triggering a self-reinforcing cascade. That margin is narrower than the averaged figure suggests, and Section 5.2 shows conformity runs stronger still once the averaging over cues is undone.

The private-preference estimates trace a sharp lifecycle pattern (Panel B). Reading up the profile ladder, a respondent's judged disposition toward FP barely moves as a woman ages from 16 to 22 (+0.08) and does not move at all when she marries (+0.01). It then jumps with the arrival of her first child, by +0.51. Marriage alone changes little, but a child changes everything, consistent with a birth-spacing motive that activates once a woman has a child to space from. By design this step also introduces the birth-spacing frame (Section 4.4), so it captures parity and the spacing motive together. As the between-respondent arm, these weights are estimated less precisely than the within-respondent parameters, yet the parity contrast is large and sharply identified.

5.2 How Religious Authority Works

The simplified model of Section 5.1 imposes two restrictions the design can relax, that the leader and scripture effects are symmetric and that the conformity weight is fixed. Neither holds, and relaxing them is where the mechanism of authority comes into view.

Table 4: Structural Estimates: Full Model

Panel A. Direct authority effects (vs. silence)	
Leader: supports (γ^+)	0.50***
	(0.02)
Leader: opposes (γ^-)	-0.12***
	(0.02)
Scripture: pro (δ^+)	0.38***
	(0.02)
Scripture: anti (δ^-)	-0.06***
	(0.02)
Panel B. Norm and crowding-out	
Conformity, no authority (α_0)	0.90***
	(0.04)
Leader salience (α_R)	0.15***
	(0.03)
Scripture salience (α_S)	0.24***
	(0.04)
Endorsement premium, leader	0.38
Endorsement premium, scripture	0.31
Respondent SD	0.32
Residual SD	0.54
Respondents	635
Stories	5,080

NOTE: Random-effects maximum-likelihood estimates of the full model, with a respondent random intercept, enumerator fixed effects, and cluster-robust standard errors (at the respondent) in parentheses. Direct authority effects are relative to a silent leader and no scripture. The crowding parameters give $\alpha(R,S) = \alpha_0 - \alpha_R|R| - \alpha_S|S|$, so a larger α_R or α_S means a speaking authority displaces more of the peer norm. The endorsement premium ($\gamma^+ + \gamma^-$ and $\delta^+ + \delta^-$) is zero under a symmetric response and is rejected in both channels ($p < 0.001$). Private-preference (profile) controls are included as in Table 3. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The norm channel comes first. With no authority speaking, the conformity weight is $\alpha_0 = 0.90$ (Table 4, Panel B), somewhat larger than the constant-weight estimate of 0.70 in Section 5.1 because it isolates conformity in the authority vacuum, with no averaging over the cues. This is a sterner benchmark than the averaged weight lets on. In the authority vacuum the social multiplier is $1/(1-\alpha_0) \approx 10$, close to the tipping threshold, so the peer norm is nearly self-reinforcing on its

own. When a cue does speak, that weight falls. Each unit of leader salience reduces it by $\alpha_R = 0.15$ and each unit of scriptural salience by $\alpha_S = 0.24$, both estimated with precision and jointly far from zero. A leader or a verse thus displaces the peer norm rather than adding to it. Authority substitutes for conformity, and the substitution is the same whether the cue points for or against, because what loosens the norm is the cue's salience, whatever direction it points.

That loosening has a direction of consequence, even though it does not depend on the direction of the cue. Because pulling behavior off the norm moves it toward the cue's own message, an authority does the most where it fights the prevailing norm and the least where it echoes one. Figure 11 shows the pattern in the raw responses. A supportive leader lifts the elicited chance most where few women in the village use FP, and by less where most already do. An opposing leader barely moves behavior against a low-use norm and suppresses it most where use is widespread. Scripture behaves the same way. A leader's leverage is therefore greatest against a strong, misaligned norm and weakest reinforcing an aligned one.

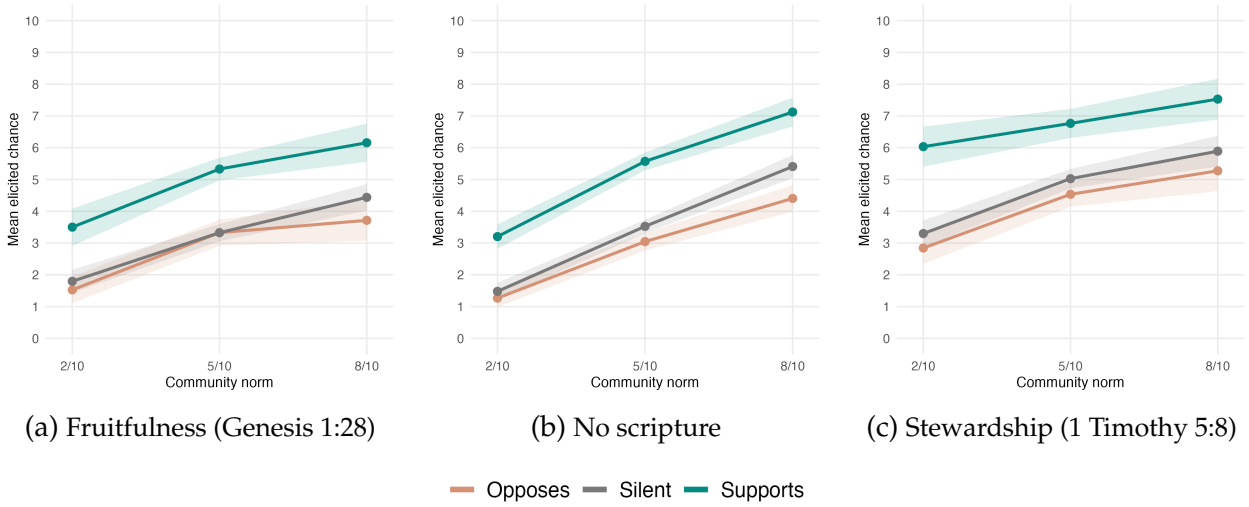
Authority is lopsided in a second, separate way. Holding the norm fixed, a supportive leader raises the chance far more than an opposing leader lowers it, $\gamma^+ = 0.50$ against $\gamma^- = -0.12$ (Panel A), and scripture shows the same tilt, $\delta^+ = 0.38$ against $\delta^- = -0.06$. The symmetric restriction that a blessing and a prohibition are mirror images is rejected in both channels ($p < 0.001$). A blessing is the far stronger lever. I call this the endorsement premium. The premium is robust across the atanh, fractional-logit, and binomial estimators, and focus-group participants judged the supportive and opposing cards to carry equal weight, so it reflects how women respond rather than an unbalanced manipulation.

The premium follows from the same suppressed demand the secret question uncovers. A supportive cue does not change what a woman wants. It licenses her to act on a preference the community norm holds down, and because that suppressed demand is widespread a blessing reaches many women, while a prohibition meets little demand left to suppress and reaches few. The heterogeneity bears this out. A supportive scriptural cue moves a woman more, not less, the more her own reading of scripture already favors provision. The effect grows by about 0.07 per standard deviation of her scriptural orientation. Were the premium instead a matter of correcting a presumed pronatalist scripture, the supportive cue would land hardest on the women who most lean pronatalist. I find the reverse. A blessing works by resonating with a conviction a woman already holds.

These pieces describe a single mechanism. Because a speaking authority displaces the peer norm, it lowers the social multiplier. A leader therefore amplifies a shift less than the constant-weight benchmark of Section 5.1 implies, and tips no community merely by talking. A leader coordinates instead. By loosening a misaligned norm and supplying a supportive frame, the leader moves a community toward the use level its members would privately choose, without changing those

private preferences. That is why the same leader matters most where the norm is entrenched and the demand it holds down is already high.

Figure 11: Elicited Chance by Community Norm, Leader Stance, and Scripture



(a) Fruitfulness (Genesis 1:28) (b) No scripture (c) Stewardship (1 Timothy 5:8)

NOTE: Mean elicited chance that Kendi uses family planning (0 to 10) against the community norm, by the religious leader's stance (see legend), with 95% confidence bands. The common upward slope is peer conformity. Within every scriptural frame a supportive leader raises the chance markedly while an opposing leader sits only modestly below a silent one, the endorsement premium. The overall level rises from the fruitfulness frame to the stewardship frame, the direct effect of scripture.

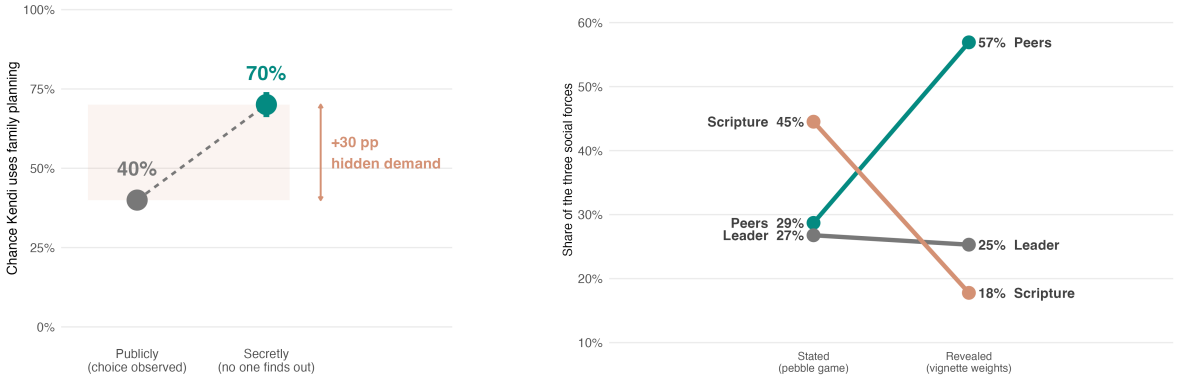
5.3 Latent Demand and the Self-Report Inversion

The results so far take latent demand as given. A supplementary question shows why. After selected vignettes, respondents are asked whether Kendi would secretly obtain a FP injection at a clinic far from her village, where it leaves no visible sign and no one, not even her husband, could find out. On the same scenarios, women judge Kendi's chance of using FP at 40% when her choice is observed, while 70% say she would obtain the secret injection (Figure 12a). Because the secret framing removes every observable consequence, it switches off both the conformity term αE and the leader signal γR and leaves only Kendi's private demand. The 30-point gap is therefore demand that observation suppresses. The force holding use down is social, not an absence of private conviction. That is exactly the premise the structural model was built to test.

That social constraint is also the force women are least willing to name. A first-person pebble game asks each respondent to divide 12 pebbles across the drivers of her own FP decision, her own conviction, scripture, the village, and her religious leader. Women assign most of their pebbles to their own conviction (52%) and, among the social forces, name scripture the loudest (21%), with the village (14%) and the leader (13%) below it. The randomized vignettes reverse this order among the three social forces (Figure 12b). The peer norm is by far the strongest, ahead of the

leader and then scripture. A woman therefore ranks peer influence at the bottom of what she says drives her, even though it is the force that moves her most.[6]

Figure 12: Latent Demand and the Self-Report Inversion



(a) Hidden demand (the secret wedge)

(b) The self-report inversion

NOTE: Panel (a) plots the mean elicited chance that Kendi uses family planning when her choice is observable against the share of respondents who say she would obtain the injection in secret, on the 635 stories carrying the secret-access question, with 95% confidence intervals. The secret framing removes every observable consequence and isolates private demand. Panel (b) plots the relative importance of the three social forces, as a share of the pebbles a woman allocates to them in the first-person game (stated) and as a share of the estimated structural weights α, γ, and δ (revealed). Own conviction, the largest single pebble category at 52%, is the private-demand channel of panel (a) and is omitted here so the two sides are comparable. Women name scripture the loudest of the three social forces, but the vignettes reveal the peer norm to be by far the strongest, inverting the order.

The inversion is not measurement error but a signature of the forces under study. It is selective. Comparing the stated and revealed shares of the three social forces (Figure 12b), women understate the village and overstate scripture, while their reported deference to the leader is about right. The bias falls precisely on the channel that conformity makes costly to admit. Naming peer pressure is itself a violation of the norm that peers enforce, so the same pressure that makes the peer norm powerful makes women unwilling to attribute their choices to it. Direct questions about social influence are therefore not merely noisy. They are systematically inverted. Recovering the true weights takes a third-person vignette.

A companion question sharpens the same point from a different angle. Alongside what Kendi *would* do, each respondent says whether she personally agrees that Kendi *should* use FP. Re-estimating the weights on this normative measure reproduces the ordering, the endorsement premium, and the crowding (Appendix Table A1), but conformity binds approval markedly less than behavior,

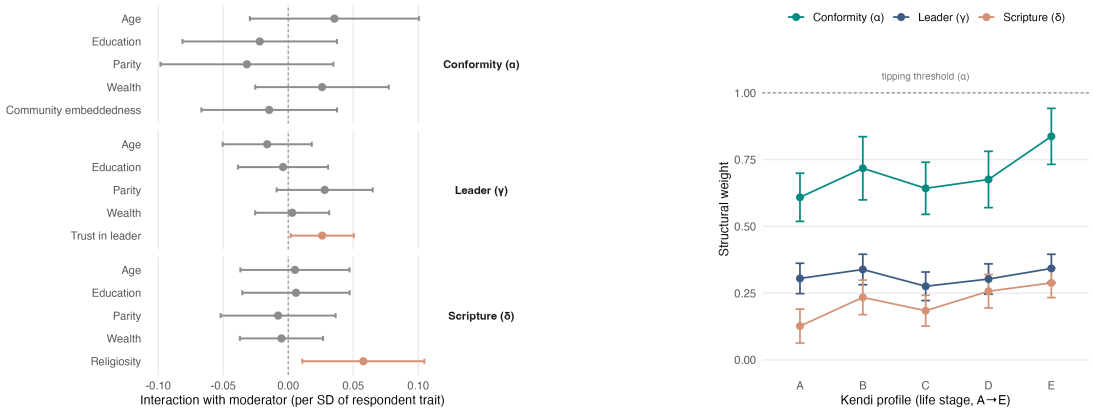
[6] A separate first-person battery points the same way. Each woman was asked how likely she would be to use FP under stated combinations of how many friends use it and what her religious leader says. Moving from no friends to half of them using it, with the leader silent, raises her stated likelihood by about as much as moving the leader from silent to approving (0.09 versus 0.10 on the five-point scale). The vignettes put the peer weight at more than twice the leader's. A ceiling blunts the battery, since most women answer near the top throughout, but asked directly they again do not register peers as the force that moves them most.

with a weight near 0.68 against 0.90. A woman's judgment of what is right is swayed less by what her neighbors do than by what she expects herself to do. The peer norm therefore presses on behavior more than on conviction. That gap is the same wedge the secret question opens in levels.

5.4 Heterogeneity

Across Respondents A single set of pooled weights should not hide sharp differences across women. The design lets every weight vary with a respondent's own traits (Section 3.4), and Figure 13(a) plots each interaction in one block per weight. Neither conformity, nor the leader effect, nor the scripture effect moves with a woman's age, education, parity, or household wealth, and every one of these interactions is indistinguishable from zero. Tested jointly across the three weights, marital status ($p = 0.15$) and employment ($p = 0.13$) leave them unchanged as well. Two traits are the exception, and both act on the same weight. Household income and religious denomination each move the conformity weight (income $p < 0.001$, denomination $p = 0.008$). I take up that pattern after the direction-of-movement results below.

Figure 13: Heterogeneity in the Structural Weights



(a) By respondent trait

(b) By Kendi's life stage

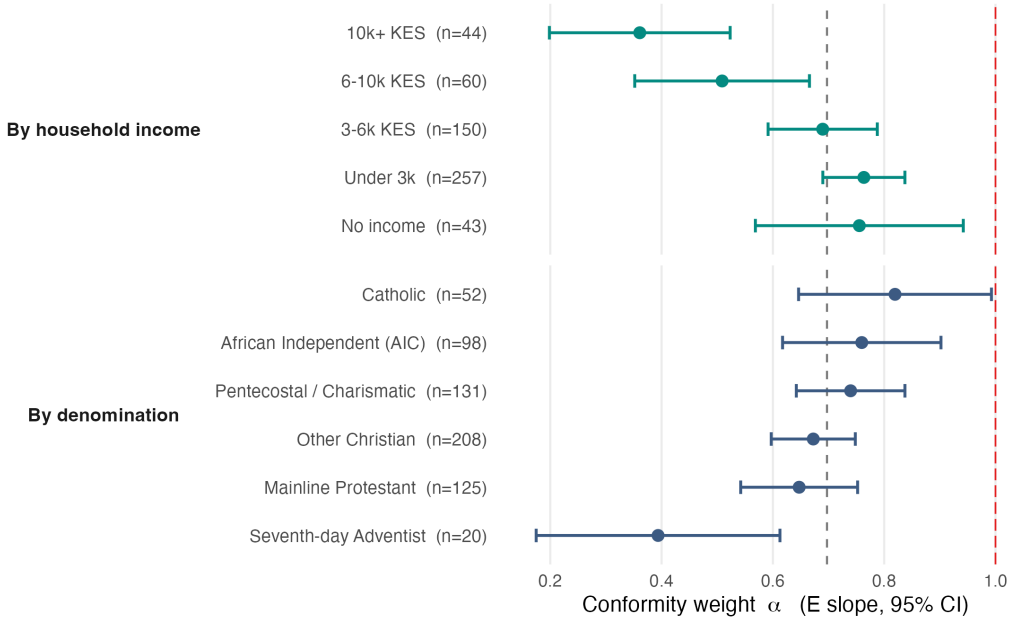
NOTE: Panel (a) plots the interaction of each structural weight (conformity, leader, scripture) with a battery of respondent traits, one estimate per row with 95% confidence intervals; a positive value means the weight is larger for higher values of the standardized trait, and estimates whose interval excludes zero are in orange. Each block pairs four demographic controls with one theory-matched disposition: community embeddedness for conformity, trust in religious leaders for the leader effect, and religiosity for scripture. Panel (b) plots all three weights estimated separately for each of Kendi's five lifecycle profiles (A–E), with a dashed line at the tipping threshold of one for the conformity weight. All estimates come from random-effects maximum-likelihood models with a respondent random intercept and enumerator fixed effects.

Where a weight moves with a disposition, it moves as the model predicts. Scripture weighs more heavily on the more religious, at about 0.06 per standard deviation of religiosity. The leader's effect shows a weaker tendency of the same kind and rises with a woman's trust in religious leaders, though that interaction is small and less firm across specifications. Conformity is flat against every

disposition of this kind. It is equally unmoved by direct measures of how embedded a woman is in her community, her membership in a local women's group, her years in the village, and the size of her close-female-friend network. The pull of religious authority is personal, concentrated among the devout and the trusting.

The pull of the community norm is broad but not uniform. Conformity tracks no single demographic gradient, yet it varies systematically with a woman's economic resources and her religious tradition (Figure 14). It falls almost monotonically with household income, from about 0.76 among women with little or no cash income of their own to 0.36 in the highest band. A woman with more money of her own is markedly less bound by what her neighbors do. It also differs across denominations, highest among Catholics, near the pooled level for the larger Protestant and Pentecostal groups, and lowest among Seventh-day Adventists, though the smallest traditions are estimated imprecisely. Susceptibility to the peer norm is common in kind but uneven in degree. This foreshadows the policy reading of Section 7, that a leader's leverage is greatest where conformity is strongest.

Figure 14: Conformity by Household Income and Denomination



NOTE: The conformity weight α, the slope of the elicited chance on the community norm, estimated separately by household income band and by religious denomination, with 95% confidence intervals and group sizes. The dashed gray line marks the pooled estimate ($\alpha \approx 0.70$) and the red line the tipping threshold of one. Income and denomination each move the conformity weight (likelihood-ratio tests, $p < 0.001$ and $p = 0.008$), and the smaller traditions, notably the Seventh-day Adventists ($n = 20$), are estimated imprecisely. Estimates come from random-effects maximum-likelihood models with a respondent random intercept and enumerator fixed effects, on the 634 Christian respondents.

Across Kendi's Life Stages How the story's own profile matters is a question about who is reached. Figure 13(b) estimates all three weights separately for each of Kendi's five lifecycle pro-

files. The conformity weight stays between 0.61 and 0.84, below the tipping threshold of one for every profile, so no Kendi type carries a self-reinforcing norm. The jump in private preference at the first child (Section 5.1) is a shift in the level of demand, not a change in the multiplier. The leader's weight is flat, close to 0.30 from the unmarried 16-year-old to the older mother, so their authority reaches every life stage alike.

Scripture behaves differently. Its weight more than doubles over the lifecycle, from 0.13 on the unmarried 16-year-old to 0.29 on the 28-year-old mother, with confidence intervals that do not overlap. The gradient tracks age and childbearing rather than marriage itself, since the cues speak to raising and providing for children. They weigh most on an older woman with children and barely touch a girl whose circumstances they do not yet address. The two authorities reach women in different ways. The leader carries personal standing over the whole congregation and moves everyone by about the same amount, while scripture moves a woman only insofar as its teaching speaks to her situation. This is the one place where the profile of the story reshapes a weight, and it does so exactly as the content of the cue would predict.

5.5 Robustness

The estimates do not depend on how the elicitation is scaled or which respondents are kept. Table A2 re-estimates the full model five ways. A fractional logit that treats the raw 0 and 10 endpoints natively, with no recoding, and a binomial random-effects model that keeps the count structure both reproduce the headline weights once placed on the same scale. Restricting the sample to the 93% of respondents who correctly recall the manipulated factor leaves every weight essentially unchanged, as does adding fixed effects for the randomized order in which the eight stories appeared. That slot is itself jointly insignificant ($\chi^2(7) = 7.8$, $p = 0.35$), so nothing in the estimates reflects fatigue or a story's position in the sequence.

The parsimonious specification is the one the data support. Table A3 walks a nested ladder. The salience terms that make the crowding asymmetric are strongly needed, while the direction-specific bilinears a richer model would add are jointly insignificant ($\chi^2(2) = 2.05$, $p = 0.36$). Against a fully saturated model that gives all 27 factor cells their own parameter, the parsimonious specification leaves only a small residual spread diffusely across 19 terms, with no systematic interaction. The randomization that underwrites this is documented in Section 4.2.

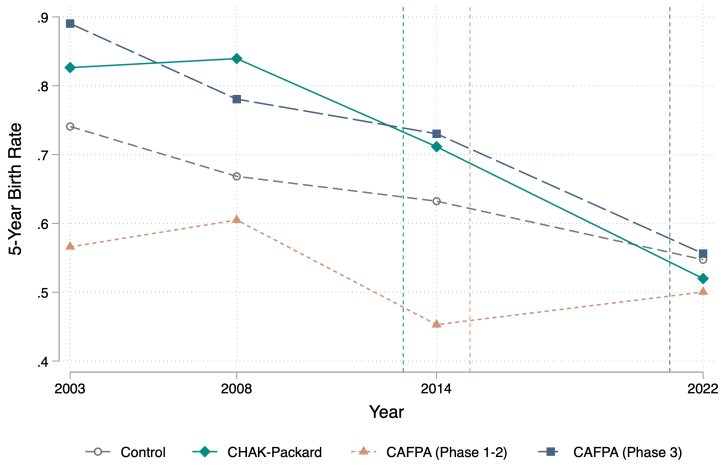
Inference holds under a coarser clustering. The headline standard errors cluster at the respondent, the level of randomization, but leaders and norms are also shared within a village. Recomputing the key tests with a wild cluster bootstrap over the 27 villages, the appropriate procedure when clusters are few, leaves the argument intact. The leader-salience term that carries the against-the-choir result stays significant at the village level ($p = 0.01$), and conformity, scripture salience, and the endorsement premium hold at $p < 0.001$. Only the small anti-scripture effect softens to

marginal significance.

6 External Validity: Evidence from Faith-Based Programs

Setting and Programs. Between 2013 and 2022, CHAK, with CCIH, ran two faith-based family planning (FP) programs that trained religious leaders as advocates (Bormet et al., 2021; Ruark et al., 2019). They shared that training but differed in what backed it. CHAK-Packard, concentrated in the western Lake Victoria basin, coupled leader engagement with community health volunteers and its own network of faith-based facilities. CAFPA, dispersed across the Central and Rift Valley regions, coupled the same engagement with county-government policy advocacy and no parallel community-distribution arm. Appendix B documents the two programs, their geography (Figure B1a), and the data.

Figure 15: Five-Year Birth Rate by Program Group, 2003–2022



NOTE: Survey-weighted county-group means of the five-year birth rate across the four KDHS waves. Groups are CHAK-Packard counties (Siaya, Kakamega, Vihiga, Homa Bay), CAFPA Phase 1–2 counties (Kiambu, Meru, Murang'a), CAFPA Phase 3 counties, and never-treated control counties, with Nairobi excluded throughout. Vertical lines mark program inception. The CAFPA Phase 1–2 series falls below trend at the 2014 baseline and partially recovers by 2022. Source: author's calculations from the KDHS.

Design. I evaluate the programs with a difference-in-differences design on four waves of the Kenya Demographic and Health Survey, identified from the change between 2014 and 2022. Because individual leaders are not observed, treatment is anchored to the church-affiliated health facilities where leader training took place. I measure exposure three ways, a binary indicator for program counties, a ward proximity tier, and a continuous within-county distance gradient (Appendix Equation B4). The religious moderators, baseline trust in leaders and the density of faith-based infrastructure, come from a ward-level small-area estimator (Appendix B.3). Because only three to four counties are treated, I rely throughout on inference robust to a small number of

clusters.

Findings. The form of engagement decides whether it works (Table 5). CHAK-Packard lowered the five-year birth rate by roughly eight percentage points. The reduction is not an artifact of the few treated counties, it holds under the synthetic control that reweights the pre-period, grows under the Callaway–Sant'Anna estimator that handles the staggered rollout, and reappears at $p <$ 0.001 as a within-county distance gradient. CAFPA carries the opposite sign, and the increase persists across the same estimators, including the synthetic control that reweights the pre-period. I do not read it as a program effect. CAFPA's counties fall below trend at the 2014 baseline and partially recover by 2022 (Figure 15), a movement consistent with the unwinding of a concurrent facility-strengthening program. With only two sparse earlier waves, no estimator can separate that recovery from a true increase.

Table 5: External Validity: Faith-Based Program Effects on Fertility and Its Channels

Panel A. Average effect on the five-year birth rate			
	County TWFE	Synthetic DiD	Callaway–Sant'Anna
CHAK-Packard	-0.077*	-0.096**	-0.114***
CAFPA Phase 1–2	0.102**	0.085**	0.136***
Panel B. Mechanism: moderated effects by outcome (per SD)			
	Births	Modern CPR	Unmet need
	CHAK-Packard		
FBO×DiD (delivery)	-0.018	0.022**	-0.022
Trust×DiD (authority)	-0.063***	0.023	-0.040**
MoH×DiD (falsification)	-0.029	-0.003	-0.016
	CAFPA Phase 1–2		
FBO×DiD (delivery)	-0.025**	-0.002	0.009
Trust×DiD (authority)	0.046*	0.014	0.006
MoH×DiD (falsification)	0.016	-0.001	-0.006

NOTE: Panel A reports the treatment effect on the five-year birth indicator from three estimators. *County TWFE* is the two-way fixed-effects 2 × 2 DiD (2014 vs. 2022); with only three to four treated counties, its significance stars are from the wild cluster bootstrap (Rademacher, 999 replications, clustered on county; p = 0.06 for CHAK, 0.02 for CAFPA). *Synthetic DiD* (Arkhangelsky et al., 2021) reweights control units and pre-periods, with adoption timed to each program's rollout (CHAK Phase 1 from the 2014 wave; CHAK Phase 2 and CAFPA from 2022) and significance from its placebo bootstrap. *Callaway–Sant'Anna* reports the overall ATT with cluster-robust (influence-function) standard errors. Panel B reports the coefficient on the triple interaction of treatment, post, and a standardized ward-level moderator: faith-based-organization health facilities (delivery), baseline trust in religious leaders (authority), and Ministry-of-Health facilities (falsification), with analytic standard errors clustered at the county. Nairobi is excluded throughout. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The mechanism follows the coordination model (Table 5, Panel B). Where faith-based delivery is denser, CHAK raises modern contraceptive use and CAFPA's upward drift in births is damped. Where trust in religious leaders is higher, each program's own effect is magnified. The density of

ordinary Ministry-of-Health facilities moves none of these outcomes, so the operative infrastructure is the faith-based network specifically, not general health-system capacity. Religious leaders move behavior on the ground when their message is bundled with a means to act on it.

Two features of the design temper these estimates. With three to four treated counties, conventional standard errors understate uncertainty, which is why the wild cluster bootstrap and synthetic control carry the inference, and a Goodman–Bacon decomposition confirms the estimate is effectively a clean two-by-two (Appendix B.6). A strong test of parallel trends is also infeasible, because the survey offers only two pre-treatment waves, sparsely sampled and unevenly covered, and 2014 is itself mid-treatment for CHAK's first phase. The program evidence therefore corroborates the structural estimates rather than identifying the mechanism on its own.

7 Discussion

This paper set out to separate three forces that observed behavior blends together and to ask whether a religious leader changes what women want or coordinates what they already do. The estimates answer plainly. Latent demand for family planning (FP) is high, the binding constraint is social, and a leader works by coordinating that suppressed demand (Section 5.2). This section draws out what that answer implies. I first put the estimated weights to work in counterfactual simulations of how far a leader can move a community and whether the movement lasts. I then turn to the design of leader engagement, its extension to other norm-governed behaviors, and the limits of the evidence.

7.1 Counterfactual Policy Simulations

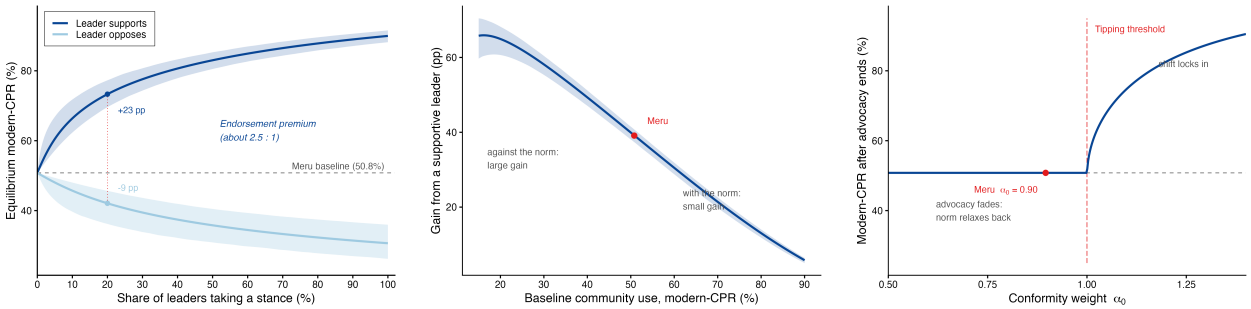
The structural weights describe how one woman's decision responds to conformity, authority, and scripture. Pushing them through the model's equilibrium map, $\bar{y} = \tanh(h + \alpha\bar{y})$, translates that private response into the community rate it implies, and because neighbors react to one another a change in any one force moves the aggregate by more than its direct effect. I anchor the exercise to Meru, whose modern contraceptive prevalence in the 2022 DHS was 50.8 percent, almost exactly the point of maximum ambiguity where conformity is most powerful. The question is how far a religious leader's stance can move this community and whether the movement lasts.

Figure 16(a) traces the equilibrium prevalence as the share of leaders taking a stance rises, and the response is both strong and asymmetric. If one in five leaders endorsed provision as clearly as the supportive cue in the vignette, the model implies equilibrium use in Meru would rise by about 23 points, with a 95% bootstrap interval of 19 to 27. The same reach of opposition would lower it by only nine, roughly two and a half to one. This is the endorsement premium of Section 5.2 carried to the community level, and conformity amplifies it. At the point of maximum ambiguity the silent-

leader multiplier is close to 10, so even a modest direct push is magnified. When the leader speaks the multiplier falls to about four, and to about two once a scriptural frame is invoked as well, because a speaking authority displaces part of the peer norm. The leader gives up amplification for a direct push. That is the price of coordinating a community.

Where a leader is worth most is a separate question from how far reach carries, and Figure 16(b) answers it. It plots the gain from a fully supportive leader against the community's own baseline use, and the gain is largest exactly where use is lowest, in the communities whose norms most oppose provision. A supportive leader adds about 40 points at a baseline near Meru's and more than 60 where use sits far below it, but only a handful of points where the norm already favors provision. This is the asymmetry that gives the paper its title, now stated as a targeting rule. A leader accomplishes the most preaching against an entrenched norm and the least preaching to one already aligned, so the lever is worth engaging precisely where the social constraint binds hardest. The crowding follows the same logic, since a leader who displaces a hostile norm gains more than one who merely echoes a favorable one.

Figure 16: Counterfactual Simulations of Religious Authority



(a) Endorsement premium (b) Preaching against the choir (c) Durability / tipping threshold

NOTE: All three panels push the estimated structural weights through the equilibrium map $\bar{y} = \tanh(h + \alpha\bar{y})$, with the private-preference intercept calibrated to Meru's 2022 modern contraceptive prevalence of 50.8 percent (KDHS, women's sample weights). Panel (a) plots equilibrium prevalence against the share of leaders taking a stance, for a supportive and an opposing cue, with the crowding of the conformity weight scaled to that share; the marked points give the effect at a one-in-five reach. Panel (b) plots the gain from a fully supportive leader against the community's baseline use. Panel (c) plots the prevalence the community settles at after a supportive campaign is withdrawn, against the conformity weight α_0. Shaded bands are 95% intervals from a parametric bootstrap over the estimated weights' joint sampling distribution.

These are upper bounds on what authority can do, not forecasts of what a program will deliver. They assume every reached woman hears a stance as salient as the vignette's and that the endorsement stays in place. The shaded bands, a parametric bootstrap over the joint sampling distribution of the estimated weights, are narrow throughout, so the curves are precise given the model. The uncertainty that should give pause is not statistical but external, whether a real sermon carries the salience of the vignette's cue. The second assumption is where the model draws its sharpest line.

Figure 16(c) follows the same community after the advocacy fades. Because Meru's conformity weight, $\alpha_0 = 0.90$, sits just below one, the community has a single equilibrium, so once the leader falls silent the norm relaxes back toward where it began and the gain is transient. Only when conformity crosses one does a second, high-use equilibrium open, and a sufficient push then locks the community into it after the advocacy ends. Meru sits just short of that threshold. The lever bends the norm firmly while it is applied but does not, on its own, tip the community into a new one.

This distinction between a large but transient response and a durable one organizes the policy discussion that follows. It separates what a burst of preaching can accomplish from what a standing religious presence can, and it points to the communities where engagement is most likely to endure.

7.2 Policy Implications

The counterfactuals turn a general call to engage religious leaders into a specific one. A leader is a coordinating device, so engagement pays off where coordination is strong and where the norm most needs moving. Three conditions follow from the estimates. Engagement works best where conformity is strong, since the multiplier that carries a leader's push is largest there and the community sits closest to a tipping point. The heterogeneity in Section 5.4 locates this concretely, with conformity running strongest among poorer women. It works best where trust in religious leaders is high, since that is what gives a leader's stance its weight. And it works best where use is low, because that is where the norm most opposes provision and a supportive leader has the most to displace (Figure 16(b)). The same effort buys a far larger shift in a high-conformity, high-trust community whose norm resists provision than in one where use is already common or where leaders carry little weight.

A fourth condition concerns time. Because the conformity weight I estimate sits just below the tipping threshold, a leader's push bends the norm while it is applied but does not hold once it stops (Figure 16(c)). Advocacy that arrives as a burst therefore fades, and a program evaluated years later can show little even after moving behavior while it ran. The two Kenyan programs bear this out. Both trained religious leaders and both carried a real message into their communities, but they differed in whether that message had a standing presence behind it. CHAK-Packard paired its leaders with community health volunteers and a network of faith-based facilities that kept the endorsement live and gave women a place to act on it, and its effect on fertility persisted to the 2022 survey. CAFPA carried religious-leader messaging and county policy advocacy but no parallel community-distribution arm, and no reduction appears in its counties. The contrast is not that one program preached and the other did not. It is that one built a presence that sustained the shift while the other relied on the preaching alone.

This points to a complementarity my structural model can only half describe. The model holds

access fixed, so it speaks to the demand side, how far authority can move what women want to do, and it is silent on supply. Yet the program evidence shows the heterogeneity concentrating in the faith-based delivery network rather than in general health infrastructure (Section 6), and a shifted norm is worth little if women cannot act on it. Engaging leaders and building the means to act on a changed norm are therefore complements, not substitutes. A blessing coordinates demand and a delivery network lets that demand be met, and neither accomplishes much without the other. My interviews with religious leaders in the study area suggest that both halves are within reach.[7]

How a leader frames the message matters as much as whether a leader speaks. Scripture is a strong force in the estimates, but it is not a lever a program can pull, since neither the text nor a congregation's reading of it can be reassigned. A leader can choose which reading to invoke. The resonance result of Section 5.2 shows that a scriptural endorsement moves a woman most when it aligns with a conviction she already holds. A leader therefore does best to frame provision in terms the congregation is already inclined to accept. Authority works by granting people permission to act on what they quietly want, and the framing that grants that permission most cleanly is the one that meets them where they already stand.

7.3 Generalization to Other Norm-Governed Behaviors

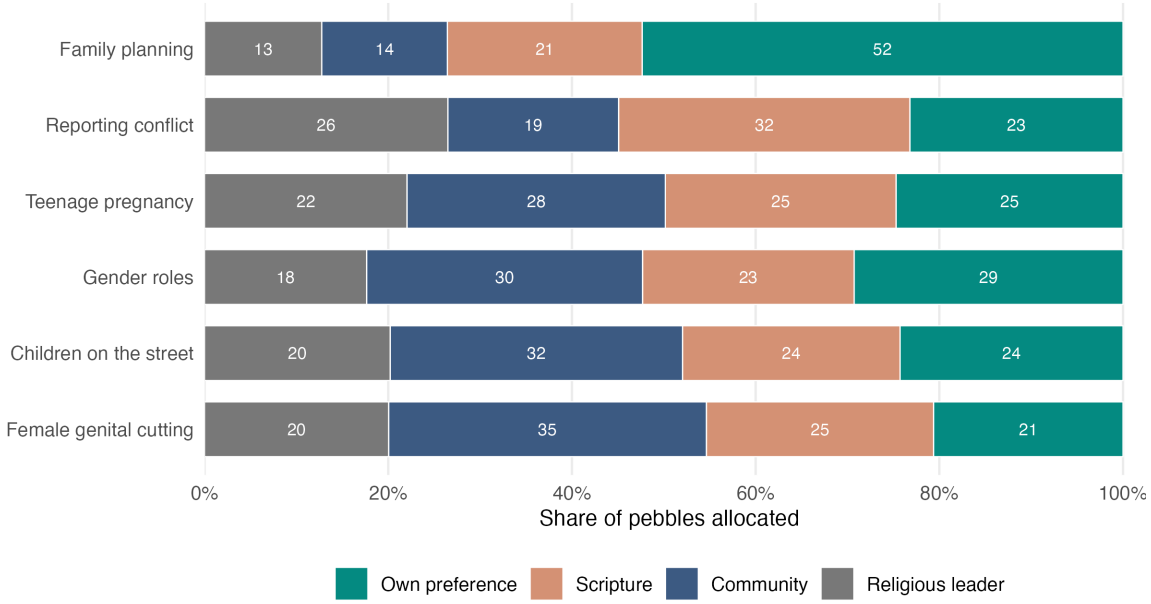
Nothing in the model is specific to FP. It describes any behavior an individual weighs against a community norm, a scriptural teaching, and a trusted authority, so it should travel to the other norms that hold women back in the region. Female genital cutting still affects one in seven women, and two in five ever-partnered women report violence from a partner (Section 2). These are no more matters of information or access than FP is. They persist as social norms.

The survey offers a first look at whether the same forces are at work. Alongside the FP game, each respondent divided pebbles over the same four forces for five other issues, from female genital cutting to gender roles to teenage pregnancy (Figure 17). The four forces organize every one of them, which is the first thing the framework requires. Their balance differs. FP is the outlier. It is the only issue a woman traces mostly to her own preference, and the only one where she gives the community the smallest share. For female genital cutting, gender roles, and children on the street, the community is the largest force. It draws more than twice the share it draws for FP.

The contrast sharpens the case for generalization. The FP pattern is the self-report inversion of Section 5.3, a woman privately owning a choice whose social constraint she will not name. For the other behaviors she names the community openly, so the social constraint the model places at the

[7] Alongside the women's survey, and over the same period, I interviewed 17 religious leaders in the Maua area of Meru County. Each interview ran about 30 minutes, was conducted in English, and paired a short structured questionnaire on traditional social norms with a closing open-ended question, so the evidence is both quantitative and qualitative. Given the small number of leaders, I read it as suggestive rather than representative.

Figure 17: The Same Four Forces Across Six Norm-Governed Behaviors



NOTE: Mean share of the 12 pebbles each respondent allocated across own preference, scripture, the community, and the religious leader, for six social issues. Bars sum to 100%. Family planning is the only issue a woman traces mostly to her own preference and the one where she gives the community the smallest share, the self-report inversion documented in Section 5.3. For the other five the community is among the largest forces. The allocations are first-person self-reports and are descriptive, the structural mechanism is estimated only for family planning.

center is, if anything, more visible there. A norm a community will openly enforce is one a trusted authority is well placed to move, and the leader who confronts it preaches against a louder choir than FP presents.

I do not push this further than the design allows. The pebble allocations are first-person self-reports, and the coordination mechanism is estimated only for FP. The cross-domain evidence shows that the ingredients the model needs are present across these behaviors, not that a leader in fact moves them. Testing whether the endorsement premium, the crowding, and the tipping threshold recur for female genital cutting or intimate-partner violence would take the same vignette design carried into those domains, which is where I would take this work next.

7.4 Limitations

Several limits bound these conclusions. The decomposition separates a woman's own preference, scripture, the community, and her religious leader, but it leaves out her husband, whom a separate battery suggests can outweigh all four. The model captures the social field around the decision more completely than the household within it. The vignette is also hypothetical and third-person, which guards against social-desirability bias but means the counterfactual magnitudes rest on a

cue that may be more salient than a sermon in the field, so I read them as bounds on potential. And the structural survey comes from a single county, so whether the same weights hold where conformity or trust differ is a question the design cannot settle.

The program evidence carries its own caveats, set out in Appendix B. With few treated counties and sparse, unevenly covered pre-treatment surveys, a strong test of parallel trends is infeasible, and CAFPA's counties differ enough from their controls that I place no causal weight on its estimate. The model is also static. It locates a community relative to the tipping threshold but does not observe how fast a norm adjusts or whether repeated advocacy loses force. The durability results describe which equilibrium a community settles at, not the pace or the path of getting there.

8 Conclusion

Across the developing world, religious institutions hold authority that often rivals that of the state, shaping culture, beliefs, and behavior. A central ambiguity in the economics of social norms is whether religious leaders merely echo the entrenched preferences of their communities, or whether their endorsement can move a binding peer norm. This paper grounds that question in the persistent challenge of reproductive health in sub-Saharan Africa. For decades, unmet family planning need in this region has been characterized primarily as a problem of physical or economic access. The evidence presented here demonstrates that the binding constraint is often social, and that a trusted authority can relax it.

A vignette experiment isolates a woman's private conviction from her community's peer norm, her scriptural frame, and her religious leader's stance. I find that latent demand for family planning is high. Conformity to an unspoken social sanction suppresses adoption. A religious leader operates on this constraint by coordinating what women already quietly prefer. A supportive leader displaces the peer norm, moves behavior most effectively where the prevailing norm is misaligned, and lifts use more with a blessing than a prohibition suppresses it. This lever is bounded. Because estimated conformity in Meru County, Kenya, sits just below the critical tipping threshold, a leader can firmly bend the norm while engaged, but a leader alone cannot move the community into a self-sustaining, high-use equilibrium.

Two faith-based programs in Kenya allow me to test whether this mechanism scales. A difference-in-differences design on four waves of Demographic and Health Survey data reveals that program effects are heterogeneous, concentrated where the faith-based delivery network is dense and trust in religious leaders is high. The program that paired leader engagement with a standing means to act on the shifted norm lowered fertility durably, while the one that carried the same engagement without that means shows no reduction I can credibly attribute to it. The aggregate data cannot separate the channels on their own, but their pattern is the mechanism's signature, and it appears

under the conditions the model marks out.

Together, these results refine the familiar policy call to "involve religious leaders" into a targeted, testable design. Authority engagement yields the highest returns where conformity is strong, institutional trust is high, and a faith-based network can deliver on the newly relaxed norm. It succeeds by permitting individuals to act on a suppressed demand. Because this framework is not specific to family planning, its logic extends to other norm-governed behaviors, such as female genital cutting or intimate-partner violence, where the binding constraint is similarly social. More broadly, this paper treats religion not as a historical residue, but as a contemporary instrument of development policy. It offers a formal framework to measure how far a trusted authority can move the norms that hold a community in place.

References

Acemoglu, D., & Jackson, M. O. (2015). History, Expectations, and Leadership in the Evolution of Social Norms. *The Review of Economic Studies*, *82*(2), 423–456.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic Difference-in-Differences. *American Economic Review*, *111*(12), 4088–4118.

Ashraf, N., Bau, N., Nunn, N., & Voena, A. (2020). Bride Price and Female Education. *Journal of Political Economy*, *128*(2), 591–641.

Bailey, M. J. (2013). Fifty Years of Family Planning: New Evidence on the Long-Run Effects of Increasing Access to Contraception. *Brookings Papers on Economic Activity*, *2013*(1), 341–409.

Bassi, V., & Rasul, I. (2017). Persuasion: A Case Study of Papal Influences on Fertility-Related Beliefs and Behavior. *American Economic Journal: Applied Economics*, *9*(4), 250–302.

Bau, N., Henning, D., Low, C., & Steinberg, B. (2024, July). *Family Planning, Now and Later: Infertility Fears and Contraceptive Take-Up* (w32735). National Bureau of Economic Research. Cambridge, MA.

Becker, S. O., Rubin, J., & Woessmann, L. (2021). Religion in economic history: a survey. In *The Handbook of Historical Economics* (pp. 585–639). Elsevier.

Blackstone, S. R., Nwaozuru, U., & Iwelunmor, J. (2017). Factors Influencing Contraceptive Use in Sub-Saharan Africa: A Systematic Review. *International Quarterly of Community Health Education*, *37*(2), 79–91.

Bongaarts, J., & Bruce, J. (1995). The Causes of Unmet Need for Contraception and the Social Content of Services. *Studies in Family Planning*, *26*(2), 57.

Bongaarts, J., & Casterline, J. (2013). Fertility Transition: Is Sub-Saharan Africa Different? *Population and Development Review*, *38*, 153–168.

Bormet, M., Kishoyian, J., Siame, Y., Ngalande, N., Erb, K., Parker, K., Huber, D., & Hardee, K. (2021). Faith-Based Advocacy for Family Planning Works: Evidence From Kenya and Zambia. *Global Health: Science and Practice*, *9*(2), 254–263.

Brock, W. A., & Durlauf, S. N. (2001). Discrete Choice with Social Interactions. *The Review of Economic Studies*, *68*(2), 235–260.

Bryan, G., Choi, J. J., & Karlan, D. (2020). Randomizing Religion: the Impact of Protestant Evangelism on Economic Outcomes. *The Quarterly Journal of Economics*, *136*(1), 293–380.

Bursztyn, L., González, A. L., & Yanagizawa-Drott, D. (2020). Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia. *American Economic Review*, *110*(10), 2997–3029.

Caldwell, J. C., & Caldwell, P. (1987). The Cultural Context of High Fertility in sub-Saharan Africa. *Population and Development Review*, *13*(3), 409.

Dupas, P., Jayachandran, S., Lleras-Muney, A., & Rossi, P. (2025). The Negligible Effect of Free Contraception on Fertility: Experimental Evidence from Burkina Faso. *American Economic Review*, *115*(8), 2659–2688.

Field, E., Pande, R., Rigol, N., Schaner, S., & Troyer Moore, C. (2021). On Her Own Account: How Strengthening Women's Financial Control Impacts Labor Supply and Gender Norms. *American Economic Review*, *111*(7), 2342–2375.

Goldin, C., & Katz, L. F. (2002). The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions. *Journal of Political Economy*, *110*(4), 730–770.

Gulesci, S., Jindani, S., La Ferrara, E., Smerdon, D., Sulaiman, M., & Young, P. (2025). A Stepping Stone Approach to Norm Transitions. *American Economic Review*, *115*(7), 2237–2266.

Jones, B. F., & Olken, B. A. (2005). Do Leaders Matter? National Leadership and Growth Since World War II. *The Quarterly Journal of Economics*, *120*(3), 835–864.

KNBS. (2019). *2019 Kenya Population and Housing Census* (IV). Kenya National Bureau of Statistics. Nairobi. Retrieved September 15, 2025, from https://www.knbs.or.ke/wp-content/uploads/2023/09/2019-Kenya-population-and-Housing-Census-Volume-4-Distribution-of-Population-by-Socio-Economic-Characteristics.pdf

KNBS & ICF. (2023). *Kenya Demographic and Health Survey 2022*. KNBS and ICF. Nairobi, Kenya, Rockville, Maryland, USA.

Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, *60*(3), 531.

Pew Research Center. (2025). *How the Global Religious Landscape Changed From 2010 to 2020* [Pew Research Center] [Section: Religious Demographics]. Retrieved March 21, 2026, from https://www.pewresearch.org/religion/2025/06/09/how-the-global-religious-landscape-changed-from-2010-to-2020/

Prummer, A., & Siedlarek, J.-P. (2017). Community leaders and the preservation of cultural traits. *Journal of Economic Theory*, *168*, 143–176.

Ruark, A., Kishoyian, J., Bormet, M., & Huber, D. (2019). Increasing Family Planning Access in Kenya Through Engagement of Faith-Based Health Facilities, Religious Leaders, and Community Health Volunteers. *Global Health: Science and Practice*, *7*(3), 478–490.

Tessema, G. A., Streak Gomersall, J., Mahmood, M. A., & Laurence, C. O. (2016). Factors Determining Quality of Care in Family Planning Services in Africa: A Systematic Review of Mixed Evidence (K. Mortimer, Ed.). *PLOS ONE*, *11*(11), e0165627.

Timæus, I. M., & Moultrie, T. A. (2020). Pathways to Low Fertility: 50 Years of Limitation, Curtailment, and Postponement of Childbearing. *Demography*, *57*(1), 267–296.

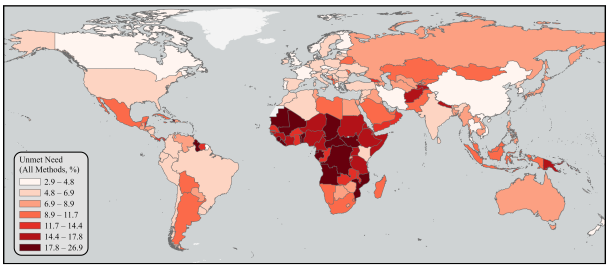
UNDESA. (2024). Model-Based Estimates and Projections of Family Planning Indicators. Retrieved September 11, 2025, from https://population.un.org/dataportal

United Nations. (2019). *Family Planning and the 2030 Agenda for Sustainable Development (Data Booklet)*.

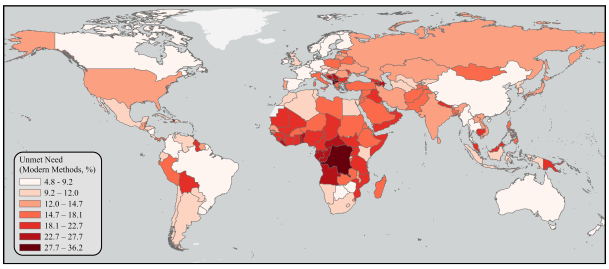
United Nations. (2022). *World Family Planning 2022: Meeting the Changing Needs for Family Planning*.

Appendix A: Motivational Statistics

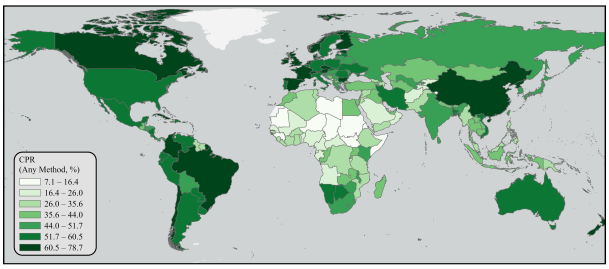
Figure A1: Global Reproductive Health and Fertility Indicators



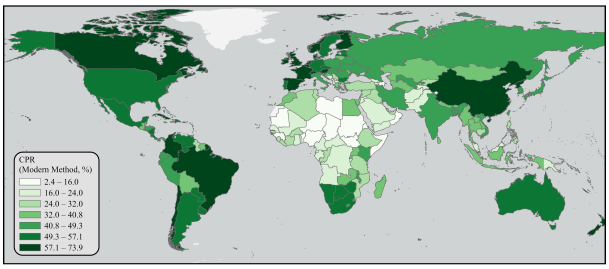
(a) Unmet Need for FP: Any Method



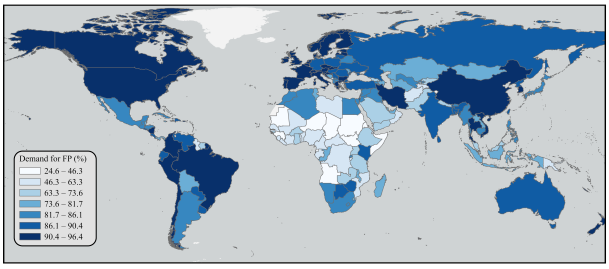
(b) Unmet Need for FP: Modern Method



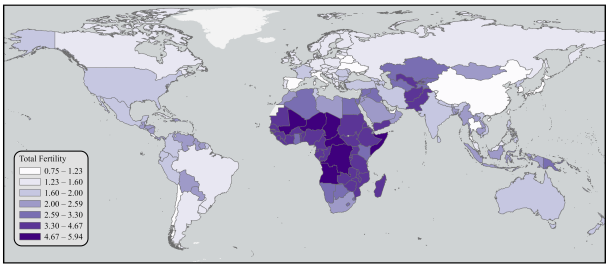
(c) CPR: Any Method



(d) CPR: Modern Method



(e) Total Demand for FP



(f) Total Fertility

NOTE: This figure presents global estimates for key reproductive health indicators among women of reproductive age (15–49). Panels (a) and (b) illustrate unmet need for family planning (FP), defined as the percentage of women who wish to stop or delay childbearing but are not using any contraceptive method, or specifically a modern method, respectively. Panels (c) and (d) display the Contraceptive Prevalence Rate (CPR), tracking the percentage of women currently using any method versus modern methods. Panel (e) measures total demand for FP, representing the fraction of women of reproductive age who are currently using any method or who have an unmet need. Panel (f) maps total fertility, representing the average number of children born per woman. Data for all panels are sourced from UNDESA (2024).

Table A1: Health Facility Landscape in Kenya

	KEPH Level 2 (Dispensary)	KEPH Level 3 (Health Centre)	KEPH Level 4 (Sub-county Hosp.)	KEPH Level 5+ (Referral)	Total
	Panel A: All Kenya Health Facilities by Ownership				
Government	5,125 (69.80%)	1,372 (77.50%)	397 (78.80%)	40 (82.50%)	6,934 (71.90%)
FBO	638 (39.70%)	249 (49.00%)	129 (48.80%)	10 (80.00%)	1,026 (43.50%)
Private	5,186 (49.70%)	2,124 (67.00%)	594 (61.60%)	25 (88.00%)	7,929 (55.40%)
Other	288 (35.10%)	60 (51.70%)	12 (58.30%)	0 (0.00%)	360 (38.60%)
Total	11,237 (57.90%)	3,805 (69.40%)	1,132 (66.20%)	75 (84.00%)	16,249 (61.30%)
	Panel B: CHAK Member Facilities by Denomination				
ACK	32 (37.5%)	4 (100.0%)	5 (60.0%)	0 (0.0%)	41 (46.3%)
Methodist	9 (77.8%)	7 (28.6%)	0 (0.0%)	0 (0.0%)	16 (56.3%)
SDA	25 (80.0%)	8 (87.5%)	3 (100.0%)	0 (0.0%)	36 (83.3%)
AIC / AIPCA	37 (59.5%)	8 (100.0%)	4 (75.0%)	1 (0.0%)	50 (66.0%)
Presbyterian/PCEA	25 (56.0%)	5 (20.0%)	2 (50.0%)	1 (100.0%)	33 (51.5%)
Orthodox	3 (0.0%)	2 (50.0%)	0 (0.0%)	1 (0.0%)	6 (16.7%)
Other FBO	82 (68.3%)	35 (71.4%)	17 (64.7%)	2 (100.0%)	136 (69.1%)
Total	213 (61.5%)	69 (69.6%)	31 (67.7%)	5 (60.0%)	318 (63.8%)

NOTE: KEPH = Kenya Essential Package for Health. Level 2 = Dispensary; Level 3 = Health Centre; Level 4 = Sub-county Hospital; Level 5+ = County and National Referral Hospitals. Percentages in parentheses show the share of facilities in each cell offering family planning services. FBO = Faith-Based Organization. CHAK = Christian Health Association of Kenya. Data from the Kenya Master Health Facility Registry. Panel B includes only CHAK member facilities; CHAK facilities also appear in Panel A under the FBO category.

Appendix B: Survey Sample Statistics

Table B1: Respondent Summary Statistics

	Mean	SD	N
Panel A. Demographics			
Age (years)	35.01	8.39	635
Living children	3.51	1.64	494
Age at first birth	19.29	3.50	494
Married or in union	0.81	0.40	635
Worked for pay (last 7 days)	0.56	0.50	635
No formal schooling	0.07	0.25	635
Primary schooling	0.69	0.46	635
Secondary schooling or higher	0.25	0.43	635
Panel B. Religion and trust			
Catholic	0.08	0.27	635
Mainline Protestant	0.20	0.40	635
Pentecostal / Charismatic	0.21	0.40	635
African Instituted Church	0.15	0.36	635
Seventh-day Adventist	0.03	0.17	635
Other Christian	0.33	0.47	635
Attends services weekly or more	0.95	0.21	635
Trusts religious leaders (somewhat or more)	0.90	0.31	634
Scriptural orientation (stewardship +, pronatal –)	2.03	1.49	635
Perceived FP social sanctions (1–5)	2.19	0.79	635
Panel C. Family planning			
Knows a family planning method	1.00	0.04	635
Has ever used family planning	0.95	0.21	635
Currently uses a modern method	0.84	0.36	603
Approves of modern family planning	0.96	0.19	635
Ideal number of children	4.30	1.55	635
Currently pregnant	0.04	0.19	635

NOTE: Sample of 635 reproductive-age Christian women surveyed in Meru County, Kenya. Means are proportions for indicator (0/1) variables. Living children and age at first birth are defined over the 494 mothers in the sample.

Table B2: Orthogonality of the Randomized Vignette Factors

	Community norm (E)	Leader stance (R)	Scripture frame (S)	Story slot	Age	Number of living children	Scriptural orientation
Community norm (E)	1.000						
Leader stance (R)	0.004	1.000					
Scripture frame (S)	0.008	-0.012	1.000				
Story slot	-0.012	-0.007	0.021	1.000			
Age	-0.007	0.000	-0.024	0.000	1.000		
Number of living children	-0.007	-0.003	-0.028	0.000	0.629***	1.000	
Scriptural orientation	0.004	-0.004	0.008	0.000	-0.012	-0.028	1.000

NOTE: Pooled pairwise correlations across all 5,080 vignette responses. The three randomized factors (community norm *E*, leader stance *R*, scripture frame *S*) are mutually uncorrelated and uncorrelated with story position and respondent characteristics, confirming the within-respondent randomization is clean.

Table B3: Balance of Respondent Characteristics Across Randomized Kendi Profiles

	Profile						
	A	B	C	D	E	All	Joint Test (p-value)
Age (years)	36.103	33.959	34.203	35.179	35.746	35.008	0.180
Living children	3.632	3.600	3.239	3.667	3.455	3.510	0.310
Married or in union	0.802	0.862	0.797	0.772	0.800	0.806	0.486
Secondary schooling or higher	0.172	0.260	0.301	0.252	0.231	0.246	0.199
Attends services weekly or more	0.957	0.959	0.958	0.959	0.931	0.953	0.780
Trusts religious leaders	0.931	0.869	0.902	0.911	0.869	0.896	0.432
Scriptural orientation	1.793	2.154	1.986	2.057	2.162	2.033	0.291
Perceived FP social sanctions	2.196	2.106	2.260	2.138	2.223	2.187	0.513

NOTE: Means of respondent characteristics by the between-respondent Kendi profile (*A*–*E*). The final column reports the p-value of a joint F-test that the profile indicators are zero in a regression of each characteristic on profile. No characteristic differs across profiles.

Appendix C: Proofs

This appendix proves the three results stated in Section 3. Throughout, $\bar{y}$ denotes an equilibrium of the mean-field map $\bar{y} = \tanh(h + \alpha\bar{y})$, with h collecting the private and directional forces defined there. Proposition 1 characterizes when this map has a single root versus a tipping regime with multiple equilibria, Proposition 2 derives the social multiplier that propagates a shift in the leader or scripture force, and Corollary 1 shows why a trusted authority moves behavior most when it opposes an entrenched norm.

$\bowtie$ *Proof* (Proposition 1)*: By* (5)*, an equilibrium is a fixed point of the map* $g(\bar{y}) = \tanh(h + \alpha\bar{y})$ *on* $[-1, 1]$*. As* g *is continuous and maps* $[-1, 1]$ *into itself, a fixed point exists. Its slope is* $g'(\bar{y}) = \alpha\left(1 - \tanh^2(h + \alpha\bar{y})\right) = \alpha\left(1 - g^2\right)$*, which lies in* $(0, \alpha]$ *and attains its maximum* α *where* $g = 0$*.*

When $\alpha \leq 1$*,* $g'(\bar{y}) \leq 1$ *everywhere, with equality only at isolated points, so* $g(\bar{y}) - \bar{y}$ *is strictly decreasing and has a single zero. The equilibrium is unique.*

When $\alpha > 1$*,* $g'(\bar{y}) > 1$ *on an interval around the point where* $g = 0$*, so* g *is S-shaped and crosses the* 45*-degree line up to three times. Let* $\phi(\bar{y}) = g(\bar{y}) - \bar{y}$*. The number of roots changes only at a tangency, where* $\phi(\bar{y}) = 0$ *and* $\phi'(\bar{y}) = 0$ *hold jointly, that is where* $\alpha(1 - \bar{y}^2) = 1$*. These two conditions define the critical value* $h^* > 0$*. For* $|h| < h^*$ *the curve crosses three times, giving two stable outer equilibria and one unstable interior equilibrium. For* $|h| = h^*$ *two of these roots coincide in a double root. For* $|h| > h^*$ *a single crossing remains. This is the mean-field characterization of Brock and Durlauf (2001), to which I refer for the complete argument.* □

$\bowtie$ *Proof* (Proposition 2)*: Differentiate the equilibrium condition* $\bar{y} = \tanh(h + \alpha\bar{y})$ *with respect to* $\bar{R}$*, using* $\partial h/\partial\bar{R} = \gamma$ *and* $\frac{d}{dz}\tanh(z) = 1 - \tanh^2(z) = 1 - \bar{y}^2$ *evaluated at the equilibrium. This gives*

$$\frac{\partial\bar{y}}{\partial\bar{R}} = (1 - \bar{y}^2)\left(\gamma + \alpha\frac{\partial\bar{y}}{\partial\bar{R}}\right).$$

Collecting terms yields $\frac{\partial\bar{y}}{\partial\bar{R}}\left[1 - \alpha(1 - \bar{y}^2)\right] = \gamma(1 - \bar{y}^2)$*, so*

$$\frac{\partial\bar{y}}{\partial\bar{R}} = \gamma\left(1 - \bar{y}^2\right)M(\bar{y};\alpha), \qquad M(\bar{y};\alpha) = \frac{1}{1 - \alpha(1 - \bar{y}^2)}.$$

Repeating the argument with $\partial h/\partial\bar{S} = \delta$ *gives* $\partial\bar{y}/\partial\bar{S} = \delta(1 - \bar{y}^2)M(\bar{y};\alpha)$*. At any stable equilibrium* $\alpha(1 - \bar{y}^2) < 1$*, so the denominator is positive and* M *is well defined. Since* $1 - \alpha(1 - \bar{y}^2) \in (0, 1]$ *at any stable equilibrium, it follows that* $M \geq 1$*, and* $M \to \infty$ *as* $\alpha(1 - \bar{y}^2) \to 1$*.* □

$\bowtie$ *Proof* (Corollary 1)*: Differentiate the extended equilibrium* $\bar{y} = \tanh(h + \alpha(\bar{R}, \bar{S})\bar{y})$ *with respect to* $\bar{R}$*, using* $\partial h/\partial\bar{R} = \gamma$ *and, for* $\bar{R} > 0$*,* $\partial\alpha/\partial\bar{R} = -\alpha_R$*. With* $\frac{d}{dz}\tanh(z) = 1 - \bar{y}^2$ *at the equilibrium,*

$$\frac{\partial\bar{y}}{\partial\bar{R}} = (1 - \bar{y}^2)\left(\gamma - \alpha_R\bar{y} + \alpha(\bar{R}, \bar{S})\frac{\partial\bar{y}}{\partial\bar{R}}\right).$$

Collecting terms and writing $M = 1/(1 - \alpha(\bar{R}, \bar{S})(1 - \bar{y}^2))$ yields the stated expression. The scripture derivative follows identically, with $\partial h/\partial \bar{S} = \delta$ and $\partial \alpha/\partial \bar{S} = -\alpha_S$, and the opposing cases ($\bar{R} < 0$, $\bar{S} < 0$) flip the sign of the salience derivative. The economic content lies in the salience term $-\alpha_R \bar{y}$. When a leader opposes an entrenched norm, $\bar{R}$ and $\bar{y}$ carry opposite signs, so this term adds to the direct push γ and the two reinforce one another, whereas a leader who merely echoes the prevailing norm sees the two offset. A trusted authority is therefore most effective preaching against the choir, which is the comparative static the main text builds on. □

Appendix D: The Vignette Stimuli

This appendix reproduces the full set of cards used in the vignette survey of Section 4. Each respondent was shown four cards per story, one of the five Kendi profiles, which stayed fixed across her eight stories, and one level of each of the three randomized treatment factors. Figure D1 shows the five profiles, and Figures D2, D3, and D4 show the three levels of the community norm, the religious leader's message, and the scriptural frame. Figure D5 reproduces the figurine scale used to elicit the perceived-behavior items.

Elicitation questions. The survey was administered in Kimeru. The items below give the verbatim English source wording of the questions asked after each vignette.

- *Primary outcome.* "On a scale of 0 to 10, where 0 is no chance and 10 is definitely. What is the chance that Kendi will go to the clinic to receive a family planning method?"
- *Normative agreement.* "Based on this story, do you personally agree or disagree that Kendi should use a family planning method?"
- *Secret access* (asked once per respondent). "Kendi can secretly receive a family planning injection (like Depo-Provera) at a clinic far from her village. The injection lasts three months, leaves no visible sign, and no one in her village, not even her husband, could find out. Do you think Kendi would get the injection?"
- *Comprehension check* (on one randomly selected factor). "In the story I just read, how many women in Kendi's village use family planning?"

Figure D1: Kendi Lifecycle Profiles

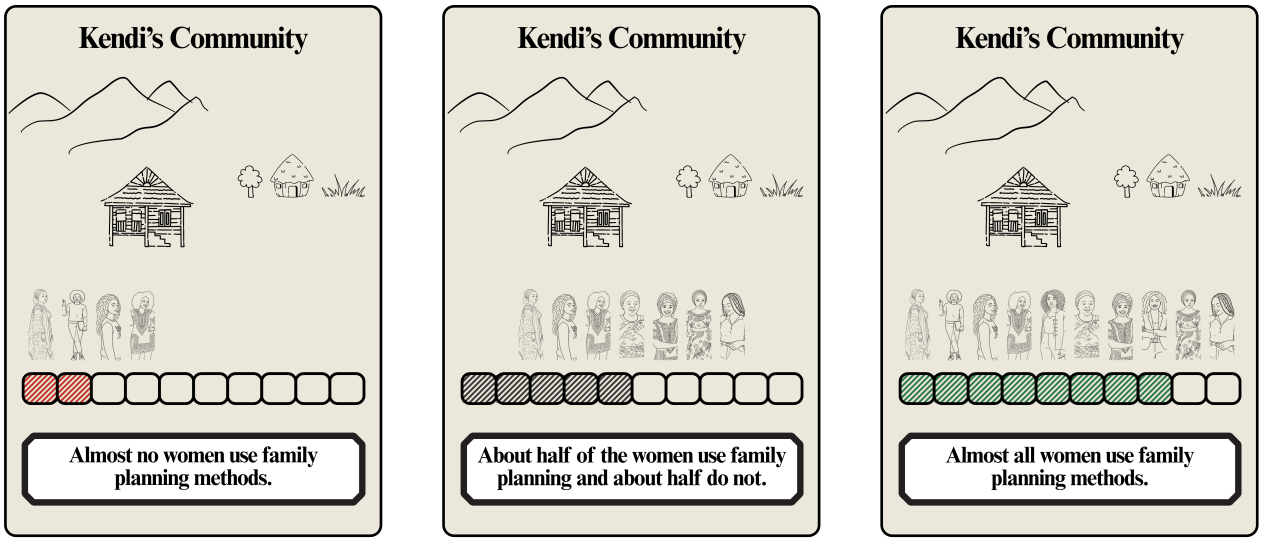


(a) Profile A (b) Profile B (c) Profile C

(d) Profile D (e) Profile E

NOTE: The five Kendi profiles. They form a ladder in which each adjacent pair differs in a single attribute, so that the four adjacent contrasts recover the age, marriage, and parity effects that identify the baseline preference parameter β.

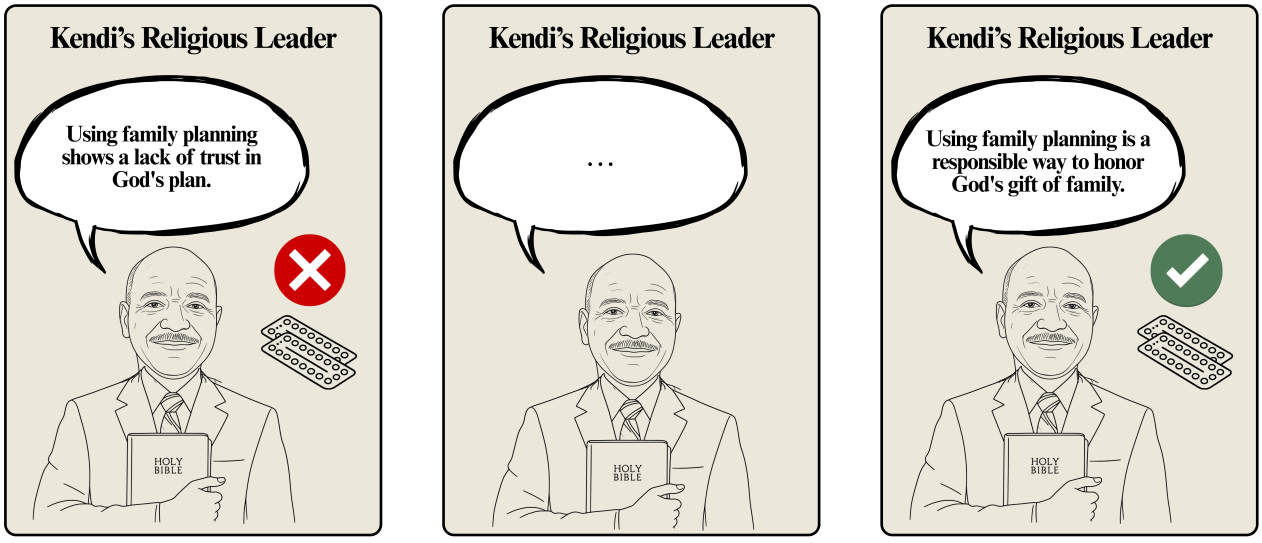
Figure D2: Community-Norm Card at Three Levels



(a) $E = -0.6$ (two of ten) (b) $E = 0$ (five of ten) (c) $E = +0.6$ (eight of ten)

NOTE: The community-norm card marks two, five, or eight of ten women as using family planning with shaded blocks, coded $E = -0.6$, 0, and $+0.6$.

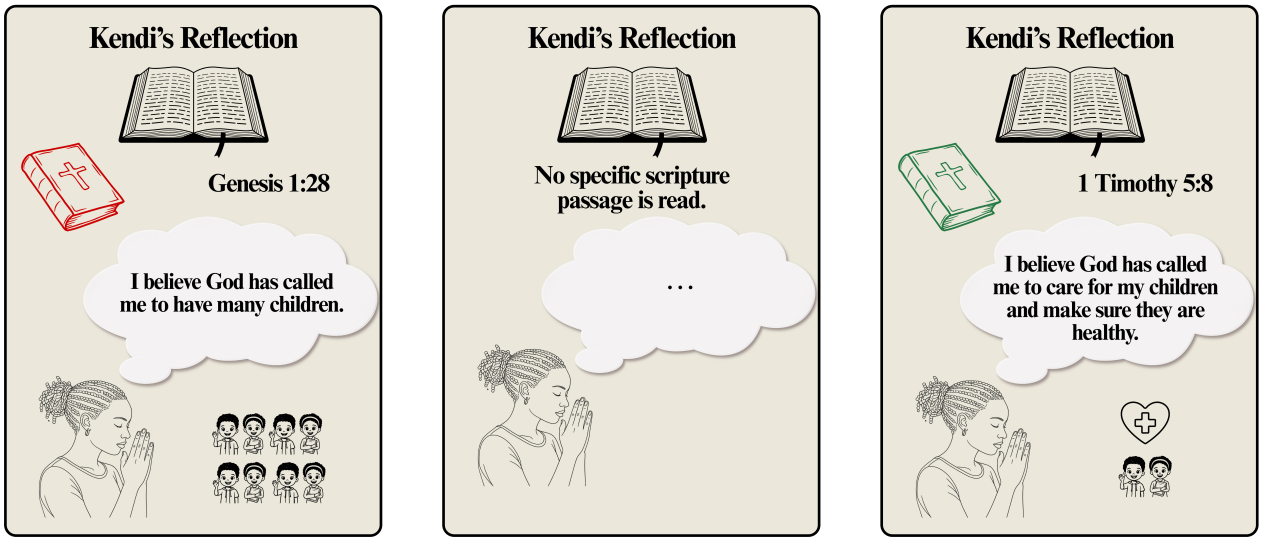
Figure D3: Religious-Leader Card at Three Levels



(a) $R = -1$ (against) (b) $R = 0$ (silent) (c) $R = +1$ (supports)

NOTE: The religious-leader card at its three levels. The leader teaches that family planning shows a lack of trust in God's plan ($R = -1$), has not spoken about it ($R = 0$), or teaches that it is a responsible way to honor God's gift of family ($R = +1$).

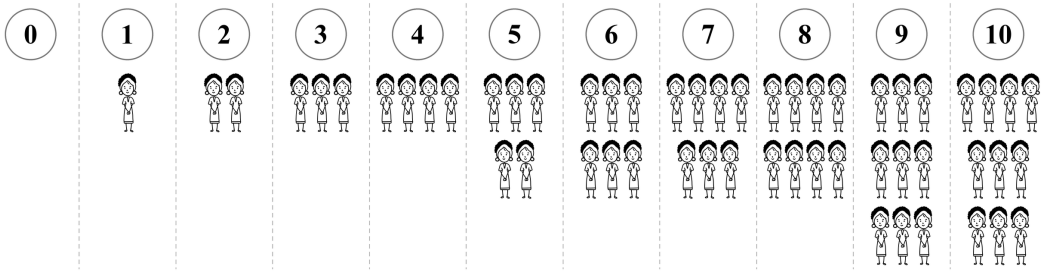
Figure D4: Scripture Card at Three Levels



(a) $S = -1$ (fruitfulness) (b) $S = 0$ (no passage) (c) $S = +1$ (stewardship)

NOTE: The scripture card at its three levels. Kendi reads Genesis 1:28 on being fruitful and multiplying ($S = -1$), no particular passage ($S = 0$), or 1 Timothy 5:8 on providing for one's household ($S = +1$).

Figure D5: The Figurine Frequency Scale



NOTE: The figurine scale was used to elicit perceived behavior, where a respondent estimated how many of ten women would act in a given way. The primary chance elicitation of Section 4.3 used the zero-to-ten chance scale of Figure 8 instead.