Resuf Ahmed
Postdoctoral Fellow
Stanford King Center on Global Development, Stanford University
Room 334, 366 Galvez Street, Stanford, CA 94305
resuf@stanford.edu | resuf.github.io

September 15, 2026

Professors Emily Breza, Davide Viviano, and Arun Chandrasekhar
Workshop on Development Economics and Networks for Early-Career Researchers
National Bureau of Economic Research
Cambridge, MA

Dear Professors Breza, Viviano, and Chandrasekhar,

I am writing to apply to the Workshop on Development Economics and Networks for Early-Career Researchers, to be held in Cambridge on December 5, 2026.

I am a Postdoctoral Fellow at the Stanford King Center on Global Development. I completed my PhD in Economics at HEC Lausanne, University of Lausanne, in 2025, advised by Dominic Rohner and Ekaterina Zhuravskaya. My work is in development and political economy, with a focus on how beliefs about group identity shape economic and social interaction in India, and I use both field experiments and large-scale administrative and spatial data.

The paper I would like to present is **"Debunking Myths: Ancestral Narratives and Intergroup Cooperation in India,"** joint with Riccardo Ghidoni, Alessandro Saia, Mariapia Mendola, and Stefano Piasenti. It is a lab-in-the-field experiment in Uttar Pradesh in which we return personalized genetic ancestry test results to Hindu respondents and randomize whether they receive the result before or after a set of incentivized and survey outcomes. The draft is a work in progress: endline data collection is ongoing, and we are currently designing a larger follow-up survey. The design of that follow-up is exactly where I would most value the workshop's input, as I describe in the accompanying statement.

Should I be shortlisted, **Katherine Casey** (Stanford Graduate School of Business, kecasey@stanford.edu) could be contacted to provide a recommendation.

Enclosed, in a single PDF, are this cover letter, my curriculum vitae including a list of all completed PhD coursework, a one-page statement of motivation and research interests, and the working paper. Thank you for considering my application.

Sincerely,

Resuf Ahmed

RESUF AHMED

Last updated on September 15, 2026

Postdoctoral Fellow, Stanford King Center on Global Development, Stanford University
Personal Webpage | resuf@stanford.edu |

CURRENT POSITION

Postdoctoral Fellow *January 2026 -*
Stanford University

EDUCATION

PhD Economics *2020 - 2025*
HEC Lausanne, University of Lausanne
Advisors: Dominic Rohner and Ekaterina Zhuravskaya

MA Economics *2018 - 2020*
Paris School of Economics

BA Economics *2014 - 2017*
University of Delhi

PhD COURSEWORK

All doctoral coursework completed as part of the PhD in Economics, HEC Lausanne (2020–2025).

University of Lausanne
Advanced Methods in Natural Experiments

University of Geneva
Microeconomics I (PhD)
Microeconomics II (PhD)

University of Zurich *(doctoral module mobility, Doctoral Program in Economics)*
Econometrics for Research Students, Part I *Fall 2020*
Macroeconomics for Research Students, Part I *Fall 2020*
Econometrics for Research Students, Part II *Spring 2021*
Workshop in Machine Learning Techniques in R *Spring 2021*

Geneva Graduate Institute (IHEID)
Advanced Development Microeconomics: Market Failures *Spring 2022*
Advanced Development Microeconomics: Other Topics *Spring 2022*

RESEARCH

Field: Development Economics, Political Economics, Cultural Economics

Working Papers

From Reservation to Transformation: The Socioeconomic Impact of Political Quotas

Ancient Epics in the Television Age: Religious Identity and the Rise of Hindu Nationalism in India (with Dean Yang, Alessandro Saia, Akhila Kovvuri and Paul Brimble) (*NBER WP 33417*) *Under Review*

Leader Marches as an Opposition Electoral Strategy (with Feyaad Allie) *Under Review*

Technology and the Production of Justice: System-Wide Effects of Indonesia's E-Court Reform (with Sultan Mehmood and Alexandra Filimonova) *Under Review*

Work in Progress

Power-sharing, Segregation and Conflict (with Ekaterina Zhuravskaya, Saumitra Jha, Dominic Rohner, and Alessandro Saia)

Debunking Myths: Ancestral Narratives and Intergroup Cooperation in India (with Riccardo Ghidoni, Mariapia Mendola, Stefano Piasenti and Alessandro Saia)

Press & Public Engagement

The Sunday morning broadcast that changed Indian democracy, March 2026, *VoxDev*
Decoding Bharat Jodo Yatra's impact on Karnataka Elections, May 2023, *Hindustan Times*
How the Ramayan TV Show Helped Shape a Singular Nationwide Hindu Identity, February 2025, *The Hindu (coverage)*

SEMINARS AND CONFERENCES

** (Presentation by co-author)*

20th Annual Conference on Economic Growth and Development, ISI-Delhi	December 2025
* BSE Summer Forum - Political Economy, BSE (Spain)	July 2025
Eleventh ECINEQ Conference, World Bank Group and George Washington University	July 2025
Tokyo Conference on Institutions, Law, and Development, GRIPS, Japan)	June 2025
European Meeting on the Political Economy of South Asia, Naples	June 2025
* BREAD Conference on Development, Princeton University	May 2025
19th Annual Conference on Economic Growth and Development, ISI-Delhi	December 2024
Lewis Lab Graduate Student Workshop, University of Manchester	November 2024
Paris School of Economics	November 2024
ASREC Europe 2024	September 2024
ACES Summer School, Hong Kong	August 2024
IRES Graduate Workshop, LA (USA)	June 2024
Paris School of Economics	May 2024
ASREC 2024, Chapman University	March 2024
Delhi School of Economics	December 2023
18th Annual Conference on Economic Growth and Development, ISI-Delhi	December 2023
ASREC 2023 Workshop, Harvard University	March 2023
Graduate Student Seminar, Brown University	March 2023
17th Annual Conference on Economic Growth and Development, ISI-Delhi	December 2022
DSE Winter School 2022, Delhi School of Economics	December 2022
Zurich Political Economy Seminar Series (Z-PESS), ETH Zurich	November 2022
ISI Seminar, Indian Statistical Institute - Delhi	August 2022
Society for Economics Research in India Workshop, ISB Hyderabad	July 2022
Applied Economics Lunch Seminar, Paris School of Economics	April 2022
ASREC 2022 workshop (Online), Chapman University	April 2022
Public Economics and Policy Seminar, University of Lausanne	March 2022
ECON GSW Seminar, University of Zurich	October 2021

TEACHING

Development Economics MSc Economics, University of Lausanne	Fall 2025, 2024, 2023, 2022
Economic Analysis and Communication MSc Economics, University of Lausanne	Fall 2024, 2023, 2022, 2021
Statistics and Econometrics I BSc Economics, University of Lausanne	Fall 2024, 2023, 2022, 2021

AWARDS & HONOURS

Virginia Gray Graduate Student Research Award 2024 by the American Political Science Association
Doctoral Mobility Fellowship by the University of Lausanne - 30,000 USD
PhD Scholarship by the University of Lausanne for the year 2020 - 2025
PPD Grant 2019 by the Paris School of Economics
Charpak BCS Scholarship (Bourse de Couverture Sociale) by the Government of France. 2018-2020
Young India Fellowship 2017 by the Ashoka University - 700,000 INR (Declined)

PROFESSIONAL SERVICE

Referee: *World Development*

SKILLS

Python, QGIS, STATA and R

Field Visits

Karnataka, India (January 2024)
Uttar Pradesh, India (March-May, 2025)

Statement of Motivation and Research Interests

Resuf Ahmed · Stanford King Center on Global Development

Research interests. I work on the political economy of group identity in developing countries, with a regional focus on India. Two questions organize my agenda. First, where do the beliefs that sustain group boundaries come from, and what happens to economic and social interaction when those beliefs are disturbed? Second, how does the spatial and social structure in which groups are embedded determine whether those boundaries harden into conflict or give way to exchange? I work with lab-in-the-field experiments and with large-scale administrative, electoral, and spatial data for India.

The paper I would present. In "Debunking Myths: Ancestral Narratives and Intergroup Cooperation in India," we return personalized genetic ancestry results to Hindu respondents in Uttar Pradesh and randomize whether they receive the result before or after a set of incentivized and survey outcomes. Priors sit at a ceiling: 86.5 percent believe they are 100 percent Indian by descent, so the test mostly delivers a downward shock to a purity claim. Respondents do not update as a Bayesian would. On learning their own ancestry is less pure than believed, they revise beliefs about a Muslim target's Indian ancestry *downward*, preserving the perceived distance between the groups. The outcomes that move are not abstract attitudes but stated willingness to enter social and market interactions with Muslims. The treatment operates on whom people will link with across a religious boundary, which makes the experiment, at bottom, about the belief foundations of homophily.

Why this workshop. The paper's central limitation is a network limitation. We measure individual willingness to form ties, but we do not observe the networks respondents are embedded in, and we randomize at the individual level. We therefore cannot address what matters most: what happens when ancestry information diffuses through a village rather than arriving at one person in isolation. Identity beliefs seem a canonical case in which individual treatment effects are a poor guide to aggregate ones. If a respondent's neighbors also learn their ancestry is less "pure" than assumed, both the meaning of the information and the scope for boundary maintenance change.

We are now designing a substantially larger follow-up survey, and the design choices are still open. I would like to build network structure into it: elicit ties at baseline, randomize the saturation of ancestry feedback within villages, and study how results and their interpretations propagate. That raises questions I cannot yet answer well. How should one power a saturation design when the spillovers travel through belief transmission rather than observed behavior? What empirically distinguishes diffusion of information from diffusion of norms? How should interference be handled when the outcome is tie formation itself? Having this design scrutinized by this group before we field it would materially change what we do.

I would also value reactions to a second project, "Power-sharing, Segregation and Conflict" (with Ekaterina Zhuravskaya, Saumitra Jha, Dominic Rohner, and Alessandro Saia), which measures Hindu–Muslim residential segregation across India and asks how it interacts with political power-sharing in shaping communal violence. Segregation is a coarse stand-in for contact structure, and I would like to learn what this literature would take to be the right object to measure. My doctoral training was in micro, macro, econometrics, and development; what I know of empirical network economics I have picked up in the course of trying to use it. A workshop pairing lectures from senior researchers with work-in-progress presentations is exactly the setting in which I would learn most, and at a point in this project where that learning would still change the research.

Debunking Myths: Ancestral Narratives and Intergroup Cooperation in India

Resuf Ahmed* Riccardo Ghidoni† Alessandro Saia† Mariapia Mendola‡ Stefano Piasenti†

Feedback draft — endline data collection ongoing. Comments very welcome.

May 2026

Abstract

We study how personalized genetic ancestry information affects Hindu–Muslim boundaries in Uttar Pradesh, India. In a lab-in-the-field experiment, we randomize the timing of the Genetic Ancestry Test (GAT) result: treated respondents ($N = 475$) receive their result before incentivized trust-game decisions and survey outcomes, while control respondents receive the same information afterward. Because 86.5 percent of respondents report a prior belief of 100 percent Indian ancestry, the treatment primarily delivers a negative shock to a near-ceiling purity prior. In the full-sample intent-to-treat, receiving the GAT result before outcomes reduces beliefs about a Muslim target's Indian ancestry (-0.17 SD), lowers willingness to engage in social interaction with Muslims (-0.13 SD), and decreases warmth toward Muslim-majority countries (-0.21 SD). Reciprocation in a trust game is muted in the pooled comparison but becomes strongly negative (-0.16 to -0.19 SD) when we exploit heterogeneity in the realized belief gap. All effects are concentrated among respondents whose ancestry realization falls below their prior belief. The pattern is highly selective. Stereotype indices and attitudes toward non-Muslim countries remain unchanged. Moreover, despite learning that their own ancestry is less pure than believed, respondents revise beliefs about Muslim ancestry *downward* rather than upward, contrary to what Bayesian updating would predict. Treated respondents are also significantly less likely to place Muslim ancestry at or above their own realized GAT score (-7.9 percentage points, $p < 0.01$), indicating an identity-preserving revision that maintains perceived group distinctiveness. These findings are consistent with identity-defensive belief updating, in which individuals respond to identity-threatening information by reinforcing group boundaries rather than revising them.

Keywords: genetic ancestry; Hindu nationalism; intergroup relations; motivated reasoning; field experiment; India

JEL Codes: C93, D72, J15, Z13

*Stanford University.
†University of Bologna.
‡University of Milano-Bicocca.

Note to Workshop Participants

Endline data collection is ongoing and will allow us to examine additional elements related to ancestral beliefs of participants. We are also designing a larger survey where participants will be exposed to information about aggregate averages of GAT scores of Hindu and/or Muslim respondents to investigate the mechanisms underlying these results more precisely. We welcome suggestions on the survey design and on which mechanism tests would be most informative.

1 Introduction

Group identities are central determinants of social, political, and economic life. A large literature in economics, sociology, and political science has studied how ethnic, religious, and national cleavages shape trust, cooperation, redistribution, public goods, conflict, innovation, and growth. Yet most existing research focuses on the consequences of group identity, rather than on the beliefs that sustain these identities and make them socially consequential. One such belief is the idea of common ancestry.

The idea of common ancestry is most often associated with ethnicity. Even though ethnicity has no single agreed-upon definition, a central element that binds individuals into an ethnic community is the belief that they descend from common ancestors. Smith famously describes the idea of common and homogeneous ancestral origin as "the sine qua non of ethnicity" (Smith, 1986). But ancestry-based narratives are not confined to ethnic groups narrowly defined. Religious and national identities can also be narrated as communities of descent, especially when political movements portray some groups as native, pure, or historically entitled to the territory, and others as foreign, mixed, or less authentic.

In this sense, ancestry is not merely a biological fact. It is also a social and political belief: a way through which groups imagine their origins, define their boundaries, and distinguish insiders from outsiders. While universal common ancestry is a pillar of modern evolutionary theory (Theobald, 2010), perceived ancestral origins are shaped by selective historical narratives, myths, and political reinterpretations. These ancestral narratives can sustain cooperation and unity within groups, but they can also create exclusion, hierarchy, and intergroup hostility. Historical examples illustrate this dual role. In eighteenth-century France, the belief that nobles and commoners descended from different ancestral roots helped legitimize social division and hierarchy. By contrast, myths of Roman descent were mobilized in Fascist Italy to portray Italians as members of a unified and powerful national community. In both cases, the force of the narrative did not depend on its historical accuracy. What mattered was that ancestry provided a language through which social and political boundaries could be drawn.

Today, ancestral narratives remain central to exclusionary politics. In many societies, claims about "pure" origins, native bloodlines, or indigenous descent are used to define who truly belongs. Invocations of Anglo-Saxon, Celtic, or Aryan ancestry have been used by nationalist and supremacist movements to assert superiority, defend nativist claims, and resist inclusion. These narratives suggest the existence of a potentially important gap between perceived and actual ancestral origins: individuals may believe that members of their own group share a unique and homogeneous ancestry, while overstating their difference from out-groups.

This paper studies such a gap in the context of religious identity. Our setting is not an ethnic cleavage in the narrow sense, but a Hindu–Muslim boundary that is often interpreted through ancestry. In rural Uttar Pradesh, India, Hindu nationalist ideology often presents Hindus as uniquely indigenous to India, while Muslims are frequently portrayed as descendants of foreign outsiders. The relevant belief is therefore not only that Hindus and Muslims differ in religion, but that they differ in origin. This narrative is inconsistent with genetic evidence showing that Hindus and Muslims from the same regions of South Asia descend from broadly shared ancestral populations.

We ask what happens when such beliefs are confronted with personalized scientific information. Can genetic ancestry information weaken exclusionary religious boundaries by revealing that ancestry is more mixed, and more shared, than people believe? Or do individuals respond defensively when such information threatens a valued religious-national identity?

We conduct a lab-in-the-field experiment in rural Uttar Pradesh, India, where Hindu nationalist ideology is politically salient and many Hindus believe that their ancestry is uniquely indigenous, in contrast to Muslims who are often perceived as descendants of outsiders. Scientific evidence shows that this belief is incorrect. Hindus and Muslims in the same region share essentially the same ancestral populations (Moorjani et al., 2013). We test whether providing Hindu males with personalized genetic ancestry information shifts beliefs and behavior in a direction consistent with this shared origin.

Our design randomizes the timing of when respondents receive the result of a genetic ancestry test. All respondents provide a DNA sample and state their prior belief about the share of their ancestry traced to India. Treated respondents receive their test result immediately before making incentivized decisions and reporting attitudes. Control respondents receive the same information only after these outcomes are measured. This design isolates the causal effect of learning one's ancestry result while holding constant eventual access to the information.

The key empirical fact is simple and striking. Most respondents begin with a near 100 percent belief in their own Indian ancestry. When the test reveals a lower share, respondents do not update toward greater perceived similarity with Muslims. Instead, they revise their beliefs in the opposite direction. They lower their estimate of Muslim ancestry, reduce their willingness to interact socially and economically with Muslims, and express colder attitudes toward Muslim-majority countries. In incentivized trust-game behavior, they return less to Muslim counterparts once the analysis isolates cases in which the information contradicts prior beliefs. These effects are concentrated among respondents whose realized ancestry falls below their prior belief.

This pattern reveals a form of identity-defensive belief updating. When individuals receive information that challenges a valued identity, they do not passively incorporate it. Instead, they adjust related beliefs and preferences in ways that preserve perceived group distinctiveness. In our setting, learning that one's own ancestry is less "pure" leads individuals to revise Muslim ancestry beliefs *downward* rather than upward, preserving perceived group distinctiveness through an identity-preserving revision. This response is not consistent with standard Bayesian updating, which predicts that evidence of shared South Asian ancestry should raise beliefs about Muslim ancestry. It is, however, a direct prediction of models in which beliefs about identity are valued and protected (Bénabou and Tirole, 2011).

A central challenge is to distinguish identity-defensive processing from alternative explanations. One possibility is that respondents misinterpret the genetic information, for example by inferring that a lower Indian ancestry share reflects greater out-group ancestry. Several features of the data argue against a purely mechanical misinterpretation. The response is highly selective. It is concentrated on beliefs and behaviors directly tied to Hindu–Muslim boundaries, while leaving stereotype indices and attitudes toward non-Muslim countries unchanged. Respondents do not report lower trust in the genetic test itself, suggesting that they do not simply reject the information. The pattern instead points to a targeted adjustment of beliefs that preserves identity-relevant distinctions.

Our findings contribute to three strands of literature. First, we provide causal evidence on how information interventions affect intergroup attitudes. Prior work shows that contact and narrative-based interventions can reduce prejudice (Allport, 1954; Pettigrew and Tropp, 2006; Broockman and Kalla, 2016a; Ghosh et al., 2026), but it remains unclear whether factual information alone is sufficient. We show that it is not. Absent a social or narrative frame, personalized information about shared ancestry can produce the opposite of its intended effect. Second, we contribute to the literature on identity in economics by documenting how individuals respond to identity-threatening information in a high-stakes social context (Akerlof and Kranton, 2000). The identity-preserving

downward revision of out-group ancestry beliefs, which moves in the opposite direction to standard Bayesian updating, provides a novel empirical signature of motivated belief maintenance. Third, we contribute to the emerging literature on the social consequences of genetic ancestry information by providing the first randomized evidence on its behavioral effects.

The setting of our study is intentionally demanding. The sample consists of rural Hindu men in a region where Hindu nationalist beliefs are widespread and politically reinforced. This context sharpens the test of identity-defensive behavior, as the prior belief in ancestral distinctiveness is both strong and socially meaningful. While this may limit external validity, it provides a clear environment in which the theoretical mechanism should operate most strongly. Ongoing follow-up work examines whether similar patterns arise in more diverse and less politically charged populations.

Taken together, the results suggest a limitation of information-based approaches to reducing intergroup conflict. When beliefs are tied to identity, providing accurate information is not enough. Individuals may respond by reshaping related beliefs in ways that preserve group boundaries, even when doing so requires moving further away from the factual implication of the information itself.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the Indian context and the study region. Section 4 explains the experimental design. Section 5 describes the data and key variables. Section 6 presents the empirical strategy. Section 7 reports the main results. Section 8 concludes.

2 Literature Review

2.1 Information Interventions and Intergroup Relations

A growing literature uses randomized experiments to test whether structured contact or information provision can reduce intergroup prejudice. The contact hypothesis (Allport, 1954) predicts that intergroup exposure under cooperative conditions reduces hostility, a prediction confirmed on average across observational studies (Pettigrew and Tropp, 2006) but more fragile in experiments. Lowe (2021) shows that cooperative contact between Hindu and Muslim cricket players reduces prejudice, while adversarial contact increases it. Rao (2019) finds that school integration in Delhi raises generosity toward poorer peers but also triggers social distancing among richer children. Scacco and Warren (2018) documents reductions in discriminatory behavior from a shared vocational training program in Nigeria. On the information side, Broockman and Kalla (2016b) show that "deep canvassing" conversations, built around empathy induction and personal narrative, produce large reductions in transphobia. Recent work also studies structured interventions that combine contact with coordinated group activities. Ghosh et al. (2026) examine integrated youth camps in West Bengal that include rituals, sports, and civics training, and find sustained reductions in ingroup bias and increases in intergroup friendships. These interventions operate through shared experience, coordination, and repeated interaction, rather than through factual information alone.

Our paper departs from this literature in a critical dimension. Prior work delivers information through social contact or tailored narrative framing; we deliver cold, personalized, biological fact with no accompanying reframe. The contrast isolates whether factual information *per se* is sufficient to shift attitudes, independent of the social and emotional channels that effective interventions typically operate through. Our finding that it is not, and that it can backfire, clarifies what the active ingredient in successful interventions is: not factual correction, but the social or narrative scaffolding around it.

2.2 Identity Economics and Motivated Belief Maintenance

Akerlof and Kranton (2000) formalise the idea that group membership is a direct argument in the utility function: individuals place positive value on beliefs and behaviors that affirm a valued social identity. A key implication is that identity-threatening information is costly to receive, and agents will resist, reinterpret, or avoid it (Bénabou and Tirole, 2011). Bénabou and Tirole (2011) develop this formally: beliefs about one's type are assets that rational agents protect through motivated cognition, particularly when the identity is socially reinforced. Experimental support comes from Möbius et al. (2022), who show asymmetric updating on good versus bad performance news, and from Taber and Lodge (2006), who document that political partisans apply scepticism selectively to identity-inconsistent arguments.

We contribute three things to this literature. First, we test the theory in an unusually high-stakes context (religiously charged political identity in rural India), where the predicted effects should be large if the mechanism is real. Second, we document a specific form of identity defense that the existing theory predicts but that has not been cleanly identified: *identity-preserving belief revision*, in which respondents who learn their own ancestry is more mixed than believed revise the Muslim target's ancestry *downward* rather than upward, preserving perceived distinctiveness by moving in the opposite direction to what Bayesian updating would predict. Third, the selectivity of effects, significant on ancestry beliefs, social interaction, and Muslim-country affect but null on stereotypes and non-Muslim evaluations, provides unusually clean evidence for boundary-hardening specifically rather than generalized negativity.

2.3 The Genetics of South Asian Populations and the Information Shock

The information shock in our experiment rests on established population genetics. Moorjani et al. (2013) demonstrate that all South Asian populations, regardless of religion, caste, or region, descend from a mixture of the same two ancestral groups that admixed between 2,000 and 4,000 years ago. The implied genetic distance between Hindus and Muslims from the same district is negligible relative to global genetic variation. A Genetic Ancestry Test (GAT) measuring "Indian" (South Asian) ancestry therefore returns essentially the same figure for Hindu and Muslim respondents in the same village.

This creates the identification opportunity we exploit. Hindu nationalist ideology (Jaffrelot, 1996) holds that Hindu identity is rooted in a unique biological heritage, a claim inconsistent with the Moorjani findings. Most of our respondents held near-100% priors about their own Indian ancestry at prior belief, a belief that the GAT result contradicts for the majority of treated respondents. The shock is therefore not generic bad news but specifically a threat to a biologically grounded purity claim, which is precisely the condition under which Bénabou and Tirole (2011) predict identity-defensive responses. The academic literature on the downstream behavioral effects of ancestry-test results is virtually nonexistent (Rajaram and Hajela, 2019); our paper provides the first randomized evidence.

3 Context

3.1 Hindu–Muslim Relations in Uttar Pradesh

Uttar Pradesh (UP) is the most populous Indian state, with approximately 200 million residents. It is also the heartland of contemporary Hindu nationalism: the BJP has dominated UP elections since 2014. Muslims constitute approximately 19 percent of UP's population. Hindu-Muslim relations in UP have been shaped by a long history of communal tension, periodic violence, and political mobilisation around religious identity. The study villages are located in the Moradabad district of UP, an area characterised by mixed Hindu-Muslim settlement patterns, agricultural livelihoods, and relatively high exposure to nationalist political messaging.

3.2 Genetic Ancestry Beliefs in India

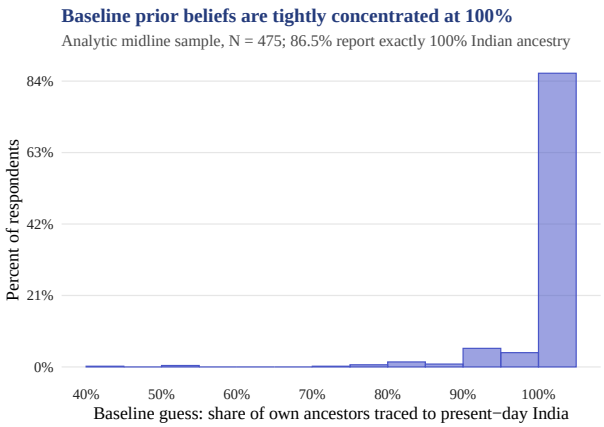
A core element of Hindu nationalist ideology is the claim that Hindus constitute a distinct civilizational and, implicitly, biological unit, often expressed through the concept of *Hindutva* ("Hindu-ness"). In the popular imagination, this is sometimes associated with a belief that Hindu ancestry is uniquely Indian, in contrast to Muslims who may be believed to descend from Central Asian or Middle Eastern invaders.

Genetic evidence sharply contradicts this narrative (Moorjani et al., 2013): the vast majority of DNA variation in both Hindus and Muslims in UP traces to the same ancestral South Asian populations.

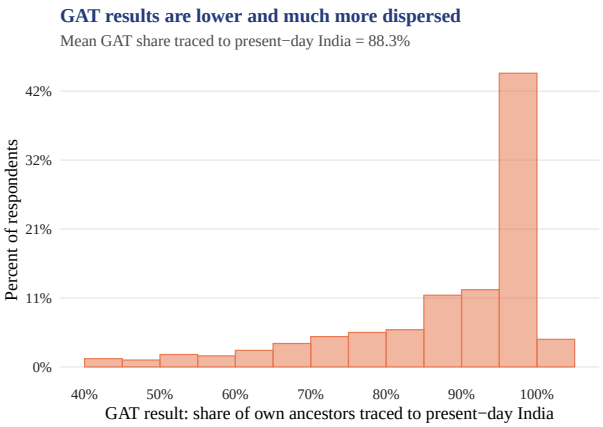
The GAT in our study measures the fraction of a respondent's genome attributable to South Asian ("Indian") ancestry, a quantity that is biologically identical in meaning for Hindus and Muslims from the same region. Respondents, however, interpret it through an identity lens: they understand the question as asking about the distinctively "Hindu" share of their heritage. At baseline, Hindu respondents guess their Indian ancestry share at approximately 98.5 percent on average; their actual GAT results average 88.5 percent, a gap of roughly 10 percentage points. The information shock for Hindu respondents is therefore that their ancestry is more mixed than believed, which they interpret as a challenge to Hindu genetic distinctiveness. Figure 1 shows the distributions of prior beliefs and actual GAT scores, the relationship between the two, and the realized gap among treated respondents.

3.3 Study Villages

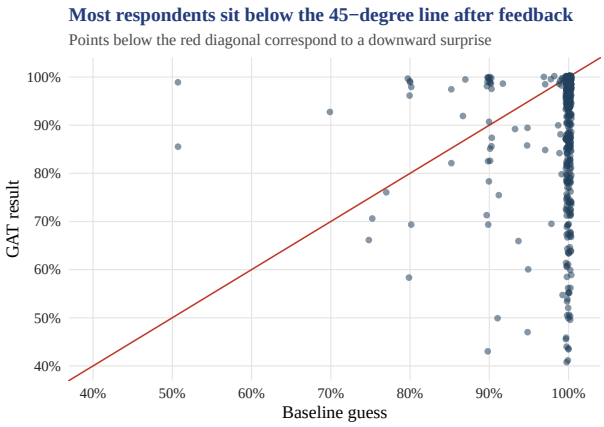
Our experiment was conducted in 51 villages in Moradabad District, Uttar Pradesh. The villages were selected to ensure variation in Muslim population share, enabling heterogeneity analysis by local out-group exposure. Figure 2 shows the location of the study villages. The villages are concentrated in a region with Muslim shares ranging from near zero to approximately 70 percent at the village level.



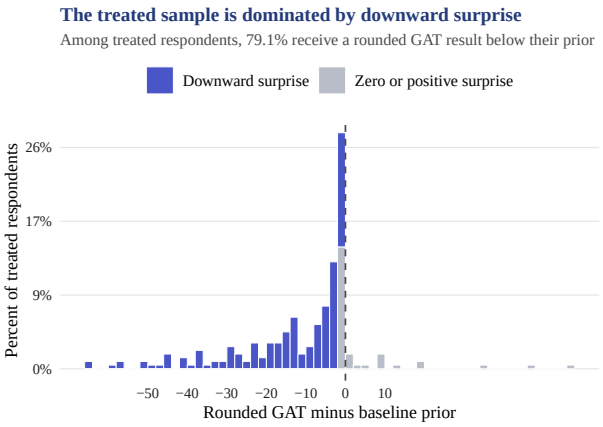
(a) Baseline prior beliefs



(b) GAT results



(c) Baseline priors versus GAT



(d) Realized gap among treated respondents

Figure 1: Beliefs, GAT results, and the realized identity threat

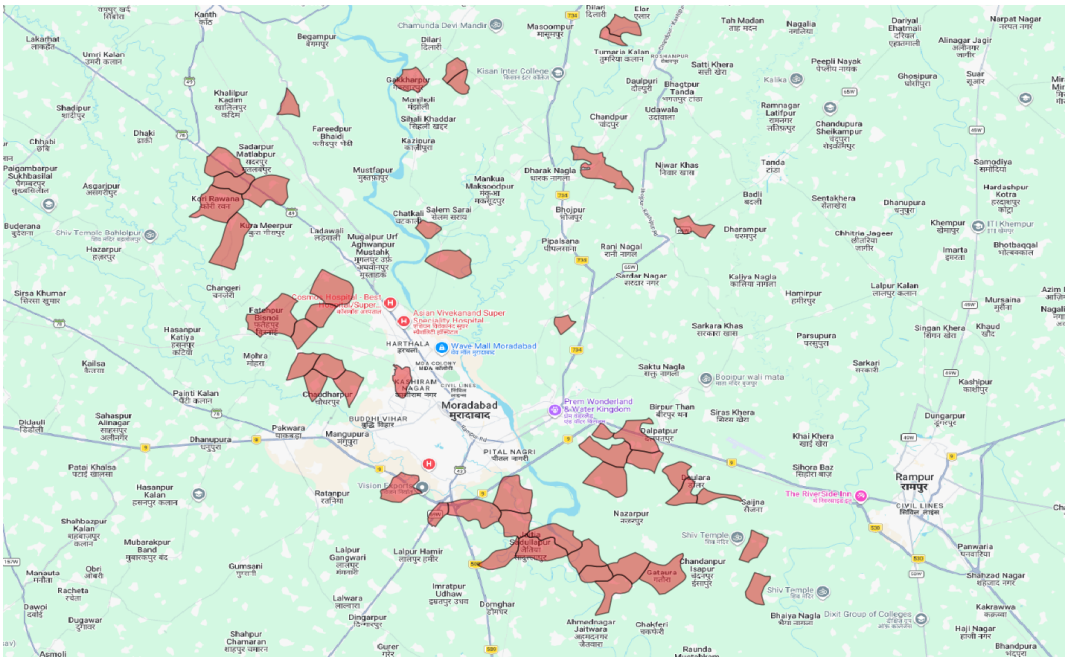


Figure 2: Study Villages and Muslim Voter Share

Note: Polygons show 51 study villages in Moradabad District, UP.

4 Experimental Design

4.1 Overview

The experiment studies whether personalized genetic ancestry information changes beliefs, social preferences, and intergroup behavior among Hindu men in rural Uttar Pradesh. The design has three connected components: a genetic ancestry test administered at baseline, a randomized timing intervention that determines when respondents receive their result, and incentivized trust-game decisions involving Hindu and Muslim counterparts.

Three samples contribute to the design. A **Hindu auxiliary sample** (100 men: 75 in Part 0A, 25 in Part 0B) made first-mover decisions in the trust game; their choices are used to match main-sample respondents with real prior decisions made by Hindu counterparts. A **Muslim auxiliary sample** (107 men with valid swabs) made the same first-mover decisions, creating real stakes for main-sample respondents matched to Muslim counterparts. The **Hindu main sample** (542 baseline records, 502 with valid swabs) stated their purity prior (their belief about the share of their ancestry traced to present-day India) and provided an oral swab for the genetic ancestry test.

The follow-up wave recontacts the 475 main-sample respondents who provided outcome data at midline. Treatment varies only the timing of the GAT result reveal: treated respondents learn their result before the outcomes are measured; control respondents learn it only after. All respondents therefore eventually receive the same information, and the design cleanly identifies the causal effect of learning the result before versus after outcome measurement. Figure 3 and Table **??** summarize the sample flow.

4.2 Randomisation

Randomization is conducted among the 475 Hindu main-sample respondents with nonmissing midline outcome data. Eligibility requires a valid nonfailed genetic ancestry test and nonmissing prior belief measures; no restriction on the direction or size of the belief gap is imposed for treatment assignment.[1]

Treatment assignment is stratified by three pre-specified baseline characteristics: caste status (high-caste versus other), whether the respondent's purity prior exceeds the median gap, and whether the village Muslim share is above the median. These three binary variables define eight randomization strata, and treatment is assigned within strata using a reproducible randomization algorithm that evaluates balance across village characteristics, political preferences, prior ancestry beliefs, age, education, occupation, and civil status.

A second stratified randomization independently assigns the order in which respondents face Hindu- and Muslim-named counterparts in the trust game, ensuring counterpart-order balance within treatment and control groups. Table 1 summarizes the design components.

4.3 Treatment: GAT Reveal

The treatment is the revelation of the respondent's genetic ancestry test result before outcome measurement. At baseline, all main-sample respondents provide an oral swab and state their prior belief about the percentage of their ancestry traced to present-day India. The test is described to respondents as nonmedical and as providing a general estimate of ancestry origins, not health information or individual genealogical identification.

In the follow-up wave, treated respondents are first reminded of their prior beleif and then shown their actual genetic ancestry result. The result is displayed on a 0 to 100 scale representing the percentage of the respondent's ancestors traced to present-day India. Enumerators are instructed to emphasize that the number shown is the respondent's true test result and to ensure comprehension of the scale.

Control respondents receive the same reminder about their prior belief but do not receive the genetic ancestry result before making the main decisions. Instead, they complete the same outcome modules first and are shown their own genetic ancestry result only at the end of the survey. Therefore, all respondents eventually receive the information, but only treated respondents receive it before the outcomes are measured. This design identifies the causal effect of ancestry-information revelation on subsequent behavior and beliefs, while holding fixed eventual access to the information.

The treatment is especially informative because most respondents held a near-100% purity prior at baseline. When the GAT result is below this prior, as it is for 79 percent of treated respondents, the reveal directly threatens the purity claim rather than providing confirmatory or neutral information. The intervention has three elements: respondents are reminded of their own prior belief; treated respondents are then shown the realized genetic ancestry result before completing the outcome modules; control respondents complete the same outcome modules first and receive the result only at the end.

[1] A subset of 426 respondents also have a positive purity gap and a nonmissing belief about Saleem's ancestry; this subset is used in mechanism robustness checks but is not the main analysis sample.

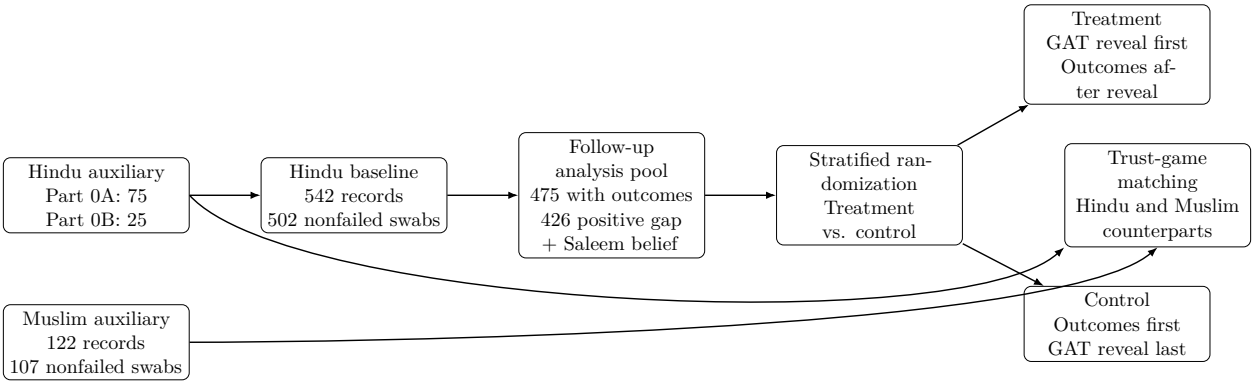
4.4 Outcomes

The main outcomes measure intergroup behavior, social distance, identity, political attitudes, and mechanisms related to beliefs about ancestry.

The primary behavioral outcome is the respondent's decision as **second mover** in an incentivized trust game. First-mover decisions were collected separately from the Hindu and Muslim auxiliary samples, who each decided whether to send INR 200 to a matched second mover. If the first mover sent the money, it tripled to INR 600 and the main-sample respondent (second mover) chose how much to return. Each respondent played two sequential second-mover roles: once matched to a Hindu first mover and once to a Muslim first mover, with the order randomized. The main trust-game outcomes are the amounts returned to each counterpart and their difference (Muslim minus Hindu return). Control-group respondents returned on average INR 237 to the Muslim counterpart out of INR 600.

The main survey outcomes are grouped as follows. **Muslim ancestry belief** is the respondent's estimate of the share of Muslims' ancestry traced to India (0–100 scale). This item uses the name "Saleem" as the Muslim target: respondents are asked how much of Saleem's ancestry they believe traces to India, and separately whether they believe Saleem's ancestry share is at or above their own GAT result. The control-group mean is 53 out of 100, substantially below the respondents' own prior of about 98. **Social interaction willingness** is an index of willingness to eat with, share a meal at the home of, and have as a neighbour a Muslim (5-point Likert scale; control mean 2.25). **Economic interaction willingness** is an index of willingness to buy from and sell to a Muslim trader (control mean 3.32). **Country feeling thermometers** measure warmth toward people from Nepal, England, Bangladesh, Pakistan, and Afghanistan (0–100); we separate Muslim-majority countries (Bangladesh, Pakistan, Afghanistan; control mean 10.4) from others. **Political and identity outcomes** include evaluations of political leaders, support for political systems (including religious guidance in politics), and stated Hindu-versus-Indian identity. **Mechanism checks** include recall of the GAT result, trust in the test, and perceived uniqueness of Hindu identity.

Figure 3: Experimental Timeline and Sample Flow



Notes: The main analysis sample is $N = 475$ (all with nonmissing midline outcomes). The 426-observation pool additionally requires a positive purity gap and nonmissing Saleem belief; it is used in mechanism robustness checks, not for treatment assignment.

Table 1: Randomization Strata and Balance Variables

Design component	Implementation
Eligibility	Hindu respondent; nonfailed GAT; nonmissing midline outcome data. Main sample: $N = 475$. Mechanism subsample (positive purity gap + nonmissing Saleem belief): $N = 426$.
Treatment strata	Caste status; above-median ancestral-belief gap; above-median village Muslim share.
Treatment assignment	Stratified random assignment. The assignment code searches candidate seeds and chooses the first assignment satisfying the pre-specified balance rule.
Balance variables	Village characteristics, political preferences, prior ancestry beliefs, interest in ancestry, age, education, occupation, civil status, and prior treatment status.
Order assignment	The order of Hindu- and Muslim-named counterpart decisions is randomized within treatment-by-stratum cells.

Notes: Randomization was conducted among all 475 respondents with midline outcome data. The positive-purity-gap subsample ($N = 426$) is used for mechanism robustness only, not for treatment assignment eligibility.

Table 2: Outcome Families

Outcome family	Variables	Scale	Incentivized
Reciprocation in the trust game	Amount returned to Hindu-named counterpart; amount returned to Muslim-named counterpart; difference in returns	INR 0–600	Yes
Social interaction	Eat with Muslim; meal at Muslim's house; Muslim neighbor	5-point Likert index	No
Economic interaction	Buy from a Muslim trader; sell to a Muslim trader	Index	No
Identity	Hindu versus Indian identity; choice of Indian-flag or Hindu-symbol magnet	Ordered scale; binary choice	Gift choice is real
Political attitudes	Feeling thermometers for political leaders; views over political systems	0–100; categorical	No
Outgroup attitudes	Feeling thermometers toward people from Nepal, Bangladesh, England, Pakistan, and Afghanistan	0–100	No
Muslim ancestry belief	Estimated share of Muslims' ancestry traced to India, elicited through the Saleem item	0–100	Yes
Mechanisms and checks	Recall of own GAT result; trust in test; uniqueness and identity challenge; responsibility for religion; risk willingness; open-ended perceived purpose	Mixed scales	No

Notes: Outcome families can be converted into pre-specified indices after standardizing component variables in the analysis sample.

5 Data and Key Variables

5.1 Sample

The Hindu main sample contains 542 baseline records. Of these, 40 provided swabs that failed quality checks, leaving 502 respondents with valid genetic ancestry results. At the midline follow-up, 27 respondents could not be recontacted, leaving the main analysis sample of $N = 475$ respondents with nonmissing outcome data. Table 4 confirms that the treatment and control groups are balanced on all observable baseline characteristics.

5.2 Descriptive Statistics

Table 3 presents descriptive statistics for the full midline sample ($N = 475$). The single most striking feature of the baseline data is the compression of purity priors: 86.5 percent of respondents guessed exactly 100 percent Indian ancestry, and the average prior belief is 98.4 percent. Actual GAT results are substantially lower and more dispersed, averaging 88.3 percent with a standard deviation of 14.5 percentage points. The median respondent therefore holds a purity prior roughly 10 percentage points above their actual result.

This gap between prior and realized ancestry is the source of the information shock. Among the 235 treated respondents, 186 (79 percent) receive a rounded GAT result below their prior, constituting the downward-surprise group in whom we expect identity-defensive responses. The remaining 49 treated respondents receive a result equal to (33) or above (16) their prior; these groups serve as within-treatment comparators for the mechanism analysis.

Table 3: Descriptive Statistics: Full Hindu Non-Failed Sample

Variable	N	Mean	SD	Median	Min	Max
Demographics						
Age (years)	502	42	15.24	41	18	80
Education (1–3 scale)	502	1.79	0.76	2	1	3
BJP supporter (share)	502	0.86	0.35	1	0	1
Ancestry beliefs						
Guessed Hindu ancestry (%)	502	98.47	5.69	100	40	100
Actual Indian ancestry (%)	502	88.5	14.35	94.9	34.1	100
Belief gap (Guessed – Actual)	502	9.96	15.34	4.3	-58	65.9
Village context						
Muslim voter share (%, E-rolls)	502	28.34	20.08	21.26	2.55	76.95
Muslim share (%, survey)	502	25.44	24.73	15	0	96

5.3 Randomisation Balance

Table 4 shows that the treatment and control groups are well balanced on observable baseline characteristics. No characteristic shows a statistically significant difference at conventional levels. The p-values from t-tests for balance on age, education, BJP support, ancestry beliefs, and Muslim share are all well above 0.10.

Table 4: Treatment Balance: Control vs. Treatment Group

Variable	**Control** *Mean (SD)*	**Treatment** *Mean (SD)*	***p*-value**
N	240	235	
Age (years)	41.51 (15.06)	42.23 (15.35)	0.603
Education (1–10 scale)	1.79 (0.77)	1.82 (0.77)	0.677
Occupation category	1.94 (0.91)	1.95 (0.92)	0.891
Upper caste	0.52 (0.50)	0.51 (0.50)	0.755
Civil status (married)	0.71 (0.45)	0.69 (0.47)	0.516
BJP supporter	0.85 (0.35)	0.85 (0.36)	0.924
Muslim voter share (%)	27.80 (20.02)	27.91 (20.31)	0.953
Total voters	1638 (1189)	1700 (1252)	0.579
Guessed Indian ancestry (%)	98.48 (4.65)	98.35 (6.82)	0.808
Actual Indian ancestry (%)	88.32 (14.39)	88.27 (14.73)	0.974
Total population	2163 (1352)	2168 (1337)	0.968
Male share (%)	52.68 (0.93)	52.77 (0.89)	0.290
SC share (%)	17.81 (14.02)	16.21 (12.45)	0.190
ST share (%)	0.01 (0.05)	0.01 (0.06)	0.549
Literacy rate (%)	50.97 (10.31)	50.33 (10.47)	0.504
Muslim share, survey (%)	24.82 (24.46)	25.89 (25.35)	0.643
Survey population	3407 (2854)	3157 (2418)	0.304

Note: Standard deviations in parentheses. p-values from two-sample t-tests. Sample: Hindu non-failed respondents appearing at midline.

6 Empirical Strategy

6.1 What We Estimate

The main causal object is the full-sample intent-to-treat effect of revealing the GAT result before outcome measurement. Treatment status is randomized among eligible midline respondents, so the pooled ITT compares respondents who learn their ancestry result before the trust-game and survey modules to respondents who complete the same modules before seeing the result.

We then use the realized direction of the ancestry shock as mechanism evidence. Define

$$\text{Gap}_i = \text{Prior}_i - \text{round}(\text{GAT}_i), \tag{1}$$

where Prior_i is the respondent's prior belief about the share of his ancestry traced to present-day India. A positive value means the prior exceeds the rounded GAT result; we refer to this as a downward surprise. A zero or negative value means the GAT result is equal to or above the prior.

Because prior belief are tightly concentrated at 100 percent, the distribution of the gap is dominated by positive values: the large majority of treated respondents receive a downward surprise, so the pooled ITT is primarily identified by identity-threatening feedback rather than a mixture of threatening and confirming realizations.

This decomposition is not a second randomized treatment. Treatment is randomized, but realized gap direction depends on the respondent's prior and laboratory result. We therefore interpret the gap-direction tables as evidence on where the randomized treatment is most likely to bite, not as a randomized comparison between downward- and non-downward-surprise respondents.

6.2 Main Specification

For the pooled ITT, we estimate

$$Y_{iv} = \alpha + \beta \, \text{Treat}_i + \mathbf{X}_i'\gamma + \delta_s + \varepsilon_{iv}, \tag{2}$$

where Y_{iv} is the outcome for respondent i in village v, Treat_i indicates early revelation of the GAT result, δ_s denotes randomization-strata fixed effects, and $\mathbf{X}_i$ includes controls for BJP support, education, occupation, and civil status. Standard errors are clustered at the village level.

For the mechanism decomposition, we replace Treat_i with indicators for treated respondents whose rounded GAT result is below, equal to, or above their prior belief. The omitted group is the control group.

6.3 Outcome Standardisation

All continuous outcomes are standardised relative to the control group at midline:

$$\tilde{Y}_{iv} = \frac{Y_{iv} - \bar{Y}^{\text{ctrl}}}{\sigma^{\text{ctrl}}}. \tag{3}$$

Thus, coefficients are measured in control-group standard deviations. This makes effect sizes comparable across trust-game allocations, ancestry beliefs, interaction indices, thermometers, and political outcomes.

6.4 Attrition and Missingness

The lead midline outcomes have little missingness. The main analytic file contains 475 respondents with nonmissing midline treatment status and outcome data for the core survey families. Some country-thermometer outcomes have fewer observations because respondents could report that they did not know the country. The tables report outcome-specific observation counts.

6.5 Heterogeneity

The central heterogeneity analysis is by realized gap direction. The identity-threat interpretation predicts stronger effects when the GAT result is below the respondent's prior belief, because this is the case in which the feedback directly contradicts a purity prior. We also report the same gap-direction decomposition among respondents whose prior belief was exactly 100 percent Indian ancestry, since this group represents the sharpest version of the purity-claim interpretation.

7 Main Results

7.1 Overview

Tables 5–7 report the lead midline results. Outcome families are defined in Section 4 (Table 2); social and economic interaction willingness are reported as separate indices since the treatment may affect everyday social contact and market contact differently.

The first result is the pooled ITT. The clearest full-sample effects are not in the trust game. They appear in ancestry beliefs, social interaction willingness, and symbolic evaluations of Muslim-majority

countries. The second result is the mechanism pattern: effects are stronger when the revealed GAT result is below the respondent's prior belief, exactly where the feedback is most identity-threatening. The third result is that the pattern becomes sharper among respondents whose prior belief was exactly 100 percent Indian ancestry.

Table 5: Pooled ITT Effects on Lead Intergroup Outcomes

	Trust			Beliefs		Intergroup attitudes			Thermometer		
	Muslim > Hindu	Muslim Amount	Hindu Amount	Guess Saleem	Gap vs. Self	Social interaction	Economic interaction	Stereotypes Index	All Countries	Non-Muslim	Muslim Countries
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Treatment	-0.069	-0.080	-0.023	-0.165**	-0.162**	-0.133*	-0.125	-0.034	-0.071	-0.072	-0.211**
	(0.076)	(0.081)	(0.119)	(0.075)	(0.076)	(0.069)	(0.083)	(0.070)	(0.088)	(0.092)	(0.094)
Observations	475	475	475	475	475	475	475	475	434	420	398

Notes: Outcomes are standardized using the control-group mean and standard deviation (full midline sample, $N = 475$). Control-group means equal zero by construction and are omitted. All specifications include randomization-strata fixed effects and controls for BJP support, education, occupation, and civil status. Standard errors are clustered at the village level. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.2 Pooled ITT Effects

Table 5 reports the pooled ITT for the full midline sample. Receiving the GAT result before outcome measurement lowers respondents' estimate of Saleem's Indian ancestry by 0.165 standard deviations and lowers the Saleem-versus-self benchmark by 0.162 standard deviations. These belief effects are statistically significant at the 5 percent level.

The interaction outcomes move in the same direction but differ in precision. The social interaction index falls by 0.133 standard deviations and is marginally significant. The economic interaction index falls by 0.125 standard deviations but is less precisely estimated. Reporting the two separately matters: the cleaner full-sample signal comes from the social-interaction items, and pooling them into a single index would obscure this difference.

The treatment also reduces warmth toward Muslim-majority countries by 0.211 standard deviations. By contrast, the all-country and non-Muslim-country thermometers are small and statistically insignificant. This pattern suggests a targeted symbolic response rather than a general decline in all country evaluations.

Trust-game outcomes are negative but imprecisely estimated in the pooled ITT. The Muslim-versus-Hindu trust gap, the amount allocated to the Muslim counterpart, and the amount allocated to the Hindu counterpart do not move significantly in the full sample. This attenuation is expected if the pooled ITT combines genuinely threatening feedback with feedback that is equal to or above the prior.

Table 6: Treatment Effects by Realized Gap Direction

	Trust			Beliefs		Intergroup attitudes			Thermometer		
	Muslim > Hindu	Muslim Amount	Hindu Amount	Guess Saleem	Gap vs. Self	Social interaction	Economic interaction	Stereotypes Index	All Countries	Non-Muslim	Muslim Countries
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Treatment: GAT below belief	-0.162**	-0.188**	-0.067	-0.199**	-0.230***	-0.144**	-0.167*	-0.024	-0.054	-0.090	-0.208**
	(0.070)	(0.089)	(0.137)	(0.086)	(0.086)	(0.071)	(0.098)	(0.075)	(0.099)	(0.096)	(0.103)
Treatment: GAT equals belief	0.249	0.265	0.148	-0.104	-0.151	-0.037	-0.006	-0.164	-0.064	0.044	-0.165
	(0.276)	(0.201)	(0.214)	(0.174)	(0.172)	(0.135)	(0.199)	(0.162)	(0.208)	(0.217)	(0.154)
Treatment: GAT above belief	0.355	0.469	0.128	0.109	0.616***	-0.191	0.129	0.116	-0.296	-0.096	-0.332**
	(0.349)	(0.299)	(0.373)	(0.180)	(0.201)	(0.182)	(0.188)	(0.161)	(0.253)	(0.329)	(0.156)
Observations	475	475	475	475	475	475	475	475	434	420	398

Notes: Outcomes are standardized using the control-group mean and standard deviation (full midline sample, $N = 475$). Control-group means equal zero by construction and are omitted. All specifications include randomization-strata fixed effects and controls for BJP support, education, occupation, and civil status. Standard errors are clustered at the village level. The three treatment rows decompose treated respondents by whether the rounded GAT result is below, equal to, or above the baseline prior. This decomposition is mechanism evidence, not a separately randomized comparison across rows. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.3 Realized Gap Direction

Table 6 decomposes treatment by whether the rounded GAT result is below, equal to, or above the respondent's prior.[2] This is the main mechanism table. It should be read as heterogeneity in the randomized treatment effect, not as a randomized comparison across gap categories.

The downward-surprise row is the most informative for the identity-threat interpretation. When the GAT result is below the respondent's prior, treatment reduces the Muslim-versus-Hindu trust measure by 0.162 standard deviations and the amount allocated to the Muslim counterpart by 0.188 standard deviations. It also reduces Saleem ancestry beliefs by 0.199 standard deviations, the Saleem-versus-self benchmark by 0.230 standard deviations, social interaction willingness by 0.144 standard deviations, and economic interaction willingness by 0.167 standard deviations. Warmth toward Muslim-majority countries falls by 0.208 standard deviations.

The equality and above-prior rows are noisier and much smaller samples (33 and 16 treated respondents, respectively) and should be interpreted as transparency checks rather than informative comparisons. They do not show the same coherent boundary-hardening pattern, consistent with the identity-threat interpretation: without a downward purity surprise, the information carries no threat and produces no defensive response.

[2]Among the 235 treated respondents in the midline sample, 186 have a GAT result below their prior, 33 have a GAT result equal to their prior, and 16 have a GAT result above their prior. Among the 240 control respondents, 183 have a GAT result below their prior, 43 have a GAT result equal to their prior, and 14 have a GAT result above their prior. Appendix Table 10 shows that treated and control respondents in the positive-gap arm are observationally similar, supporting the mechanism interpretation.

Table 7: Treatment Effects by Realized Gap Direction Among Respondents with 100 Percent Prior Beliefs

	Trust			Beliefs		Intergroup attitudes			Thermometer		
	Muslim > Hindu (1)	Muslim Amount (2)	Hindu Amount (3)	Guess Saleem (4)	Gap vs. Self (5)	Social interaction (6)	Economic interaction (7)	Stereotypes Index (8)	All Countries (9)	Non-Muslim (10)	Muslim Countries (11)
Treatment: GAT below belief	-0.163**	-0.186**	-0.024	-0.216**	-0.215**	-0.161**	-0.204**	-0.018	-0.047	-0.117	-0.197*
	(0.075)	(0.094)	(0.136)	(0.098)	(0.098)	(0.076)	(0.091)	(0.081)	(0.111)	(0.111)	(0.105)
Treatment: GAT equals belief	0.258	0.282	0.165	-0.166	-0.165	-0.055	-0.105	-0.203	-0.091	0.033	-0.183
	(0.272)	(0.206)	(0.217)	(0.160)	(0.158)	(0.146)	(0.207)	(0.174)	(0.209)	(0.216)	(0.155)
Control mean	-0.023	-0.007	-0.012	0.016	-0.035	0.000	0.050	0.007	-0.015	0.012	-0.012
Observations	411	411	411	411	411	411	411	411	377	365	347

Notes: Outcomes are standardized using the control-group mean and standard deviation. All specifications include randomization-strata fixed effects and controls for BJP support, education, occupation, and civil status. Standard errors are clustered at the village level. The sample is restricted to respondents whose baseline prior was exactly 100 percent Indian ancestry. In this sample, the nonnegative realized-gap arm is the equality arm. Stars: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.4 Respondents with 100 Percent Prior Beliefs

Table 7 restricts the sample to respondents who guessed exactly 100 percent Indian ancestry at baseline. This group comprises 411 respondents in the midline analysis sample (205 treatment, 206 control), and it is theoretically important because the GAT result most directly challenges a strong purity prior. Appendix Table 9 confirms that the treatment and control groups are well balanced within this subgroup.

Within this group, the downward-surprise treatment effects are again negative and concentrated on the same families. The Muslim-versus-Hindu trust measure falls by 0.163 standard deviations, the amount allocated to the Muslim counterpart falls by 0.186 standard deviations, Saleem ancestry beliefs fall by 0.216 standard deviations, and the Saleem-versus-self benchmark falls by 0.215 standard deviations. Both interaction indices decline: social interaction willingness falls by 0.161 standard deviations, while economic interaction willingness falls by 0.204 standard deviations. The economic-interaction effect is now statistically significant at the 5 percent level, which suggests that market contact responses are strongest among respondents with the sharpest purity prior.

7.5 Informative Nulls and Political Spillovers

The null effects are substantively useful. The stereotype index is small and insignificant across the pooled ITT and the gap-direction tables. Non-Muslim country thermometers are also small and insignificant. These nulls argue against a simple story in which receiving an unexpected ancestry result makes respondents generically negative or more hostile on every survey margin.

Instead, the effects are concentrated on boundaries that are close to the ancestry information itself: beliefs about Muslim ancestry, willingness to interact socially or economically with Muslims, warmth toward Muslim-majority countries, and, in the downward-surprise group, incentivized behavior toward the Muslim counterpart.

One political outcome does move significantly. Treatment increases support for religious guidance in politics by 0.29 standard deviations ($p = 0.002$; control mean 2.43 on a 1–5 scale). This effect is specific: strong-leader, expert-rule, army-rule, and anti-democracy items are all flat. The finding extends the boundary-hardening story beyond private attitudes into institutional preferences. A respondent whose purity claim has been undermined cannot repair the biological loss, but can

demand stronger institutional affirmation of religious boundaries. This political result sits outside the headline intergroup tables and should be read as an exploratory finding pending replication in the online survey; it is included here because it disciplines the interpretation and is available in the current data.

7.6 Interpretation and Mechanisms

The results are most consistent with identity-defensive processing. The diagnostic evidence is the direction of the Saleem ancestry revision. Standard Bayesian updating predicts that respondents who learn their own ancestry is shared with South Asian populations should revise beliefs about Muslim ancestry *upward*: if my ancestry traces to the same populations as Muslims in the same region, their ancestry should also be high. The estimates move in the opposite direction. The treatment reduces beliefs about Saleem's Indian ancestry by 4.9 percentage points in levels (SE 2.2), against a mean own-belief gap of 10.1 percentage points. The Saleem revision is smaller in magnitude than the own-belief shock, but its direction is the key finding: Saleem beliefs go down when Bayesian updating predicts they should go up.

The relative-positioning measure reinforces this. In the control group, 17.5 percent of respondents believe Saleem's ancestry is at or above their own GAT result. Treatment reduces this share by 7.9 percentage points (column 5 of Tables 6–7, $p < 0.01$). Treated respondents are actively placing Saleem *below* their own realized ancestry score, directly asserting Hindu distinctiveness despite having just learned their own ancestry is more mixed than believed. This identity-preserving revision is a direct prediction of motivated belief maintenance (Bénabou and Tirole, 2011) and is difficult to explain by standard Bayesian updating or generic bad-news responses.

Three alternative mechanisms can be largely set aside. **Credibility:** treated respondents show no increase in distrust of their GAT result (treatment effect on self-reported trust in the test: 0.06 SD, $p = 0.54$), making a simple information-rejection story implausible. **Demand effects:** if the response were driven by respondents performing Hindu identity for the enumerator, we would expect broad hostility across all Muslim-related items. The stereotype index and political-leader favorability are both flat, while interaction willingness and ancestry beliefs move sharply; this selectivity is inconsistent with a uniform response-style shift. **Noisy updating:** respondents may misinterpret a lower-than-expected ancestry score as implying Muslim ancestry contamination rather than shared South Asian heritage. We cannot fully rule this out at present; the endline and online surveys will include a direct belief-elicitation item asking respondents what they believe their GAT result implies about Hindu–Muslim genetic similarity.

Pooled ITT trust-game effects are attenuated because they average the 186 downward-surprise respondents (for whom the treatment is identity-threatening) with the 49 equal-or-above-prior respondents (for whom it is neutral). Once the treatment is decomposed by realized gap direction, trust-game returns to the Muslim counterpart fall by 0.19 SD in the downward-surprise arm, consistent with the broader mechanism and not contradictory to the pooled survey effects.

8 Conclusion

This paper provides randomized evidence that personalized genetic ancestry information backfires in a context where ancestral beliefs are politically charged. The full-sample intent-to-treat shows that receiving a GAT result before outcome measurement reduces beliefs about a Muslim target's Indian ancestry, lowers social and economic interaction willingness, and reduces warmth toward Muslim-

majority countries. These effects are concentrated among the 79 percent of treated respondents whose GAT result fell below their near-100% purity prior. When the treatment is restricted to this downward-surprise group, trust-game behavior also turns negative and economic interaction effects sharpen. Respondents who received neutral or above-prior feedback show no coherent response pattern, confirming that the purity threat, not the information per se, is the active ingredient.

The findings make three contributions. First, they provide the first randomized evidence on how personalized genetic ancestry feedback affects intergroup attitudes, extending the identity-defensive-processing literature to a biologically grounded belief context. Second, the identity-preserving downward revision of out-group ancestry beliefs is a novel empirical signature of motivated belief maintenance that standard Bayesian updating cannot produce: treated respondents lower their estimate of Muslim Indian ancestry and are less likely to place Muslim ancestry at or above their own realized GAT score, even though learning that one's own ancestry is shared with South Asian populations implies Muslim ancestry should be revised upward. Third, the results clarify what makes prejudice-reduction information interventions succeed or fail: the biological fact of shared ancestry, absent a cooperative or narrative reframe, is not sufficient and can produce the opposite of the intended effect.

Limitations and next steps. The midline results capture a 12-week effect window; we do not yet know whether the identity-defensive response persists or attenuates over time. Because both groups receive the GAT result by endline, the forthcoming endline comparison will identify early versus delayed exposure rather than a clean persistence test. The sample is also all-male, rural, and UP-specific, limiting external validity.

Two planned extensions address these gaps. Endline data collection is currently ongoing and will examine the durability of midline effects. A larger online survey in a nationally representative sample will replicate the core treatment with greater statistical power, include women, allow heterogeneity by region and caste, and implement anonymous self-administered outcome modules to rule out enumerator demand effects. Both surveys will elicit respondents' own interpretation of their GAT result, a direct test of whether respondents correctly understand what the result implies about Hindu–Muslim genetic similarity. We welcome feedback on the empirical strategy, the mechanism interpretation, and the design of these extensions.

References

Akerlof, George A. and Rachel E. Kranton, "Economics and Identity," *Quarterly Journal of Economics*, 2000, *115* (3), 715–753.

Allport, Gordon W., *The Nature of Prejudice*, Reading, MA: Addison-Wesley, 1954.

Bénabou, Roland and Jean Tirole, "Identity, Morals, and Taboos: Beliefs as Assets," *Quarterly Journal of Economics*, 2011, *126* (2), 805–855.

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

_ **and** _ , "Durably Reducing Transphobia: A Field Experiment on Door-to-Door Canvassing," *Science*, 2016, *352* (6282), 220–224.

Ghosh, Arkadev, Prerna Kundu, Matt Lowe, and Gareth Nellis, "Creating cohesive communities: a youth camp experiment in India," *Review of Economic Studies*, 2026, *93* (1), 438–475.

Jaffrelot, Christophe, *The Hindu Nationalist Movement in India*, New York: Columbia University Press, 1996.

Lowe, Matt, "Types of Contact: A Field Experiment on Collaborative and Adversarial Caste Integration," *American Economic Review*, 2021, *111* (6), 1807–1844.

Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat, "Managing Self-Confidence: Theory and Experimental Evidence," *Management Science*, 2022, *68* (11), 7793–7817.

Moorjani, Priya et al., "Genetic Evidence for Recent Population Mixture in India," *American Journal of Human Genetics*, 2013, *93* (3), 422–438.

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

Rajaram, Shyamala and Kashyap Hajela, "Genetic Ancestry Tests: The Science, Accuracy, and Consumer Use," *Frontiers in Genetics*, 2019, *10*, 1–10.

Rao, Gautam, "Familiarity Does Not Breed Contempt: Generosity, Discrimination, and Diversity in Delhi Schools," *American Economic Review*, 2019, *109* (3), 774–809.

Scacco, Alexandra and Shana S. Warren, "Can Social Contact Reduce Prejudice and Discrimination? Evidence from a Field Experiment in Nigeria," *American Political Science Review*, 2018, *112* (3), 654–677.

Smith, Anthony D, "The ethnic origins of nations," *(No Title)*, 1986.

Taber, Charles S. and Milton Lodge, "Motivated Skepticism in the Evaluation of Political Beliefs," *American Journal of Political Science*, 2006, *50* (3), 755–769.

Theobald, Douglas L, "A formal test of the theory of universal common ancestry," *Nature*, 2010, *465* (7295), 219–222.

Appendix: Design Integrity

Table 8 verifies that swab failure at baseline is not systematically related to observable respondent characteristics, supporting the assumption that the 502 Hindu respondents with valid genetic ancestry results are not a selected subset of the 542 Hindu baseline participants.

Table 8: Swab Quality Balance: Valid vs. Failed Swabs (Hindu Sample)

Variable	**Valid swab** *Mean (SD)*	**Failed swab** *Mean (SD)*	***p*-value**
N	502	40	
Age (years)	42.00 (15.24)	41.38 (14.85)	0.799
Education (1–3 scale)	1.79 (0.76)	1.70 (0.72)	0.450
Occupation category	1.96 (0.91)	1.80 (0.94)	0.298
BJP supporter	0.86 (0.35)	0.88 (0.33)	0.740
Guessed Indian ancestry (%)	98.47 (5.69)	98.97 (3.79)	0.437
Muslim voter share (%, census)	25.44 (24.73)	30.93 (26.27)	0.209

Note: Standard deviations in parentheses. p-values from two-sample t-tests. Sample restricted to Hindu main-sample respondents ($N = 542$). No covariate differs significantly between valid and failed swabs, supporting the assumption that swab failure is not systematically related to baseline characteristics.

Table 9 reports balance within the 100 percent prior subgroup used in Table 7. Table 10 reports a descriptive comparison of treated and control respondents whose rounded GAT result fell below their prior belief; because gap direction is a post-randomization realization, this table is a descriptive check rather than a formal balance test.

Table 9: Balance Check: Respondents with 100% Prior Belief (Hindu Sample)

Variable	**Treatment** *Mean (SD)*	**Control** *Mean (SD)*	***p*-value**
N	205	206	
Age (years)	42.69 (15.54)	41.66 (14.82)	0.489
Education (1–3 scale)	1.79 (0.77)	1.78 (0.78)	0.910
Occupation category	1.98 (0.93)	1.94 (0.90)	0.668
Upper caste	0.54 (0.50)	0.53 (0.50)	0.958
BJP supporter	0.87 (0.34)	0.86 (0.35)	0.789
Muslim voter share (%, census)	28.09 (20.39)	27.80 (19.93)	0.884

Note: Standard deviations in parentheses. p-values from two-sample t-tests. Sample restricted to Hindu main-sample respondents at midline with a baseline prior of exactly 100% Indian ancestry ($N = 411$). No covariate differs significantly between treatment and control groups, confirming that randomization achieved balance within this subgroup.

Table 10: Descriptive Comparison: GAT-Below-Prior Respondents by Treatment Status

Variable	**Treatment** *Mean (SD)*	**Control** *Mean (SD)*	***p*-value**
N	186	183	
Age (years)	43.43 (15.20)	43.03 (15.13)	0.799
Education (1–3 scale)	1.84 (0.76)	1.84 (0.78)	0.974
Occupation category	1.90 (0.94)	1.87 (0.91)	0.721
Upper caste	0.59 (0.49)	0.61 (0.49)	0.612
BJP supporter	0.84 (0.36)	0.86 (0.35)	0.710
Muslim voter share (%, census)	27.86 (19.99)	26.77 (19.46)	0.596

Note: Standard deviations in parentheses. p-values from two-sample t-tests. This is a *descriptive* comparison, not a balance check: gap direction is a post-randomization realization (it depends on the respondent's GAT result), so conditioning on it breaks the experimental design. The table documents that treated and control respondents with a positive belief gap are observationally similar, which supports the interpretation of the mechanism decomposition in Table 6. Sample: Hindu main-sample respondents at midline whose rounded GAT result fell below their baseline prior ($N = 369$).

Appendix: Baseline Survey Transcript (Part 1)

This appendix reproduces the English-language transcript of the baseline survey administered to the Hindu main sample ($N = 542$). The survey was implemented on tablets in Hindi. Enumerator instructions are shown in *italics*. Programmer and implementation notes are omitted.

Screen 1 — Informed Consent

Enumerator: Read the informed consent form aloud, ensuring clarity and comprehension. Take your time to make sure participants fully understand the information.

Module 1: Project Overview

Dear Sir/Madam, you are invited to participate in a research study conducted as part of project "ROOTS." Before deciding whether to participate, it is important that you have all the necessary information to make an informed and responsible choice.

Module 2: What Does Participation Involve?

Participation in the ROOTS research project includes: (1) completing questionnaires that cover sociodemographic questions and questions related to ethnic origins; (2) undergoing an oral swab to determine ethnic ancestry; (3) agreeing to be recontacted for follow-up questions; (4) authorising the use of your ethnic ancestry data in an aggregated and anonymous form for future research purposes.

Module 3: Benefits, Discomforts, and Risks

Participation in the study is voluntary. By taking part, you may become aware of personal information that differs from your expectations. This may have an impact on your mood or social interactions. In the event of any emotional discomfort arising from participation, you will have access to a psychological support service provided by us.

Module 4: Withdrawal

You have the right to withdraw your consent to participate at any time, without prior notice or specific reason.

Module 5: Feedback

You have the right to request information about the findings and results of the research.

Module 6: Anonymity

The processing of collected data will be conducted in a way that removes any references that could link individual responses to a specific person. All participants will remain anonymous, and once data collection is complete, it will no longer be possible to link the collected data to individual identities. Research results will be published in summary form.

Consent button text: *"By clicking 'Continue,' you confirm that you have read and understood both the 'Informed Consent Form' and the 'Informational Form on Data Processing,' and that you give us your permission to proceed with the study."*

Screen 2 — Welcome and Participation Fee

Thank you for participating, and welcome to this study! You will receive *Rs.*300 for your participation. Additionally, you can earn extra money depending on your and other participants' decisions during the study. You will be paid your earnings at the end of the study.

Screen 3 — Sociodemographic Questions

To start, we would like to ask a few simple questions about you.

1. In which year were you born?
2. What is your gender? *[Male / Female / Other]*
3. What is your religion? *[Hindu / Muslim / Other]*
4. What is your occupational status? *[Government job / Salaried employee / Own business / Student / Farming / Farm labour / Daily wage earner / Housewife / No work / Other]*
5. What is your highest education level completed? *[No formal education / Primary / Middle / Secondary / Higher secondary / Diploma / Undergraduate / Postgraduate / Professional degree / Doctoral / Other]*
6. What is your current civil status? *[Single / Married or cohabiting / Widowed / Separated or divorced / Other]*
7. Where were you born? (state/union territory)
   A. Where was your mother born?
   B. Where was your father born?
   C. Where was your paternal grandfather born?
   D. Where was your paternal grandmother born?
   E. Where was your maternal grandfather born?
   F. Where was your maternal grandmother born?
8. To which caste do you belong? *[GEN / ST / SC / OBC]*
9. Which party did you vote for in the last elections? *[BJP / INC / SP / BSP / Other INDIA alliance / Other NDA alliance / Did not vote / . . .]*
10. Which issues should the Government of India be focused on most today? Please rank the following in order of importance (1 = most important): welfare programmes for the poor; harmony between communities; corruption; employment opportunities.

Screen 4 — Swab Test Explanation

The next part of the study involves a swab test. We will collect a small sample of saliva from inside your mouth. This sample helps us get an idea of where your family originally came from — that is, the regions your ancestors (the people in your family from many generations ago) may have belonged to.

Please note:

- This is not a medical test and has no effect on your health. It does not provide any medical information.
- The test compares your saliva sample with those from people in different parts of the world. It gives a general idea of where your ancestors can be traced back to, but it does not identify specific individuals or provide a full history of your family.

The test involves rubbing a cotton swab on the inside of your cheeks for 30 seconds on each side to collect a sample. Additionally, it is important that you confirm you have not drunk, eaten, or smoked in the last 30 minutes.

The outcome of the test will tell us what percentage of your ancestors will be traced back to India (that is, the entire territory of India as it exists today — including all its current states and union territories), based on an estimate drawn from your saliva sample.

Enumerator: When reading the part that defines what we mean by India, point to the map and explain that the areas shaded in grey represent all countries outside of India, and the area marked in red is the entire territory of India as it exists today. Emphasise that Uttar Pradesh, and hence both you and the participants, are part of this red territory.

Screen 5 — Prior Elicitation (Incentivised)

Before we start the test, we want you to make your best guess about what percentage of your ancestors will be traced back to India by the test. You can earn extra money if your guess is close to the result.

What percentage of your ancestors will be traced back to India?

Look at the slider bar below. The bar represents the range of possible values from 0 to 100, where:

- 0 means you think none of your ancestors will be traced back to India
- 100 means you think all of your ancestors will be traced back to India

Use the slider to choose a percentage. If your guess is within 5 points (higher or lower) of the actual result from the swab test, you will receive an additional $Rs.50$.

Enumerator: Once the instructions have been clearly explained and understood (make sure they understand the "percentage" concept), ask the participant to drag the slider themselves. Do not give examples on specific numbers; avoid any kind of anchoring.

[Slider: integer 0–100]

Screen 6 — Swab Collection

Thank you for answering the questions. We will now quickly collect the saliva sample for the test, as explained earlier. Once the sample is collected, the study will be complete. We will inform you of the result in a few days.

Enumerator: At this point you must administer the GAT (oral swab).

Screen 7 — Closing

Thank you for participating in this study. The test that was administered to you will be analysed. In the second part of the study, we will show you the result of the test. If you have any questions or concerns, feel free to ask us. Have a great day!

Appendix: Midline Survey Transcripts (Part 2)

This appendix reproduces the English-language transcripts of the midline (follow-up) surveys for the control and treatment groups. The surveys were implemented on tablets in Hindi. Control and treatment groups received identical outcome modules; the only structural difference is **when** the GAT result is revealed: after all outcomes for the control group (Screen 14) and before all outcomes for the treatment group (Screen 3). The Muslim target in the ancestry belief elicitation was named Saleem. Enumerator instructions are shown in *italics*. Programmer and implementation notes are omitted.

Part 2A — Control Group

Screen 1 — Welcome

Enumerator: Read aloud and make sure they understand.

Welcome to the second part of our study. You will receive *Rs.*200 for your participation. Additionally, you can earn extra money depending on your and other participants' decisions during the study.

In this second part, you will participate in a simple activity, followed by a series of questions.

At the end of the study, we will randomly select 200 participants to receive payment based on their choices in the activity. If you are selected, you will be paid at the conclusion of the study.

At the end of the study, we will also show you the result of the swab test you took in Part 1. That is, we will show you the actual percentage of your ancestors which has been traced back to India by the test.

Screen 2 — Prior Recall

Recall that in the first part of the study, you used a slider to guess what percentage of your ancestors would be traced back to India.

You guessed that [**X%**] of your ancestors would be traced back to India.

[Participant's own baseline guess displayed here]

Screens 3–4 — Trust Game (Two Rounds)

Enumerator: Read aloud and make sure they understand. Each participant plays the trust game twice, matched with [SELECTED NAME 1] in Screen 3 and [SELECTED NAME 2] in Screen 4. The two names are either Rizwan and Rajendra (or Rajendra and Rizwan), assigned in random order for each participant.

We will now ask you to take part in a simple activity. You can earn extra money from this activity.

In this activity, you are matched with another participant, [**SELECTED NAME**]. You will not know who [SELECTED NAME] is, and he will not know who you are.

[SELECTED NAME] played first. He was given *Rs.*200 and had to choose whether to KEEP the *Rs.*200 or SEND it to you.

- If [SELECTED NAME] kept the money, he earned *Rs.*200 and you received nothing. In that case, the activity ends here.
- If [SELECTED NAME] sent the *Rs.*200 to you, the amount was tripled. That means *Rs.*200 became *Rs.*600, and you receive the full *Rs.*600.

Now, suppose that [SELECTED NAME] sent you the money. You have *Rs.*600, and it is your turn to decide how much to send back to [SELECTED NAME], and how much to keep for yourself. You can choose any amount between *Rs.*0 and *Rs.*600 to send back.

Please note: your decision will only matter if [SELECTED NAME] actually chose to send the money. You will learn at the end of the study whether this happened.

Please make your decision now. There is no right or wrong choice. It is completely up to you.

[Integer input: 0–600]

Screen 4: Identical structure; participant matched with [SELECTED NAME 2].

Screen 5 — Hindu–Indian Identity

Now we would like to ask you some additional questions.

Suppose you must choose between being Indian and being Hindu. Which of these two groups do you feel more strongly connected to?

[Scale: 0 = Only Hindu ... 4 = Only Indian]

Screen 6 — Social and Economic Interaction Willingness

Please answer the following questions on a scale from 0 to 10, where 0 means "completely unwilling or completely uncomfortable" and 10 means "completely willing or completely comfortable":

1. Would you eat with a Muslim?
2. Would you go to dinner at a Muslim's house?
3. Would you feel comfortable having a Muslim as your neighbour?
4. Would you buy from a Muslim merchant?
5. Would you sell to (or do business with) a Muslim buyer?

Screen 7 — Magnet Choice (Real Gift)

As an extra gift, we would like to give you a magnet. Please choose your preferred magnet:

- Indian flag
- Hindu religious symbol

Enumerator: Distribute the physical magnet according to the participant's choice.

Screen 8 — Political Leader Feeling Thermometer

We would like to get your feelings toward some national leaders, past and present. I will read a name aloud, and I would like you to rate that person using something we call a 'feeling thermometer.'

Ratings between 50 and 100 mean that you feel favourable and warm toward the person. Ratings between 0 and 50 mean that you do not feel favourable toward the person. A rating of 50 means you don't feel particularly warm or cold. If you don't recognize a name, just tell me and we'll move on.

- Narendra Modi *[0–100]*
- Rahul Gandhi *[0–100]*
- Akhilesh Yadav *[0–100]*
- Yogi Adityanath *[0–100]*
- Mahatma Gandhi *[0–100]*

Screen 9 — Country Feeling Thermometer

We would like to get your feelings toward people from other countries using the same feeling thermometer. I will read the name of the country, and I would like you to rate your feeling towards people from there.

- People from Nepal *[0–100]*
- People from Bangladesh *[0–100]*
- People from England *[0–100]*
- People from Pakistan *[0–100]*
- People from Afghanistan *[0–100]*

Screen 10 — Political System Attitudes

We are going to describe several types of political systems and ask what you think about each as a way of governing this country. For each one, would you say it is a very good, fairly good, fairly bad, or very bad way of governing?

- Having a strong leader who does not have to bother with parliament and elections.
- Having experts (bureaucrats or babus), not government, make decisions according to what they think is best for the country.
- Having the army rule.

- Having a democratic political system.
- Having a system governed by religious law in which there are no political parties or elections.

[Very good / Fairly good / Fairly bad / Very bad]

Screen 11 — Muslim Stereotype Index

Now we are going to ask you how well each of these words describes most Muslims.

[Rating scale per trait: e.g., friendly, intelligent, good citizens, honest, peaceful, arrogant, dependable, fanatical, ungrateful, unreasonable]

Screen 12 — Muslim Ancestry Belief (Incentivised)

We now have one additional question for you, and you can earn extra money if your answer is accurate.

Think about the test you took in the first part of the survey. We administered the same test to another person, **Saleem**.

What percentage of Saleem's ancestors do you think will be traced back to India by the test?

Look at the slider bar below. It shows values from 0 to 100, where:

- 0 means you think none of his ancestors will be traced back to India
- 100 means you think all of his ancestors will be traced back to India

Use the slider to choose a percentage. If your guess is within 5 points (higher or lower) of the actual result from his swab test, you will receive an additional *Rs.*50.

[Slider: integer 0–100]

Screen 13 — Mechanism and Debriefing Questions

Finally, we would like to ask you a few questions:

1. On a scale from 0 to 10, how much do you trust the accuracy and reliability of the swab test results? *[0 = not at all, 10 = completely]*
2. How much do you agree with the following statements? *[0–10]*
   - The result of the test makes me feel that Hindus are not unique.
   - The result of the test challenges my sense of being Hindu.
3. To what extent do you think people are responsible for their religion because they could have chosen a different one? *[0 = not at all, 10 = completely]*
4. How willing are you to take risks in general? *[0 = not at all, 10 = very willing]*

Open Why do you think the research team decided to run this study? *[Audio recorded]*

Screen 14 — GAT Result Revealed [Control: occurs AFTER all outcomes]

Enumerator: Take some time to show them the bar and make sure they understand what it means, and that this is the true result from the test!

Now we are going to show you the result of the test.

[Slider display: 0% ancestors from India *100% ancestors from India]*

According to the test, **[X%]** of your ancestors have been traced back to India, as shown in the slider above.

[Participant's individual GAT result displayed here]

Screen 15 — Closing

The study is now over. If you have any questions or concerns, feel free to ask us. Have a great day!

Part 2B — Treatment Group

The treatment group received an identical set of outcome modules (Screens 4–14 below correspond to Screens 3–13 of the control). The sole structural difference is that the GAT result is revealed **before** any outcome is measured (Screen 3 below), rather than at the end.

Screen 1 — Welcome

Enumerator: Read aloud and make sure they understand.

Welcome to the second part of our study. You will receive *Rs.*200 for your participation. Additionally, you can earn extra money depending on your and other participants' decisions during the study.

In this second part, you will participate in a simple activity, followed by a series of questions.

At the end of the study, we will randomly select 200 participants to receive payment based on their choices in the activity. If you are selected, you will be paid at the conclusion of the study.

Before we start the activity, we will show you the result of the swab test you took in Part 1. That is, we will show you the actual percentage of your ancestors which has been traced back to India by the test.

Screen 2 — Prior Recall

Recall that in the first part of the study, you used a slider to guess what percentage of your ancestors would be traced back to India.

You guessed that **[X%]** of your ancestors would be traced back to India.

[Participant's own baseline guess displayed here]

Screen 3 — GAT Result Revealed [Treatment: occurs BEFORE all outcomes]

Enumerator: Take some time to show them the bar and make sure they understand what it means, and that this is the true result from the test!

Now we are going to show you the result of the test.

[Slider display: 0% ancestors from India *100% ancestors from India]*

According to the test, **[X%]** of your ancestors have been traced back to India, as shown in the slider above.

[Participant's individual GAT result displayed here]

Screens 4–14 — Outcome Modules (identical to Control Screens 3–13)

Following the GAT reveal, the treatment group completed the same outcome modules as the control group, in the same order:

- **Screen 4:** Trust game, Round 1 (matched with [SELECTED NAME 1]; see Control Screen 3)
- **Screen 5:** Trust game, Round 2 (matched with [SELECTED NAME 2]; see Control Screen 4)
- **Screen 6:** Hindu–Indian identity (see Control Screen 5)
- **Screen 7:** Social and economic interaction willingness (see Control Screen 6)
- **Screen 8:** Magnet choice (see Control Screen 7)
- **Screen 9:** Political leader feeling thermometer (see Control Screen 8)
- **Screen 10:** Country feeling thermometer (see Control Screen 9)
- **Screen 11:** Political system attitudes (see Control Screen 10)
- **Screen 12:** Muslim stereotype index (see Control Screen 11)
- **Screen 13:** Muslim ancestry belief for Saleem — incentivised (see Control Screen 12)
- **Screen 14:** Mechanism and debriefing questions — includes an additional recall item: "Are you able to recall the exact percentage of your ancestry that was traced back to India, as revealed earlier?" (see Control Screen 13)

Screen 15 — Closing

The study is now over. If you have any questions or concerns, feel free to ask us. Have a great day!