LEADER IDENTITY AND COORDINATION

Sonia Bhalotra, Irma Clots-Figueras, Lakshmi Iyer, and Joseph Vecci*

Accepted for publication at *The Review of Economics and Statistics*

Abstract: This paper examines policy effectiveness as a function of leader identity. We experimentally vary leader religious identity in a coordination game implemented in India, and focus upon citizen reactions to leader identity, controlling for leader actions. We find that minority leaders improve coordination, while majority leaders do not. Alternative treatment arms reveal that affirmative action for minorities reverses this result, while intergroup contact improves the effectiveness of leaders of both identities. We also find that minority leaders are less effective in towns with a history of intergroup conflict. Our results demonstrate that leader and policy effectiveness depend upon citizen reactions, conditioned by social identity and past conflict.

Keywords: Leader identity, in-group bias, religion, coordination failure, affirmative action, intergroup contact, conflict, India

JEL classification: P16, D70, D91, J78

*Bhalotra: University of Essex; Clots-Figueras: University of Kent; Iyer: University of Notre Dame; Vecci: University of Gothenburg. We are immensely grateful to Aditya Kuvalekar for his help, insights, and guidance. We acknowledge financial support from the International Growth Center, ESRC Grant ES/L009153/1 awarded to the Research Centre for Micro-Social Change at ISER, University of Essex (Bhalotra), the Ramon y Cajal Fellowship (Clots-Figueras), the University of Notre Dame (Iyer) and the Swedish Research Council Project no. 348-2014-4030 (Vecci). We thank Debayan Pakrashi and the Indian Institute of Technology (Kanpur) for their hospitality and logistical support. Nidhila Adusumalli, Kunal Dixit, Gunjita Gupta, Ashwin Nair and Aradhana Singh provided exceptional field supervision. We thank three anonymous referees, Amrit Amirapu, Jordi Brandts, Antonio Cabrales, Utteeyo Dasgupta, Inge van den Bijgaart, Saumitra Jha, Anandi Mani, Dilip Mookherjee, Eva Ranehill, Farah Said, Zaki Wahhaj, and seminar and conference participants at several universities for extremely useful comments. This study is registered in the AEA RCT Registry and the unique identifying number is "AEARCTR-0002353."

1. Introduction

We examine the role of leader identity in influencing economic outcomes and policy effectiveness in societies marked by social diversity. In a sharp departure from previous work on leader identity that focuses upon leader preferences and actions, we use a field experiment to isolate the role of citizen reactions, uncovering a new mechanism by which leader identity influences economic outcomes. We further show that citizen reactions to leader identity influence the success of affirmative action and intergroup contact, policies designed to empower or integrate minority groups. We also demonstrate that the success of identity-based policies varies with the history of conflict between the two groups, a result that is relevant to many post-conflict settings.

The outcome we investigate is economic coordination.[1] Coordination can be critical to resolving collective action problems and market failures. It has been shown to be important for conflict prevention, halting the transmission of viruses, changing social norms, escaping poverty traps, optimizing resource use on common land, and raising the provision of public goods.[2] While coordination problems are rife in most societies, they are of particular importance in developing countries where formal institutions to coordinate individual actions are weak and externalities from infectious disease or pollution are large. In the Indian state of Uttar Pradesh, the site of our study, previous work has observed that citizens fail to coordinate on simple tasks of community value such as sanitation or the timing of the planting of crops, with severe welfare consequences (Dreze and Gazdar, 1997; Dreze and Sharma, 1998).

[1] Coordination is distinct from cooperation as measured, say, in public goods games. Cooperation relies on voluntary contributions by individuals, driven by prosocial motivations such as altruism or trust, or specific strategies such as conditional cooperation. Coordination, on the other hand, relies on individually rational (self-interested) choices to collectively act to achieve a common goal, which may or may not result in Pareto-superior outcomes.

[2] See Schelling, 1980; Coleman, 1987; Kremer, 1993; Hoff, 2000; Hoff and Stiglitz, 2001; Adda, 2016; Bowles and Halliday, 2020. Global public goods can be characterized as having weakest-link properties and so being susceptible to coordination failure (Sandler, 1998; Nordhaus, 2006).

Using observational data to identify the impact of leader identity on coordination is difficult. First, leader identity is typically not randomly assigned and will tend to be correlated with citizen preferences. Second, with observational data there is no straightforward way of disentangling the effects of leader preferences from those of citizen reactions. Third, it can be difficult to obtain secondary data on coordination outcomes.

To overcome these issues, we implement a large laboratory-style experiment in a field setting in India, and experimentally vary the religious identity of the leader. We use a weakest link coordination game, as in Brandts and Cooper (2006). This is similar to a 5x5 stag hunt game with multiple Pareto-ranked Nash equilibria, in which individual payoffs depend positively upon the minimum effort in the group and negatively on own effort. Coordination is measured as the minimum level of effort achieved in the group. In most settings, coordination tends to occur at the lowest effort level, in a Pareto-inefficient equilibrium. Leaders can potentially create a shift to a better equilibrium by proposing a (non-binding) effort level. We conducted the game with 1028 participants in mixed-religion groups, across 44 towns in India's largest state of Uttar Pradesh. Our sample includes Muslims, a religious minority in India, and Hindus, the religious majority.

To guide interpretation of our findings, we develop a stylized theoretical framework that allows for two types of individuals--rational types, who simply maximize their economic payoff in the game, and behavioral types, who additionally respond to leader identity. The model predicts that minimum group effort will vary with leader identity if the expected fraction of behavioral types differs by citizen identity; individual choices are also guided by beliefs about how others will respond to leader identity.

Our main finding is that the introduction of Muslim leaders increases minimum group effort by 31%, coordinating group outcomes to a Pareto-superior equilibrium, while the introduction of Hindu leaders has no significant impact on coordination. This result is robust to the inclusion of town fixed effects and to several specification checks. This novel finding identifies a potential gain in aggregate output associated with having minority leaders and establishes the relevance of social identity to coordination outcomes.

Investigating mechanisms, we show, consistent with our theoretical framework, that the coordination outcome is primarily driven by citizen reactions. In particular, minority citizens are more responsive to leader identity than majority group citizens, and majority group citizens anticipate this. Thus identity moderates behavior and, once coordination is involved, beliefs over

the behavior of other members of the group matter. We rule out other possibilities: we establish that higher coordination under Muslim leaders does not arise from their proposing higher effort (a proxy for leader preferences), or from a perception that Muslims are more competent leaders, or a perception that Muslim leaders themselves supply higher effort.

The results discussed thus far describe reactions to leader identity in the control arm. We extend the analysis by randomizing two policy treatments designed to improve social integration, namely affirmative action and intergroup contact, offering the first analysis of these policies in the same experimental setting. Towns were randomly assigned to either treatment or control arms, and half the groups were randomly assigned to have Muslim leaders within each town.

In the affirmative action (AA) treatment arm, participants were told that the leadership position was subject to a quota system, and group members with a Muslim leader were informed that their leader is in a reserved position, while groups with a Hindu leader knew that their leader was not in a reserved position. We find that our main result is reversed under AA: coordination improves under Hindu but not Muslim leaders. This is most likely because the AA treatment makes salient for the Hindu majority that the Muslim minority is being favored, and this primes Hindu identity. We verify this by showing that individual effort levels of Hindus increase under Hindu relative to Muslims leaders. This result demonstrates that citizen reactions to leader identity can play a large role in the success or failure of quotas.

We facilitated intergroup contact by having mixed-identity groups collaborate in solving a puzzle before the coordination game was played. We find that this improves coordination in groups led by leaders of both religions, and more so in Hindu-led groups. Indeed, this difference counterbalances the control group difference so that, following intergroup contact, Hindu and Muslim leaders achieve similar levels of coordination. This is consistent with identity becoming salient after contact with a member of the other religion, and more so among majority group members, possibly because they have more limited exposure to minorities in a pre-treatment setting.

By virtue of conducting the experiment in areas with varying levels of previous intergroup conflict, we further identify a role for the history of intergroup conflict in shaping the effectiveness of minority group leaders. Comparing coordination across arms *within* each district, so that district-specific unobservable characteristics do not contaminate our findings, we find that a history of intergroup conflict has a detrimental effect on the effectiveness of minority (Muslim) leaders under

all policy regimes. Specifically, AA decreases Muslim leader effectiveness significantly in high conflict areas and intergroup contact increases Muslim leader effectiveness only in low conflict areas. In contrast, conflict history does not significantly influence Hindu leader effectiveness.

Our study makes a unique contribution by connecting the literature on leader identity and policy outcomes with the literature on leadership and coordination and introducing the relevance of social identity in achieving coordination. Previous work on leader identity typically identifies the combined impact of leader preferences (or actions) and citizen reactions.[3] Our experimental design uniquely allows us to isolate citizen reactions, and we find that they are a very important determinant of leader effectiveness. Our focus on coordination provides a measure of the aggregate economic impact of leader identity, in contrast to previous work that has tended to focus on whether leaders redistribute public goods or transfers towards their own group (Pande, 2003; Burgess et al., 2015). Further, our study reveals that citizen reactions to leader identity can be changed by commonly implemented policies such as AA or intergroup contact, with important implications for the success of these policies. No previous work has investigated the effectiveness of these policies as a function of leader identity. Finally, we contribute novel evidence on the relevance of conflict history, a marker of intergroup hostility, in shaping the effectiveness of leaders and of policies.[4] Overall, our work pushes forward the frontier on the open question of

[3] Many papers have examined the impact of a politician's personal identity (gender, ethnicity, caste, religion) on policy outcomes (see Chattopadhyay and Duflo, 2004; Chin and Prakash, 2011; Bhalotra and Clots-Figueras, 2014; Bhalotra et al., 2014; Brollo and Troiano, 2016; Bhavnani, 2017; Bhalotra et al., 2019). Citizen-candidate models (Osborne and Slivinski, 1996; Besley and Coate, 1997) allow leader identity to affect policy outcomes, in contrast to Downsian models where only the identity of the median voter matters.

[4] Previous research primarily focuses on the origins of conflict (see Blattman and Miguel, 2010, for a review) and its impact on growth (Rodrick, 1999), credit markets (Fisman et al., 2020), human capital (Miguel and Roland, 2011) and cooperation (see Bauer et al., 2016 for a review).

what makes some leaders more effective than others (Brandts et al., 2007; Brandts et al., 2015; Weber et al., 2001).

Our study also takes forward the experimental literature on coordination failure. While previous laboratory experimental evidence suggests that leaders are effective in improving economic outcomes (Guth et al., 2007; Levy et al., 2011; Brandts et al., 2015; Jack and Recalde, 2015; Brandts et al., 2016; Heursen et al., 2019), we provide the first evidence that leadership effects on coordination are significantly mediated by social identity in diverse societies. No previous research has examined how the effectiveness of leaders may be modified by policies designed to promote minority group representation or assimilate minorities. Previous leader-coordination studies have almost entirely been conducted in the laboratory, and the few studies of coordination games in the field have not investigated the role of leaders (Brooks et al., 2018; Afridi et al., 2020; Polania-Reyes and Echeverry, 2020). Our implementation of the coordination game in the field enables us to use real social identities rather than lab-assigned identities, to conduct the analysis in a developing country where strong formal institutions to enable coordination are less present, and to investigate how past conflict influences group behavior.

We contribute to the literature on AA policies, first by identifying that they can hamper coordination by priming majority group identity and, second, by analyzing religious identities, the existing literature having been dominated by gender quotas.[5] Of particular interest here are Gangadharan et al. (2016) and Bagues et al. (2017), who find that male identity is primed by gender quotas and committee-level exposure to women respectively, results similar to our findings that Hindu identity is primed by quotas for Muslim leaders and by intergroup contact with Muslims. Similarly, no previous paper has analyzed the effectiveness of intergroup contact in improving

[5] Many papers have examined how gender quotas influence policy outcomes, gender norms, women's aspirations and political participation (Chattopadhyay and Duflo, 2004; Adams and Ferreira, 2009; Beaman et al., 2009; Ahern and Dittmar, 2012; Beaman et al., 2012; Iyer et al., 2012; Matsa and Miller, 2013). Experimental research has primarily focused on whether gender quotas encourage women to take part in tournaments (Balafoutas and Sutter, 2012; Niederle et al., 2013; Leibbrandt et al., 2018).

coordination, nor how this varies with leader identity (and conflict history). A previous literature on intergroup contact finds that contact can change attitudes and prejudice towards the out-group, including in the case of anti-Muslim prejudice in India (Barnhardt, 2009).[6] We differ by measuring group coordination outcomes rather than individual attitudes, noting that coordination takes into account others' actions and attitudes.

Finally, we contribute to research in psychology, sociology and economics showing that social identity affects individual economic choices, and that the influence of identity on behavior varies with primes that make group identity salient (see, among others, Akerlof and Kranton, 2000; Benjamin et al., 2016; Hungerman, 2014; Chen and Chen, 2011). We take this literature forward by studying group outcomes rather than individual economic choices, using real identities rather than laboratory-primed ones, and assessing the impact of different policy regimes on identity priming.

The rest of the paper is structured as follows. Section 2 provides contextual information and describes the data collection. Section 3 delineates the experimental design. Section 4 lays out a theoretical framework to structure and interpret our results, Sections 5 and 6 present the empirical results, and Section 7 concludes.

2. Context, Site Selection and Subject Recruitment

India is a religiously diverse country, with Hindus constituting the majority religion (79.8% of the population in the 2011 census) and Muslims forming the largest religious minority (14.2% of the population). Compared to the Hindu majority, Muslims in India are less educated, more likely to live in urban areas and more likely to be victims of inter-religious violence. Relative to their population share, Muslims are under-represented in political office, the police and the judiciary (see Appendix B1 for details).

We conducted field work with 1028 subjects across 44 towns from four districts in Uttar Pradesh, India's largest state with more than 200 million inhabitants. Each district-pair was

[6] Allport (1954) articulates the potential effects of intergroup contact, and Paluck et al. (2019) review the empirical literature. Only a few recent studies use random assignment of groups to examine intergroup contact in a developing country setting (Corno et al., 2018; Lowe, 2020; Rao, 2019; Scacco and Warren, 2018).

composed of one high conflict and one low conflict district. The 44 towns were randomly assigned to three different treatment arms: 14 towns were retained as control, and 15 each were assigned to the intergroup contact and the affirmative action treatments (Appendix Table A1). The assignment was performed within each district (implicitly stratifying by high/low conflict status); we further stratified by Muslim population proportion and total population. We then recruited study participants from both Hindu and Muslim sections of each town. Participants' religion was identified from their names by our research assistants and verified using the pre-experiment questionnaire (see Appendix B2 for details of site selection and subject recruitment). Participants were given no information about the identities of other participants in the experiment.

3. Experimental Design

Each experimental session contained a pre-experiment survey and three tasks: a puzzle task, a weakest link coordination task (run across six rounds), and a social norms elicitation task. Out of the three tasks, one was chosen randomly for payment. The average payoff was ₹ 610 including a ₹ 200 show up fee. This constitutes about 2.5 days' wage for a semi-skilled laborer. The session concluded with a survey of attitudes and respondent characteristics. Subjects knew that the session had multiple stages but were not given instructions about any particular stage until reaching that stage (see Appendix B3 for complete experiment instructions).

3.1. Pre-Experiment Survey and Puzzle Task. Prior to commencement of the incentivized tasks, subjects answered a brief survey about their personal characteristics, namely height, eye color, hair color and religion. The questionnaire was primarily designed to check the religion of the participant (see layout in Appendix B3), and other questions were included to avoid making the research question explicitly salient, which might induce socially desirable response bias or experimenter demand effects. Using surveys and lists to make identity implicitly salient without making it explicitly salient is standard procedure in the identity salience literature (Steele and Aronson, 1995; Shih et al., 1999; Benjamin et al., 2016; Cohn et al., 2015). While it is difficult to know how participant responses would change if they were explicitly aware that the experiment was about religion, we should note that all subjects filled out the same questionnaire. So, even if the questionnaire made the religious focus of our research explicitly salient, this would be the same across different types of leaders or different policy arms.

After the pre-experiment survey, all subjects participated in a 12-piece jigsaw puzzle task. Participants completed the task individually in the control and affirmative action treatment groups,

and in pairs in the intergroup contact treatment group. Our objective was to suppress competitiveness and have cooperative intergroup contact, since the latter has been shown to reduce prejudice (Paluck et al., 2019; Lowe, 2020). The time given for the puzzle assembly was twelve minutes and almost all participants were able to complete the task successfully in this time.

3.2. Weakest Link Coordination Task. The task structure is closely related to the minimum effort corporate turnaround game designed by Brandts et al. (2006), which is based on the minimum effort or weakest link coordination game of Van Huyck et al. (1990). It was conducted after the puzzle task.

3.2.1. Group Formation. Individuals were assigned to four-member groups (each called a "firm"), comprised of two Hindu and two Muslim "employees." Participants did not know, nor could they observe, who the other three in their group were. We did not provide information to participants on the identity of their group members, and they were explicitly told that the people sitting on their mat were not part of their firm.

3.2.2. Effort Choices and Payoffs. The task is run across six periods. In each period employees decide how many hours (x) to devote to firm activities. Their choices vary between 0 and 20 in intervals of 5: $x_i \in \{0,5,10,15,20\}$. It is noteworthy that choices are not actual hours worked but effort choices with payoff consequences. Employees' payoffs for each period depend negatively on their own effort choice and positively on the minimum effort of all individuals in the group:

$$(1)\ p_i = 500 - 25x_i + [min\ (x_i, X_{-i}) * 40]$$

where x_i is player i's own effort (number of hours) and X_{-i} is the vector of all other players' effort choices. The payoff table is illustrated in Appendix Table A3, where the units are Indian rupees (1USD ≈ ₹68). Participants were shown the payoff table but not the payoff equation. Under this payoff structure, coordinating on any of the available effort levels is a Nash equilibrium.

Note that it is only worthwhile for profit maximizing employees to raise their own effort level if this will increase the minimum effort of the firm.[7] Given this, previous work has found that play often evolves towards the payoff-dominated equilibrium in which all players choose the lowest

[7] For a profit maximizing employee to increase their effort by 1 unit e.g. from 0 to 5 hours, they must believe there is an 85.5% probability that each of the other three employees increase their effort. To derive this probability, we solve for p where $500 = 375(1 - p^3) + 575(p^3)$.

possible effort level (Brandts et al., 2006). The task is split into two stages. The first stage repeats the coordination game described above across four periods (rounds). Employees work in the same firm across all periods. At the end of each period subjects are informed of the firm's minimum effort. Employees are never informed of individual firm members' effort choices. The first stage is designed to induce coordination on an inefficient equilibrium with low levels of effort, which we label "coordination failure" (Brandts et al., 2015).

3.2.3. Leader Assignment. The second stage introduces a leader, and runs for two further periods. The leader's role within each firm is to suggest a non-binding number of hours to work. Leaders do not have the scope to communicate with their employees beyond proposing an effort level, similar to other papers in the "leading by example" literature (Güth et al., 2007, Gächter et al., 2012; Levy et al., 2011). All leaders are appointed and participants cannot elect or change the leader. Firm employees are informed of the leader's proposal, but not the actual effort choice of the leader. Employees are also provided information about the characteristics of their firm leader taken from the pre-experiment survey, namely height, eye color, hair color and religion. The only characteristic that varies across leaders is their religion.

Half of the firms in each session are assigned Hindu leaders and half are assigned Muslim leaders. Leader identity is randomly allocated, and the player in each group who will be the leader is also randomly selected conditional on their religion. Our estimates for the impact of introducing a leader are thus specific to leader identity. It is important to remember that the religious composition of firms is the same across all firms, regardless of the leader's religion. Participant characteristics with regard to demographics, education, income and religiosity are balanced across groups with Hindu or Muslim leaders (Appendix Table A4). Leader characteristics other than religion, in particular, gender, age, and family income are balanced across Hindus and Muslims, though Muslim leaders are less likely to have gone to college and more likely to pray several times a day, similar to the overall population.

After being informed of the leader's effort proposal and leader characteristics at the start of the fifth period, similar to the earlier periods, employees are informed of their group's minimum effort in the previous period. All employees including the leader must then decide how many hours to work. The leader's effort, just like the effort of other employees, is not visible to the group. The coordination game is repeated for two periods with the same leader, but with a new effort proposal by the leader in each period. If this task is selected for payment, players are paid their coordination

game payoffs from two randomly selected periods. We have two additional treatment arms (described below) where the same weakest link game is played, but with changes to the environment in which the leaders operate, designed to mimic commonly proposed policy interventions.

3.2.4. Affirmative Action (AA) Treatment. Affirmative action policies, such as quotas, are common in both government and business to increase participation of disadvantaged or minority groups. As described earlier, 15 of our 44 towns were randomly assigned into an AA treatment arm. The game is conducted exactly as in the control arm described above, with one important exception. Upon the introduction of a leader at the beginning of period 5, subjects are told that "similar to many government positions, 50% of the leadership positions in this game will be reserved. Reservation will be made based on some characteristic in the initial survey."

Along with information on the leader's characteristics (height, eye color, hair color, religion), employees with a Muslim (Hindu) leader are also informed that their leader is in a reserved (unreserved) position.[8] While participants are not explicitly told that the reservation is based on religion, about 70% of respondents in the post-experiment survey correctly identified that the reservation was done on the basis of religion, with the rest citing other leader characteristics or saying "don't know."[9] It is important to reiterate that here, as in the control arm, by design all groups contain two Hindu and two Muslim participants. By comparing the control and the AA treatment arms (and thus effectively comparing a Muslim leader with a Muslim leader who is leading through a quota), we can measure whether people behave differently when they believe their leader is in a position due to an affirmative action policy. In our setting, AA does not change the composition of leaders, it only makes the reservation policy salient in participants' minds.

3.2.5. Intergroup Contact Treatment. We also investigate the impact of a randomized intervention that increases intergroup contact on citizen responses to leaders of different religious identity. The key difference between this treatment and the control arm is the implementation of the puzzle task.

[8] Leadership positions could in theory also be reserved for Hindus. However, we do not analyze this possibility since we are not aware of the existence of such a policy in any part of India.

[9] Our intent-to-treat estimates will therefore under-estimate the impact of religion-based reservation.

Unlike in the control arm and AA treatment, where puzzles are assembled individually, subjects in the contact treatment assemble the jigsaw puzzle with a partner from the other religion. The puzzle partner is a person sitting on the same mat as the participant, and therefore not a member of the same firm (see Appendix Figure A3). Subjects are encouraged to talk with their partner during the 12 minutes allowed for the puzzle. By comparing outcomes across the control and the contact treatment arm, we can infer whether intergroup contact changes the impact of leader identity on coordination.

Subjects in the sample towns often live in separate Hindu and Muslim neighborhoods, which limits interaction between the two communities. Nevertheless, our survey confirms that only 14% of participants incorrectly identified the religion of their puzzle partner. A potential concern is that interacting with anyone, not necessarily from a different religion, prior to the coordination game may affect coordination. As we shall see below, we can reject this concern because we see no differences in minimum effort in the contact vs the control arms in the periods *before* the leader is introduced.

3.3. Norms Elicitation Task

To measure Hindu and Muslim religious norms we follow Krupka and Weber (2013) and Gangadharan et al. (2016) and conduct a social norms coordination task. Participants are asked a set of questions related to behavior in the weakest link coordination task. Participants are asked to rate the social appropriateness of a Hindu or Muslim employee working 0, 10 or 20 hours under a Hindu or Muslim leader. Appropriateness ratings are measured on a 4-point ordered scale, consisting of the following options: very socially inappropriate, somewhat socially inappropriate, somewhat socially appropriate and very socially appropriate. These questions can thus be used to evaluate what people within our sample towns believe are the appropriate behaviors between Muslims and Hindus.

4. Theoretical Framework

In order to explain our results, we describe a stylized theoretical framework to help us understand the role of leader identity in changing individual behavior and hence the group outcome in the coordination game. Specifically, our model incorporates the role of citizen reactions to leader identity in shaping aggregate outcomes. Akerlof and Kranton (2000) introduced the role of identity in economic decision making. In a setting broadly related to ours, Benjamin et al (2016) show that priming religious identity can change individual economic choices, but utility maximization in

their model does not involve strategic interactions as in our coordination game. We build on an extensive reputational literature by introducing "behavioral types"—players who choose their effort based not only on their economic payoff but also on leader identity. Rational players, in contrast, maximize their payoff, taking into account the presence of behavioral types. As is standard in the reputational literature, we assume that behavioral types constitute only a tiny fraction of the population.

Recall that individuals in our field experiment are either Hindus or Muslims, and four players are randomly chosen from the population to constitute a group. Individuals do not know the identity or religion of other individuals in their group. The group plays the coordination game as described earlier, with individuals choosing among possible effort choices over a continuum [0,W].[10] All players selecting any effort level in this range produces a Nash equilibrium. We assume that in situations where there is no leader, all individuals will use the concept of *risk-dominance* as an equilibrium selection device in the coordination game. This means that players have uniform beliefs over the others' effort strategies and that this guides their own effort choice. This assumption is similar to that made in coordination games with investment decisions under incomplete information.[11] The coordination game payoff is maximized when the individual player matches the minimum effort of the other players. So, the optimal effort choice involves calculating the expected value of the minimum of the other three players. Under the risk-dominant criterion, players assume that all the other players are randomizing uniformly over $[0, W]$, and hence the equilibrium effort choice can be calculated as $x^* = \frac{W}{4}$.

The optimization is different for behavioral individuals. Behavioral individuals will follow the leader's proposal if the leader is from their own religion and will choose effort level x^* when the leader is from the other religion. We assume that the leader's proposal will be greater than the

[10] In our experiment, we have discrete effort choices for ease of implementation in the field.

[11] See the literature on equilibrium selection and global games (Carlsson and van Damme, 1993; Morris, Shin and Yildiz, 2016).

minimum effort in the no-leader equilibrium x^*, which is true in our data.[12] Rational players will choose an effort level that optimizes their coordination game payoffs, keeping the responses of behavioral types in mind. A fraction a_H of Hindu individuals are behavioral types, as are a fraction a_M among Muslims. The overall fraction of behavioral types in the population is thus $\tilde{a} \coloneqq pa_M + (1-p)a_H$, where p is the population share of Muslims. a_H and a_M are small enough that the probability of more than one behavioral individual in a group of 3 or 4 is close to zero, so that rational players can behave as though there is at most one behavioral player in the rest of the group.[13]

The optimal effort choices for a rational player when there is a Hindu leader (x_H), and when there is a Muslim leader (x_M), can then be derived as follows:

$$x_H = (1-\tilde{a})^3 x^* + 3(1-\tilde{a})^2 \left((1-p)a_H \frac{W}{3} + pa_M x^{**} \right) \quad (2)$$

$$x_M = (1-\tilde{a})^3 x^* + 3(1-\tilde{a})^2 \left(pa_M \frac{W}{3} + (1-p)a_H x^{**} \right) \quad (3)$$

Here the first term reflects the probability of all three other players being rational (and the optimal choice is therefore x^*); the second term is the expected optimal choice under the assumption of one player in the group being behavioral. Under Hindu leaders, if the behavioral player is Hindu (which happens with probability $(1-p)a_H$), that player will follow the leader and choose effort above x^*, and hence the rational player only needs to match the expected minimum effort of the other two rational players, which is $\frac{W}{3}$ ($> x^*$). On the other hand, if the behavioral player is Muslim (which happens with probability pa_M), that player chooses x^*. Then the rational player needs to match the expected minimum value of x^* and the choices of the other two players, which can be calculated as $x^{**} = \frac{9}{16}x^* + \frac{3}{16}\frac{x^*}{2} + \frac{1}{16}\frac{x^*}{3} < x^*$. A similar logic applies to computation of x_M.

Examination of equations (2) and (3) tells us that optimal effort in the presence of a leader will be higher than the optimal effort without a leader only under certain conditions, and that the

[12] Less than 2% of all leader proposals in rounds 5 and 6 are lower than the minimum group effort in round 4 (prior to leader identity being announced).

[13] If $\tilde{a}$ is 0.10, then the probability of more than one behavioral player is only 5.23% in a group of size 4 and 2.8% in a group of size 3. If $\tilde{a}$ is 0.05, these probabilities are 1.4% and 0.7% respectively.

change in optimal effort will depend on the identity of the leader as well as the fraction of behavioral individuals in each religion. This is a contribution we make to the coordination literature that typically overlooks social diversity within groups.

We can be more specific in our predictions: as long as pa_M is sufficiently greater than $(1-p)a_H$, minimum group effort will be higher under Muslim leaders after leader identity is made public, and the increase in minimum group effort under Muslim leaders (compared to the situation without a leader) will be greater than the increase in minimum group effort under Hindu leaders. Note that one way this sufficient condition is satisfied is if the fraction of behavioral types among Muslims (a_M) is much greater than the fraction among Hindus (a_H). This is likely to be the case, since previous literature has established that members of population minority groups are more likely to display "in-group bias" and majority groups are unlikely to do so (Bisin and Verdier, 2001; Gupta et al, 2018; Berge et al., 2019). We now proceed to test these hypotheses using the data from our field experiment.

5. The Impact of Leader Identity on Coordination Outcomes

5.1. Regression Specification

Our main outcome variable is the minimum effort exerted in the group in each round. This is the key determinant of player payoffs and is the standard measure of coordination in the weakest link literature. We test whether leaders improve coordination by estimating the following specification:

$$(4) \quad MinGroupEffort_{kjt} = \alpha + \beta Leader_{kjt} + G'_{kj}\gamma + \varepsilon_{kjt} \text{ ; t=1,2,...6}$$

where $MinGroupEffort_{kjt}$ is the minimum effort exerted by group k in town j in period t, and $Leader_{kjt}$ is a dummy variable that takes value one for periods 5 and 6, when a leader is introduced. This regression therefore compares the group's minimum effort in periods with a leader to periods without a leader. G_{kj} is a suite of control variables that includes town fixed effects, demographic controls (average age, education, gender mix and monthly household income of the group members) and a control for religiosity based on prayer frequency. Standard errors are clustered at the group level to account for within-group correlation in outcomes across different periods.

We run specification (4) separately for Hindu and Muslim leaders to test whether leader identity matters in achieving better coordination. Since comparisons between later and earlier rounds maybe influenced by round effects and because our main interest is in the comparison between Hindu and Muslims (where pre-leader rounds have similar rates of coordination) we also

run a regression on the combined data, and include an interaction term $Leader_{kjt} * MuslimLeader_{kjt}$ in order to test whether the increase in minimum effort under Muslim leaders is higher than under Hindu leaders. Recall that half of all groups *within* each town are randomly assigned to have Muslim leaders. We also run a robustness test in which we restrict the analysis to periods 5 and 6, and control for the group-specific minimum effort in period 4 and the leader's proposals in periods 5 and 6. This enables us to test whether the impact of the policies can be attributed to differences in the leader's proposal, or to differences across groups in the coordination outcome in previous rounds.

5.2. Leader Identity and Coordination in the Control Group

Consistent with the corporate turnaround game literature, we find that groups coordinate on the low-effort equilibrium in the absence of a leader, which we call "coordination failure." The average minimum group effort is less than three hours at the end of period 4 (Figure 1A). We find that introducing Muslim leaders significantly improves minimum group effort in periods 5 and 6 (Figure 1A). The efficiency gain in Muslim-led groups is large: minimum group effort increases by 1.07 hours, compared to the pre-leader average of 3.45 hours in periods 1 through 4 (Table 1, column 1).[14] The estimates are robust to controlling for the demographic and religious characteristics of group members, consistent with our randomized assignment of leader identity (column 3). In contrast, the introduction of Hindu leaders does not improve minimum group effort (Figure 1A), leading to a statistically insignificant decline of 0.488 hours (Table 1, column 2). The difference in coordination gains between Muslim and Hindu leaders is statistically significant (column 5). Our results are consistent with a higher expected fraction of behavioral types among Muslims-- our model predicts that Muslim leaders will have a bigger impact than Hindu leaders in this situation.

The results are robust to using town random effects instead of town fixed effects, using an ordered probit specification rather than OLS, and to controlling for town*mat fixed effects to

[14] The effect of Muslim leadership on minimum effort is lower than in the lab experiment of Brandts et al. (2015), which could be due to differences in the context (developed country lab experiment vs developing country field setting) or because leaders in our experiment could only communicate a numerical proposal rather than more detailed messages to participants.

ensure that participants are correctly responding to the effort choices of their firm members rather than the effort choices of those seated on the same mat during the experiment. In a post-game survey, we elicited each participant's trust of people of the other religion by asking whether they would like to have a neighbor of a different religion. We find no evidence that our results are driven by differences in cross-religion trust, see Appendix Table A10.

5.3. Mechanism: Citizen Reactions to Leader Identity

Our model generates predictions based on individual responses to leader identity (among the "behavioral types") and other players' optimal strategies in light of these expected reactions. We show that our results cannot be explained by alternative hypotheses, such as the reactions to leader proposals or to the perceived competence of leaders. First, we verify that our results cannot be explained by differences in leader proposals. Our data reveal that Muslim leaders propose 10.5 hours on average, compared to 9.4 hours for Hindu leaders (Appendix Table A5). This difference is not statistically significant (Appendix Table A6, column 1), and a Kolmogorov-Smirnov test shows that the distributions of proposals by leader identity are not statistically different (p value= 0.452; Appendix Figure A4). We confirmed that the significantly different results under Muslim leaders compared to Hindu leaders hold even when we control for leader proposals and for minimum effort in period 4 by restricting the sample to periods 5 and 6 (Table 1, column 6). Second, we find that our results are not driven by perceptions of higher competence of Muslim leaders, by greater or lesser exposure to real-life Muslim leaders, or by beliefs about the hours worked by the leader (see Appendix Table A7).

We now show two pieces of evidence consistent with our theoretical framework. Our theoretical framework predicts that we should see increased effort in Muslim-led groups when pa_M is sufficiently greater than $(1-p)a_H$. When is this more likely to happen? Based on previous research, we expect a lower fraction of Muslim behavioral type individuals (a_M) in towns where Muslims form a higher fraction of the population, i.e. where p is high. This means that the inequality is less likely to be satisfied when p is very high or very low, and more likely to be satisfied at intermediate values of p. In Appendix Figure A5, we graph the increase in minimum effort under Muslim leaders against the town's Muslim population share. The results are in line with this prediction, namely that towns with very high or very low Muslim population share show smaller increases. This is suggestive rather than conclusive as, with a small number of towns, the differences are not statistically significant.

Importantly, we show that our results are consistent with the extent to which individuals expect others to react to leader identity. As part of the social norms task, we asked respondents to rate (on a scale of 1 to 4, with 4 being the highest) how "socially appropriate" others would consider it to be for a Hindu or Muslim employee to choose the maximum effort level under a Hindu or a Muslim leader. We find that Hindu subjects rate a maximum effort as less "socially appropriate" from a Muslim employee when faced with a Hindu leader instead of a Muslim leader, a statistically significant difference (Appendix Table A8, panel A). In other words, Hindus *expect* Muslims to significantly change their behavior based upon leader identity. In contrast, Muslim employees do not expect any difference in effort from themselves under Hindu vs Muslim leaders. They do rate Hindus providing maximum effort under Muslim leaders to be less socially appropriate than under Hindu leaders, but the difference is not statistically significant (Appendix Table A8, panel B).

Given this structure of beliefs, our model predicts that rational Hindu employees are much less likely to increase effort under Hindu leaders (since they are more likely to expect the Muslims to be "behavioral"), as compared to Muslim employees under Hindu leaders. The reverse is unlikely to be true, since Muslim subjects do not assign a statistically different rating to Hindu employees' appropriateness under different leaders. We can investigate these predictions by examining how individual effort decisions respond to leader identity:

$$(5)\ IndividualEffort_{ikjt} = a + bLeader_{kjt} + G'_{ikj}g + w_{ikjt};\ t = 1,2,\ldots,6$$

where $IndividualEffort_{ikjt}$ is the effort choice of individual i in group k (of town j) and period t. As before, G_{ikj} includes town fixed effects, demographic controls and religion, and standard errors are clustered at the group level. Note that individual effort choices depend not only on their expectations of how other individuals in the group will react to the leader's proposal and the leader's identity, but also the individual's own reaction to leader identity (if they are a behavioral type).

Consistent with our model predictions, we find that Hindu employees show a decline in effort when a Hindu leader is introduced, while Muslim employees do not show any significant change in effort (Table 2, columns 1 and 2). The difference in response between Hindu and Muslim employees is significant at the 10% level. In contrast, Muslim employees exhibit a statistically significant increase in effort under Muslim leaders, while Hindu employees show a non-significant increase in effort (Table 2, columns 3 and 4). However, the difference between Hindu and Muslim employee response to a Muslim leader is not statistically significant. To summarize, our

investigations indicate that our main finding, that Muslim leaders induce greater coordination towards Pareto-superior equilibria, is primarily driven by citizen reactions that are shaped by expectations of Muslims having stronger in-group bias (i.e., they are more likely to expect the Muslims to be behavioral).

6. Coordination Responses to Leader Identity under Policy Assignment

6.1. Regression Specification

We examine whether leader effectiveness varies across policy environments by comparing coordination outcomes for a given leader identity across the different treatment groups as follows:

(6) $MinGroupEffort_{kjt} = b_0 + \beta_1 Leader_{kjt} + \beta_2 Leader_{kjt} * AA_j + \beta_3 Leader_{kjt} * Contact_j + X'_{kj}\gamma + \varepsilon_{kjt}$; $t = 1,2,\ldots,6$

In equation (6), AA_j is a dummy that equals one if town j was randomly assigned to the affirmative action treatment and $Contact_j$ equals one if the town was randomly assigned to the contact treatment. We estimate equation (6) separately for Muslim and Hindu leaders. β_1 then captures the impact of the leader on coordination in the control arm, β_2 estimates the differential impact of the leader in an AA environment and β_3 estimates the differential impact of the leader in an environment with pre-game contact between members of the different religions.

6.2. *Affirmative Action*

Our empirical results show that affirmative action (AA) policies, a commonly suggested solution to improve integration and opportunities for minorities, can in fact strongly reduce the effectiveness of minority leaders. Minimum effort increases by a statistically insignificant 0.227 hours (1.067-0.840) under Muslim leaders (Table 3, column 1). In contrast, minimum group effort increases by a large and statistically significant 2.391 hours under Hindu leaders (column 2).

Our model suggests that this may result from an increase in the expectation that Hindus include behavioral types (a_H), which would lead to increased minimum effort under Hindu leaders. This will be more likely to happen if a_H increases so much that the sufficient condition for our hypothesis is reversed, i.e. $(1-p)a_H \gg pa_M$. Such a change in behavior is consistent with previous research which finds that AA policies result in a strengthening of in-group bias among AA non-recipients (Gangadharan et al., 2016 show this in the context of gender quotas in India). This effect is likely to be heightened in contexts where people believe that quota recipients are not suitable for leadership roles because they are less skilled or not truly disadvantaged (Ip et al.,

2020). To investigate the behavior of Muslim and Hindu individuals, we examine individual effort as a function of the religion of the group leader across different treatment arms:

(7) $IndividualEffort_{ikjt} = f_0 + f_1 Leader_{kjt} + f_2 Leader_{kjt} * AA_j + f_3 Leader_{kjt} * Contact_j + X'_{ikj} g + z_{ikjt}$; $t = 1,2,\ldots,6$

Our estimates show that Hindu employees chose significantly higher effort levels in Hindu-led groups under the AA treatment compared to the control group i.e. $f_2 > 0$ (Appendix Table A9, column 4). The results are sharper when we restrict to periods 5 and 6, and include controls for the leader's proposal and for minimum group effort in the pre-leader period, similar to columns 3 and 4 of Table 3. We see that both Hindu and Muslim employees reduce effort under Muslim leaders in the AA environment compared to the control group (Appendix Table A9, columns 5 and 6). Similarly, both Hindu and Muslim employees increase effort under Hindu leaders in the AA environment (columns 7 and 8). The change in effort is statistically significant only for Hindu employees but the fact that effort choices of both sorts of citizens move in the same direction is consistent with a generalized perception of greater a_H and lower a_M under AA. The higher response of Hindu employees to Hindu and Muslim leaders is consistent with a greater fraction of behavioral types among Hindus under AA.

As with the control group results, we verify that these differences do not arise because of differences in leader proposals across treatment arms. Muslim leaders do not make statistically different proposals from Hindu leaders across any of the treatment arms (Appendix Table A6, column 2). To address the possible concern that the groups assigned to the different policy environments undergo different rates of learning over the course of the game, we restrict the estimation sample to the last two periods and control for leader proposals and minimum effort in the previous period of the game, period 4 (Table 3, columns 3 and 4). The concern is allayed, and the coefficients on Leader*AA are now statistically significant for both Muslim and Hindu leaders.

6.3. Intergroup Contact

Another common policy to improve integration is to encourage interaction between groups. We find that intergroup contact improves minimum group effort under both Muslim and Hindu leaders compared to the control group, but the difference is larger for Hindu leaders. For Muslim leaders, intergroup contact increases minimum group effort by an additional 1.007 hours compared to the control group, but this difference is not statistically significant (Table 3, column 1). Under Hindu leaders, intergroup contact results in a large and statistically significant increase of 2.755

hours of minimum group effort, compared to the control group (column 2). This difference counterbalances the better performance of Muslim leaders in the control group--minimum group effort in the contact treatment is almost the same across Muslim and Hindu leaders (see Figures 1B and 1C).

This result in the context of our framework, corresponds to an increase in both a_H and a_M, and a particularly large increase in a_H. In other words, both Hindus and Muslims are more responsive to leader identity after intergroup contact, but the effect is stronger for Hindu individuals. Examining potential mechanisms and in particular individual effort again, we find no significant increase in individual effort levels under Muslim leaders, but a strong and significant increase in individual effort for both Hindus and Muslims under Hindu leaders (Appendix Table A9, columns 5-8). Our results suggest that expectation of change in Hindu behavior after contact with Muslims is higher than the expectations of change in Muslim behavior after contact with Hindus. This asymmetry probably arises because Muslims are a minority community overall, and hence interactions with Hindus may be more common for them than the reverse. The mechanism our results indicate—namely the increase in identity-based response—is different from that emphasized in previous research on intergroup contact which has focused on whether such contact can reduce prejudice or affect attitudes such as pro-sociality, trust or egalitarianism (Rao, 2019, Finseraas et al. 2020, Paluck et al., 2019).

As before, we verify that leader proposals are not driving our results. We again find no significant differences in leader proposals under intergroup contact compared to the control group (Appendix Table A6, column 2). The coefficients of interest are larger in magnitude and statistically significant after controlling for leader proposals (Table 3, columns 3 and 4). In fact, once we control for leader proposals, the improvement in coordination in Muslim-led groups becomes statistically significant.[15]

6.4. Does Conflict History Matter?

[15] We also re-estimated equation (4) for the contact group participants, interacting leader with female puzzle partner. We find that participants with female puzzle partners do not respond differentially to the introduction of a leader (Appendix Table A11).

The results discussed so far average across areas with different histories of Hindu-Muslim conflict. In this section, we examine how the relationship between coordination outcomes and leader identity varies in the control vs treatment arms with whether the district had high or low exposure to intergroup conflict over the period 1980-2010. This is pertinent since such policies may be implemented to ameliorate historical grievances. The aim of this exercise is not to compare outcomes in high and low conflict areas directly, since many other observable and unobservable characteristics may differ across high conflict and low conflict areas. Instead, we take advantage of the fact that we randomly allocated towns to different treatments within districts characterized by high vs low conflict and we compare behavior across these randomized treatments *within* each conflict setting. Thus, we do not aim to identify the causal impact of conflict history, but rather the casual impact of the two policies under Muslim vs Hindu leaders across areas with different conflict histories.

We find a consistent pattern of results that suggest a lower responsiveness of Muslims to leader identity in high conflict areas, and a higher responsiveness of Hindus to leader identity in those same areas (Figure 3). In particular, we see that the increases in minimum group effort under Muslim leaders are uniformly higher in low conflict areas for all the three policy environments (see Table 4, columns 1 and 2). In particular, this means that our earlier finding that Muslim leaders improve coordination in the control and contact treatments emerges mostly from low conflict areas (column 1), and there is a significant decline in coordination in the AA arm in high conflict areas (column 2). This is consistent with a_M being uniformly lower in high-conflict areas. One reason for this may be that Muslims are less willing to exercise their religious preferences precisely because of the history of religious conflict, in which existing evidence suggests that they are usually the victims (Mitra and Ray, 2014).

The impact of Hindu leaders on coordination is less sensitive to conflict history, with the exception that coordination improvements in the contact arm are higher in high conflict areas (Table 4, columns 3 and 4). This is consistent with higher a_H, potentially because of behavioral type Hindu reactions to leader identity being primed by intergroup contact in areas where religious tensions are stronger. Our earlier findings that Hindu leaders do not improve coordination in the control arm, but do under both the AA and contact treatments, holds in high and low conflict areas. These results are robust to controlling for leader proposals, and thus are driven by citizen reactions to leader identity (Appendix Table A12). Overall, our results suggest that the history of intergroup

conflict matters for leader effectiveness and policy effectiveness. Of note is the result that AA, which often aims to increase leadership roles for minorities, may be detrimental for minority leader effectiveness.

7. Conclusions

We provide the first investigation of how leader effectiveness in achieving economic coordination in diverse societies varies with leader identity. We implemented a lab-in-field experiment in India's largest state, where we randomly assigned towns to two policy treatments and a control group. The random assignment was stratified by district, to allow us to compare the impact of policy treatments across districts with a history of high vs low intergroup conflict.

We find that minority leaders improve coordination (measured as minimum group effort), but majority leaders do not. This is primarily the result of citizen reactions to leader identity, rather than differences in leader actions. The mechanism driving citizen reactions appears to be higher expectations of responsiveness to leader identity among members of the minority group. In contrast to the greater responsiveness of Muslims to leader identity in the control group, we find that identity-based responsiveness to leaders appears to increase for members of the majority group in the presence of policies designed to improve social integration of the minority group. In fact, the control group results are reversed under affirmative action that is perceived to assign leadership posts to Muslims, with coordination becoming higher under Hindu leaders. A policy of intergroup contact leads to higher coordination under all leaders, but the gains are larger under Hindu leaders.

We find that social integration policies enable minority leaders to be more effective in low conflict areas. Specifically, in low conflict areas and *only* in low conflict areas, Muslim leaders serve to improve coordination outcomes in each of the three experimental arms. The effectiveness of Hindu leaders remains evident in the AA and contract arms but does not vary significantly by conflict history.

Pulled together, our findings provide compelling evidence that social identity influences behavior in a way that leads to different economic choices than would emerge from maximization of individual economic payoffs. Our particular contribution is to demonstrate this in the context of leadership. We show that citizen reactions to leader identity influence coordination outcomes. In general, we find that the minority group shows stronger responsiveness to leader identity, except in areas with a history of intergroup conflict. However, policies designed to integrate minorities

tend to activate social identity responses in the majority group in ways that can potentially reduce the effectiveness of these policies.

Our findings contribute novel evidence to research on leader identity, coordination failure, social integration policies and religious conflict, bridging key aspects of these domains of research. They provide unique evidence on the role of citizen reactions to leader identity in heterogeneous communities using a non-student sample. They further provide useful guidance for policy, identifying conditions under which leaders of minority vs majority groups may be constrained in resolving coordination problems.

Our results suggest many directions for future research in field settings. These include allowing leaders to use unrestricted communications rather than specific proposals (as in Brandts et al., 2016), examining whether raising the monetary stakes or changing group sizes affects leaders' effectiveness, and whether our results for Hindus and Muslims in India generalize to other social majority or minority groups. Recent research suggests that social identity itself may be amenable to policy initiatives (Miguel, 2004; Blouin and Mukand, 2019). While this is beyond the scope of our paper, it can also be a fruitful direction for future research.

References

Adams, R. B., and Ferreira, D., "Women in the Boardroom and Their Impact on Governance and Performance," *Journal of Financial Economics*, 94:2 (2009), 291-309.

Adda, J., "Economic Activity and the Spread of Viral Diseases: Evidence from High Frequency Data," *Quarterly Journal of Economics*, 131:2 (2016), 891-941.

Afridi, F., Dhillon, A., Li, S.X., and Sharma, S., "Using Social Connections and Financial Incentives to Solve Coordination Failure: A Quasi-Field Experiment in India's Manufacturing Sector," *Journal of Development Economics*, 144 (2020), 102445.

Ahern, K.R., and Dittmar, A., "The Changing of the Boards: The Impact on Firm Valuation of Mandated Female Board Representation," *Quarterly Journal of Economics* 127:1 (2012), 137–197.

Akerlof, G., and Kranton, R., "Economics and Identity," *Quarterly Journal of Economics*, 105:3 (2000), 715–753.

Allport, G., *The Nature of Prejudice* (New York: Basic Books, 1954).

Bagues, M., Sylos-Labini, M., and Zinovyeva, N., "Does the Gender Composition of Scientific Committees Matter?" *American Economic Review*, 107:4 (2017), 1207-1238.

Balafoutas, L., and Sutter, M., "Affirmative Action Policies Promote Women and Do Not Harm Efficiency in the Lab," *Science*, 335 (2012), 579-582.

Barnhardt, S., "Near and Dear? Evaluating the Impact of Neighbor Diversity on Interreligious Attitudes," Working paper (2009).

Bauer, M., Blattman, C., Chytilová, J., Henrich, J., Miguel, E., and Mitts, T., "Can War Foster Cooperation?" *Journal of Economic Perspectives,* 30:3 (2016), 249–274.

Beaman, L., Chattopadhyay R., Duflo E., Pande R., and Topalova P., "Powerful Women: Does Exposure Reduce Bias?" *Quarterly Journal of Economics,* 124:4 (2009), 1497-1540.

Beaman, L., Chattopadhyay R., Duflo E., Pande R., and Topalova P., "Female Leadership Raises Aspirations and Educational Attainment for Girls: A Policy Experiment in India," *Science*, 335:6068 (2012), 582-586.

Benjamin, D.J., Choi, J.J., and Fisher, G., "Religious Identity and Economic Behavior," *Review of Economics and Statistics*, 98:4 (2016), 617–637.

Berge, L, Bjorvatn, K., Galle, S., Miguel, E., Posner, D., Tungodden, B., and Zhang, K., "Ethnically Biased? Experimental Evidence from Kenya," *Journal of the European Economic Association*, 17 (2019), 1-31.

Besley, T., and Coate, S., "An Economic Model of Representative Democracy," *Quarterly Journal of Economics*, 112:1 (1997), 85-114.

Bhalotra, S., and Clots-Figueras, I., "Health and the Political Agency of Women" *American Economic Journal: Economic Policy*, 6:2 (2014), 164-197.

Bhalotra, S., Clots-Figueras, I., and Iyer, L., "Religion and Abortion: The Role of Politician Identity," Working paper (2019).

Bhalotra, S., Cassan, G., Clots-Figueras, I., and Iyer, L., "Religion, Politician Identity and Development Outcomes: Evidence from India," *Journal of Economic Behavior and Organization*, 104 (2014), 4–17.

Bhavnani, R., "Do the Effects of Temporary Ethnic Group Quotas Persist? Evidence from India," *American Economic Journal: Applied Economics*, 9:3 (2017), 105-123.

Bhavnani, R., Donnay, K., Miodownik, D., Mor, M., and Helbing, D., "Group Segregation and Urban Violence," *American Journal of Political Science*, 58 (2014), 226-245.

Bisin, A., and Verdier, T., "The Economics of Cultural Transmission and the Dynamics of Preferences," *Journal of Economic Theory*, 97 (2001), 298 – 319.

Blattman, C., and Miguel, E., "Civil War," *Journal of Economic Literature*, 48:1 (2010), 3-57.

Blouin, A., and Mukand, S., "Erasing Ethnicity? Propaganda, Nation Building, and Identity in Rwanda," *Journal of Political Economy,* 27:3 (2019), 1008-1062.

Bowles, S., and Halliday, S., *Microeconomics: Competition, Cooperation and Conflict*, (Oxford: Oxford University Press, 2020).

Brandts, J., and Cooper, D., "A Change Would Do You Good... An Experimental Study on How to Overcome Coordination Failure in Organizations," *American Economic Review*, 96:3 (2006), 669-693.

Brandts, J., and Cooper, D. J., "It's What You Say, Not What You Pay: An Experimental Study of Manager-Employee Relationships in Overcoming Coordination Failure," *Journal of the European Economic Association*, 5:6 (2007), 1223-1268.

Brandts, J., Cooper, D. J., and Fatas, E., "Leadership and Overcoming Coordination Failure with Asymmetric Costs," *Experimental Economics*, 10:3 (2007), 269-284.

Brandts, J., Cooper, D. J., and Weber, R. A., "Legitimacy, Communication, and Leadership in the Turnaround Game," *Management Science*, 61:11 (2015), 2627-2645.

Brandts, J., Rott, C., and Solà, C., "Not Just Like Starting Over - Leadership and Revivification of Cooperation in Groups," *Experimental Economics*, 19 (2016), 792–818.

Brollo, F., and Troiano, U., "What Happens When a Woman Wins an Election? Evidence from Close Races in Brazil," *Journal of Development Economics*, 122 (2016), 28-45.

Brooks, B., Hoff, K., and Pandey, P., "Cultural Impediments to Learning to Cooperate: An Experimental Study of High- and Low-caste Men in Rural India," *Proceedings of the National Academy of Sciences*, 115:45 (2018), 11385-11392.

Burgess, R., Jedwab, R., Miguel, E., Morjaria, A., and Padró i Miquel, G., "The Value of Democracy: Evidence from Road Building in Kenya," *American Economic Review*, 105:6 (2015), 1817-1851.

Carlsson, H., and van Damme, E., "Global Games and Equilibrium Selection," *Econometrica* 61 (1993), 989-1018.

Cartwright, E., Gillet, J., and Van Vugt, M., "Leadership by Example in the Weak-Link Game," *Economic Inquiry*, 51:4 (2013), 2028-2043.

Chattopadhyay, R., and Duflo, E., "Women as Policymakers: Evidence from a Randomized Policy Experiment in India," *Econometrica,* 72:5 (2004), 1409–1443.

Chen, R., and Chen, Y., "The Potential of Social Identity for Equilibrium Selection," *American Economic Review,* 101 (2011), 2562– 2589.

Chin, A., and Prakash, N., "The Redistributive Effects of Political Reservation for Minorities: Evidence from India," *Journal of Development Economics*, 96:2 (2011), 265-277.

Cohn, A., Marechal, M., and Noll, T., "Bad Boys: How Criminal Identity Salience Affects Rule Violation," *Review of Economic Studies,* 82 (2015), 1289–1308.

Coleman, J. S., "Norms as Social Capital," in G. Radnitzky and P. Bernholz, eds., *Economic Imperialism* (New York: Paragon, 1987), 133-155.

Corno, L., La Ferrara, E., and Burns, J., "Interaction, Stereotypes and Performance: Evidence from South Africa," Working paper (2018).

Finseraas, H., Hanson, T., Johnsen, A. A., Kotsadam, A., and Torsvik, G., "Trust, Ethnic Diversity, and Personal Contact: A Field Experiment," Working paper (2020).

Fisman, R., Sarkar, A., Skrastins, J., and Vig, V., "Experience of Communal Conflicts and Intergroup Lending," *Journal of Political Economy,* 128:9 (2020), 3346-3375.

Gächter, S., Nosenzo, D., Renner, E., and Sefton, M., "Who Makes a Good Leader? Cooperativeness, Optimism, and Leading-by-Example," *Economic Inquiry,* 50 (2012), 953-967.

Gangadharan, L., Jain, T., Maitra, P., and Vecci, J., "Social Identity and Governance: The Behavioral Response to Female Leaders," *European Economic Review*, 90 (2016), 302-325.

Gupta, G., Mahmud, M., Maitra, P., Mitra, S., and Neelim, A., "Religion, Minority Status and Trust: Evidence from a Field Experiment," *Journal of Economic Behavior and Organization*, 146 (2018), 180 - 205.

Güth, W., Levati, M. V., Sutter, M., and Van Der Heijden, E., "Leading by Example with and without Exclusion Power in Voluntary Contribution Experiments," *Journal of Public Economics*, 91 (2007), 1023–1042.

Heursen, L., Ranehill, E., and Weber, R., "Are Women Less Effective Leaders Than Men? Evidence from Experiments Using Coordination Games," Working paper, (2019).

Hoff, K., "Beyond Rosenstein-Rodan: The Modern Theory of Coordination Problems in Development," in *Proceedings of the Annual World Bank Conference on Development Economics* (Washington D.C.: World Bank, 2000), 145-188.

Hoff, K., and Stiglitz, J., "Modern Economic Theory and Development," in G. M. Meier and J. E. Stiglitz, eds., *Frontiers of Development Economics: The Future in Perspective* (New York: Oxford University Press, 2001).

Hungerman, D. M., "Do Religious Proscriptions Matter? Evidence from a Theory-Based Test," *Journal of Human Resources*, 49 (2014), 1053–1093.

Ip, E., Leibbrandt, A., and Vecci, J., "How Do Gender Quotas Affect Hierarchical Relationships? Complementary Evidence from a Representative Survey and Labor Market Experiments," *Management Science*, 66:2 (2020), 805–822.

Iyer, L., Mani, A., Mishra, P., and Topalova, P., "The Power of Political Voice: Women's Political Representation and Crime in India," *American Economic Journal: Applied Economics,* 4:4 (2012), 165–193.

Jack, K., and Recalde, M., "Leadership and the Voluntary Provision of Public Goods: Field Evidence from Bolivia," *Journal of Public Economics* 122 (2015), 80-93.

Kremer, M., "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics*, 108:3 (1993), 551-575.

Krupka, E., and Weber, R., "Identifying Social Norms using Coordination Games: Why Does Dictator Games Haring Vary?" *Journal of the European Economic Association*, 11:3 (2013), 495–524.

Leibbrandt, A., Wang, L.C., Foo, C., "Gender Quotas, Competitions, and Peer Review: Experimental Evidence on the Backlash Against Women," *Management Science*, 64:8 (2018), 3501–3516.

Levy, D.M., Padgitt, K., Peart, S.J., Houser, D., and Xiao, E., "Leadership, Cheap Talk and Really Cheap Talk," *Journal of Economic Behavior and Organization*, 77:1 (2011), 40-52.

Lowe, M., "Types of Contact: A Field Experiment on Collaborative and Adversarial Caste Integration," *American Economic Review*, forthcoming (2020).

Matsa, D. A., and Miller, A. R., "A Female Style in Corporate Leadership? Evidence from Quotas," *American Economic Journal: Applied Economics*, 5:3 (2013), 136-169.

Miguel, E., and Roland, G., "The Long-Run Impact of Bombing Vietnam," *Journal of Development Economics*, 96 (2011), 1-15.

Miguel, E., "Tribe or Nation? Nation Building and Public Goods in Kenya versus Tanzania," *World Politics* 56:3 (2004), 327–362.

Morris, S., Shin, H.S., and Yildiz, M., "Common Belief Foundations of Global Games," *Journal of Economic Theory,* 163 (2016), 826-848.

Niederle, M., Segal, C., and Vesterlund, L., "How Costly is Diversity? Affirmative Action in Light of Gender Differences in Competitiveness," *Management Science*, 59:1 (2013),1-16.

Nordhaus, W. D., "Paul Samuelson and Global Public Goods," in M. Szenberg, L. Ramrattan, and A. Gottesman, eds., *Samuelsonian Economics* (Oxford: Oxford University Press, 2006), 88–98.

Osborne, M., and Slivinski, A., "A Model of Political Competition with Citizen-Candidates," *Quarterly Journal of Economics*, 111:1 (1996), 65-96.

Paluck, E., Green, S., and Green, D., "The Contact Hypothesis Re-evaluated," *Behavioral Public Policy,* 3:2 (2019), 129-158.

Pande, R., "Can Mandated Political Representation Provide Disadvantaged Minorities Policy Influence? Theory and Evidence from India," *American Economic Review*, 93:4 (2003), 1132-1151.

Polania-Reyes, S., and Echeverry, D., "Coordination as Unintended Benefit: Lab-in-the-Field Evidence from a Conditional Cash Transfer Program," Working Paper, (2020).

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

Sandler, T., "Global and Regional Public Goods: A Prognosis for Collective Action," *Fiscal Studies*, 19:3 (1998), 221–247.

Rodrik, D., "Where Did All the Growth Go? External Shocks, Social Conflict, and Growth Collapses," *Journal of Economic Growth* 4 (1999), 385–412.

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

Schelling, T. C., *The Strategy of Conflict* (Cambridge: Harvard University Press, 1980).

Shih, M., Pittinsky, T. L., and Ambady, N., "Stereotype Susceptibility: Identity Salience and Shifts in Quantitative Performance," *Psychological Science*, 10:1 (1999), 80–83.

Steele, C. M., and Aronson, J., "Stereotype Threat and the Intellectual Test Performance of African Americans," *Journal of Personality and Social Psychology*, 69:5 (1995), 797–811.

Van Huyck, J. B., Battalio, R. C., and Beil, R. O., "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80 (1990), 234-248.

Weber, R., Camerer, C., Rottenstreich, Y., and Knez, M., "The Illusion of Leadership: Misattribution of Cause in Coordination Games," *Organization Science*, 12:5 (2001), 582-598.

Table 1: Leader Identity and Minimum Effort (Control Group)

Dependent variable: Minimum Effort in the Group

	(1) Muslim Leaders	(2) Hindu Leaders	(3) Muslim Leaders	(4) Hindu Leaders	(5) All Leaders	(6) All Leaders
Leader (Period>4)	1.067	-0.488	1.067	-0.488	-0.488	
	(0.494)	(0.381)	(0.508)	(0.392)	(0.379)	
Muslim Leader * (Period>4)					1.555	
					(0.620)	
Muslim Leader					-0.492	1.272
					(0.616)	(0.618)
Observations	246	246	246	246	492	164
R-squared	0.281	0.258	0.477	0.435	0.309	0.536
Town FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic Controls	No	No	Yes	Yes	Yes	Yes
Religious Controls	No	No	Yes	Yes	Yes	Yes
Experimental Controls	No	No	No	No	No	Yes

Standard errors in parentheses, clustered at group level. Demographic controls include gender, age, education and monthly household income; religious controls include dummies for whether the participant prays several times a day or once a day; experimental controls include the leader's proposal and the group minimum effort in period 4. Columns 1-5 include data from all periods; Column 6 is restricted to periods 5 and 6 only.

Table 2: Leader Identity and Individual Effort (Control Group)

	(1)	(2)	(3)	(4)
	Individual Effort			
	Hindu Leaders		Muslim Leaders	
	Muslim Employees	Hindu Employees	Muslim Employees	Hindu Employees
Leader (Period>4)	0.015	-1.156	1.157	0.377
	(0.517)	(0.408)	(0.434)	(0.531)
p-value (M employee = H employee)	[0.093]		[0.281]	
Observations	486	480	486	498
R-squared	0.242	0.280	0.309	0.278
Town FE	Yes	Yes	Yes	Yes
Demographic Controls	Yes	Yes	Yes	Yes
Religious Controls	Yes	Yes	Yes	Yes

Standard errors in parentheses, clustered at group level. Demographic controls include gender, age, education and monthly household income; religious controls include dummies for whether the participant prays several times a day or once a day. Data includes effort choices of both leaders and employees in columns (1)-(4).

Table 3: Policy Environments and Leader Effectiveness

Dependent variable: Minimum Effort in the Group

	(1) Muslim Leaders	(2) Hindu Leaders	(3) Muslim Leaders	(4) Hindu Leaders
Leader (β_1)	1.067	-0.488		
	(0.495)	(0.382)		
Leader *AA (β_2)	-0.840	2.391	-1.625	2.370
	(0.814)	(0.671)	(0.837)	(0.627)
Leader * Contact (β_3)	1.007	2.755	1.738	2.850
	(0.752)	(0.651)	(0.801)	(0.676)
p-value for $\beta_2 = \beta_3$	0.034	0.634	0.000	0.476
Observations	774	768	258	256
R-squared	0.258	0.340	0.365	0.399
Town FE	Yes	Yes	No	No
Demographic Controls	Yes	Yes	Yes	Yes
Religious Controls	Yes	Yes	Yes	Yes
Experimental Controls	No	No	Yes	Yes

Standard errors in parentheses, clustered at group level. "Leader" is a dummy that equals one for periods 5 and 6, when a leader has made a proposal in each group. Demographic controls include gender, age, education and monthly household income; religious controls include dummies for whether the participant prays several times a day or once a day; experimental controls include the leader's proposal and the group minimum effort in period 4. Columns 1 and 2 include data from all periods; Columns 3 and 4 are restricted to periods 5 and 6 only. Leader dummy is always equal to one in columns 3 and 4. β_1, β_2 and β_3 are as defined in equation 6.

Table 4: Does a History of Conflict Matter for Leader Effectiveness across Policy Environments?

Dependent variable: Minimum Effort in the Group

	(1) Muslim Leaders Low Conflict Areas	(2) Muslim Leaders High Conflict Areas	(3) Hindu Leaders Low Conflict Areas	(4) Hindu Leaders High Conflict Areas
Leader	1.181	0.978	-0.486	-0.489
	(0.577)	(0.770)	(0.555)	(0.537)
Leader *AA	0.757	-2.176	2.236	2.520
	(1.194)	(1.036)	(1.020)	(0.912)
Leader * Contact	2.014	0.320	1.663	3.470
	(1.100)	(1.027)	(0.780)	(0.939)
Observations	336	438	330	438
R-squared	0.417	0.282	0.401	0.328
Town FE	Yes	Yes	Yes	Yes
Demographic Controls	Yes	Yes	Yes	Yes
Religious Controls	Yes	Yes	Yes	Yes

Standard errors in parentheses, clustered at group level. "Leader" is a dummy that equals one for periods 5 and 6, when a leader has made a proposal in each group. Demographic controls include gender, age, education and monthly household income; religious controls include dummies for whether the participant prays several times a day or once a day.

Figure 1: Leader Identity and Minimum Group Effort in Different Policy Environments

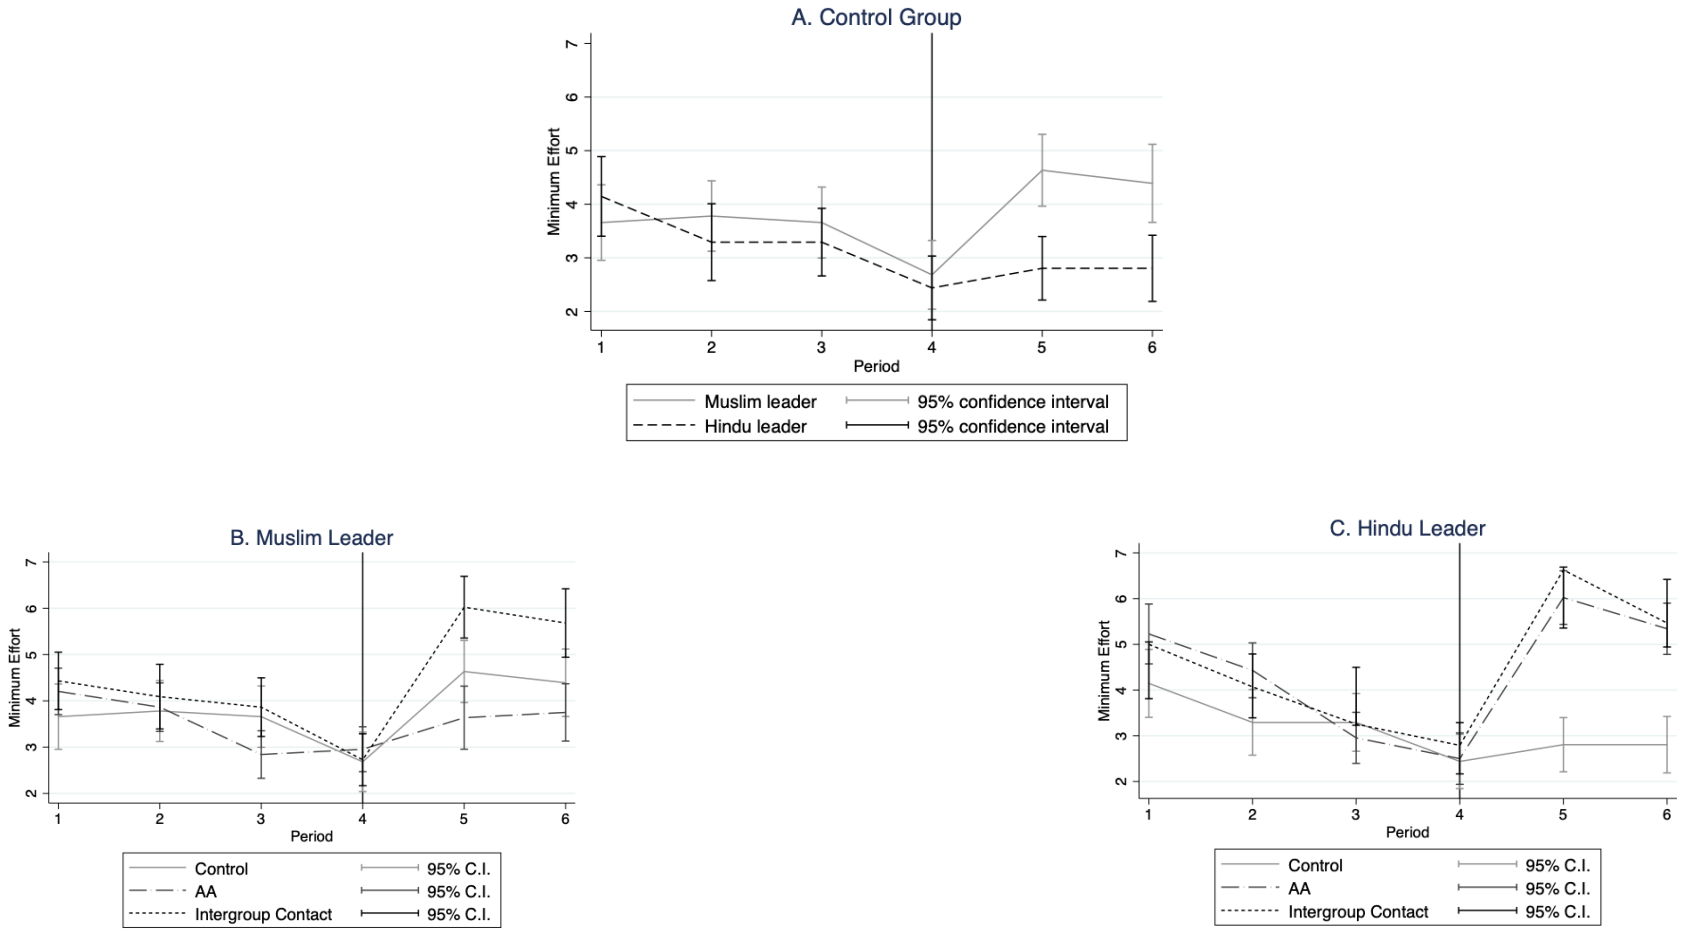

The figure shows the average minimum effort for groups with Muslim leaders and groups with Hindu leaders. Periods 1-4 are prior to leader assignment, periods 5 and 6 show outcomes after leader identity and leader proposals are revealed to participants.

Figure 2: Leader Effectiveness Across Policy Environments and Conflict Histories

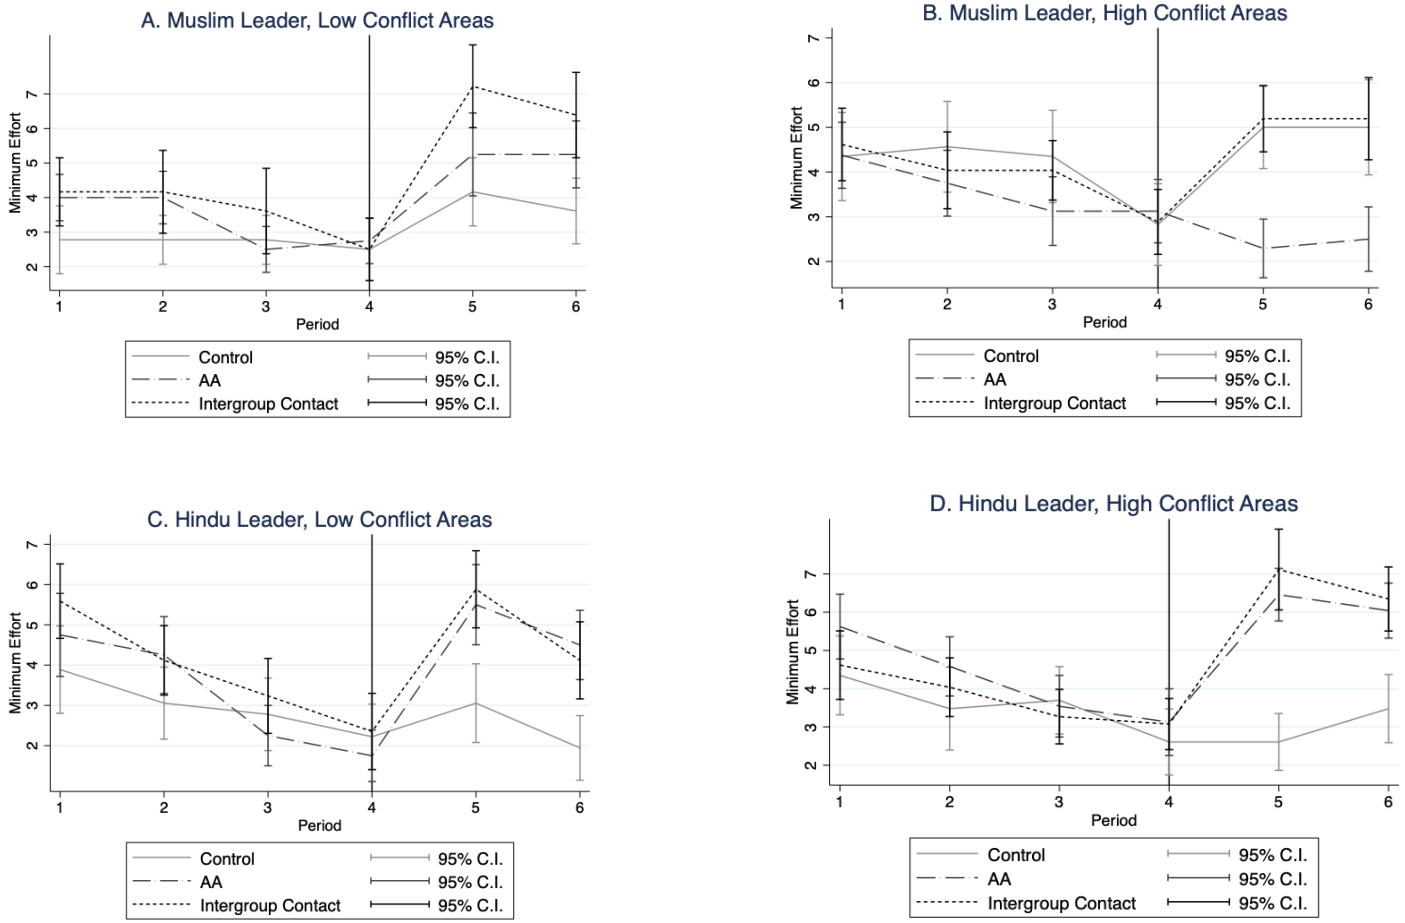

The figure shows the average minimum effort for groups with Muslim leaders and groups with Hindu leaders. Periods 1-4 are prior to leader assignment, periods 5 and 6 show outcomes after leader identity and leader proposals are revealed to participants.