

Leader Identity and Coordination*

Sonia Bhalotra
University of Essex

Irma Clots-Figueras
University of Kent

Lakshmi Iyer
University of Notre Dame

Joseph Vecci
University of Gothenburg

June 28, 2019

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 coordination irrespective of leader identity. Both policies are less effective in minority-led groups in towns with a conflict history. Our results demonstrate that leader and policy effectiveness depend upon citizen reactions, conditioned by group identity and past conflict.

JEL codes: P16, D70, D91, J78

Keywords: Leader identity, religion, coordination failure, affirmative action, intergroup contact, conflict, India

* Email addresses: srbhal@essex.ac.uk; I.Clots-Figueras@kent.ac.uk; liyer@nd.edu; joseph.vecci@economics.gu.se. 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 Jordi Brandts, Antonio Cabrales, Utteeyo Dasgupta, Saumitra Jha, Anandi Mani, Eva Ranehill, Farah Said, 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

Ethnic and religious diversity is often associated with conflict and lower economic growth. In a survey paper, Alesina and La Ferrara (2005) emphasize the importance of having “rules of the game” to channel diversity towards productive ends. Leaders could potentially play an important role in setting these rules and establishing norms. 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, we focus upon citizen reactions. While there is experimental evidence that leaders are effective in improving economic outcomes (Guth et al., 2007; Levy et al., 2011; Brandts et al., 2015; Jack and Recalde, 2015), this evidence typically pertains to a laboratory setting. It is unclear whether these results would generalize to leader effectiveness in diverse societies, where the influence of leaders may depend not only on their preferences and actions, but also on how citizens react to the leader’s social identity. We find that citizen responses to who they are led by, rather than how they are led, can significantly impact group outcomes. In a novel extension of the analysis, we test how citizens react to leaders under two randomized policy regimes that aim to improve social integration. We find that citizen reactions to leader identity also play a significant role in determining the effectiveness of such policies.

The outcome we investigate is economic coordination. Coordination can be critical to resolving collective action problems and market failures. It has been shown to be important for changing social norms, escaping poverty traps, optimizing resource use on common land, and raising the provision of public goods (Coleman, 1987; Kremer, 1993; Hoff, 2000; Hoff and Stiglitz, 2001; Bowles et al., forthcoming).¹ Coordination may be particularly

¹ Coordination is different from cooperation. Cooperation is usually measured by the willingness to contribute in a standard public goods game, in which payoffs depend on the sum of other players’ contributions and there is only one pure-strategy Nash equilibrium. Instead, in the “weakest link” coordination game where payoffs depend on the minimum of other people’s contributions, there are multiple pure-strategy Nash equilibria that can be Pareto-ranked. Conceptually, non-zero contributions in a public goods game can be related to altruism, trust or other prosocial motivations, and the leader’s role can be thought of as encouraging such motivations. In contrast, non-zero effort in the coordination game is individually rational and the leader’s role is mostly to guide the group towards a specific equilibrium i.e. to act as a focal point.

difficult to achieve in diverse societies with a history of intergroup conflict. To allow for this, our analysis is conducted in regions with both high and low levels of past conflict.

We experimentally vary the religious identity of the leader to examine their role in overcoming coordination failure. Addressing this research question with observational data would be difficult for several reasons. 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 address these challenges, we implement a large laboratory-style experiment in a field setting in India. We use a weakest link coordination game (as in Brandts and Cooper, 2006), a key feature of which is that individual payoffs depend positively upon the minimum effort in the group and negatively on their own effort. Coordination is measured as the minimum level of effort achieved in the group. Payoffs are designed such that coordination tends to occur at the lowest effort level, in a Pareto-inefficient equilibrium. Leaders can induce coordination by proposing a (non-binding) effort level. We conducted this coordination game with 1028 participants in mixed-religion groups. Our sample includes Muslims, who are a religious minority in India, and Hindus who form the majority religion.

In addition to the analysis of leader identity in the control arm, we have two treatment arms, where we introduce policy interventions that are commonly used to improve integration and opportunities for minorities: intergroup contact and affirmative action. We thus present the first analysis of affirmative action and intergroup contact in the same experimental setting. In both the treatment and the control arms, we randomly assigned some groups to have Muslim leaders and others to have Hindu leaders. Furthermore, we stratify the samples by the intensity of recent religious conflict. Our experimental design allows us to uniquely examine how leader identity influences outcomes, how such influence varies with affirmative action and intergroup contact, and how a history of intergroup conflict shapes these impacts.

We report four main findings. First, 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 is a novel finding, in terms of identifying a potential gain in aggregate output associated with having minority leaders. Investigating mechanisms, we find that this does not arise from Muslim leaders proposing higher effort (a proxy for leader preferences) but, rather, from citizen reactions to leader identity.

We further investigated potential mechanisms for the observed citizen reactions using experimental data on individual effort, and survey data on attitudes and beliefs. Examining the data on individual effort levels, we find that the results emerge from stronger in-group bias among the minority group (Muslims); this effect is absent in towns where Muslims form a population majority. This is consistent with previous literature that documents greater in-group bias for minority groups (Bisin and Verdier, 2001; Gupta et al., 2018). We find no evidence in support of alternative explanations such as a higher perceived competence for, or prior exposure to, Muslim leaders.

The main findings are reversed in the sites that were randomly assigned to receive the affirmative action (AA) treatment. In particular, when Muslim leaders are perceived to be the beneficiaries of AA, there is a deterioration of coordination in Muslim-led groups, alongside increased coordination within Hindu-led groups. The primary mechanism appears to be the priming of group identity among the majority community (Hindus). Data on individual effort levels shows Hindu participants increasing their effort in Hindu-led groups under the AA policy, a citizen response that we do not see in the control group. Gangadharan et al. (2016) similarly find that gender quotas result in the priming of male identity.

Our third result is that intergroup contact (achieved by having mixed-identity groups collaborate in solving a puzzle before the coordination game is played) improves coordination in groups led by leaders of both religions. In a similar vein, Dahl et al. (2018) finds that intergroup contact of men and women improves gender attitudes. In our study,

the gain is larger for Hindu-led groups. This is driven by increased effort of Hindu individuals in Hindu-led groups, similar to the increase observed in the AA treatment arm, suggesting that intergroup contact may also lead to the priming of group identity among the majority group. Bagues et al. (2017) similarly find that intergroup contact of men and women primes male identity.

Our final set of results emerge from re-estimating the impact of leader identity in the control group and in the two policy arms of the experiment, distinguishing towns in neighboring districts with a high vs low intensity of past Hindu-Muslim conflict. We compare coordination under different treatments *within* each district, so that district-specific unobservables do not contaminate our findings. This novel design allows us to investigate how the efficacy of our policy interventions varies with a baseline measure of intergroup hostility. We find that a history of conflict has a detrimental effect on the effectiveness of Muslim leaders under both policy regimes: intergroup contact increases Muslim leader effectiveness only in low conflict areas, and AA decreases Muslim leader effectiveness primarily in high conflict areas. In contrast, conflict history does not significantly influence leader effectiveness in Hindu-led groups. Again, these results are consistent with Hindu group identity becoming more salient in areas with a high conflict history, thereby lowering the effectiveness of Muslim leaders.

This paper contributes to the literature on leader identity in several ways. First, in a departure from previous work that typically identifies the combined impact of leader preferences (or actions) and citizen reactions, our experimental design allows us to identify citizen reactions to leader identity.² It turns out that they are a significant determinant of leader effectiveness. Second, previous work on leader identity has tended to focus on whether leaders redistribute public goods or transfers towards their own group (Pande, 2003; Burgess et al., 2015), while our focus on coordination provides us a measure of the

² The political economy literature has mainly focused on the identity of political leaders. 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. Empirical analysis of the role of leader identity in political settings include Pande, 2003; Chattopadhyay and Duflo, 2004; Chin and Prakash, 2011; Bhalotra & Clots-Figueras, 2014; Meyerson, 2014; Burgess et al., 2015; Brollo and Troiano, 2016; and Bhavnani, 2017.

aggregate economic impact of leader identity. Third, our study of potential mechanisms reveals that in-group bias among the majority group can be activated by commonly implemented policies such as AA or intergroup contact, with important implications for the success of these policies.³ No previous work has investigated effectiveness of these policies as a function of leader identity. Fourth, our design allows us to compare impacts of AA and intergroup contact in the same experimental setting, which previous work has tended to assess independently. 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 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, which shows that introducing a group leader can facilitate the transition to a more efficient outcome (Brandts et al., 2015).⁵ No previous study has examined whether leader identity matters for leader effectiveness in improving coordination, or how the effectiveness of leaders may be modified by policies designed to assimilate minorities or promote minority group representation. Also, previous leader-coordination studies have almost entirely been conducted in the laboratory,⁶ with a few exceptions (Brooks et al., 2015; Polania-Reyes, 2016). Our implementation of the coordination game in the field enables us to use real social identities rather than lab-assigned identities to investigate how past conflict influences group behavior, and to conduct the analysis in a developing country setting, where coordination failures may be more common given the absence of strong formal institutions.

³ We do not find such in-group bias among the majority community in the control group, similar to Berge et al. (forthcoming) who also find no evidence of co-ethnic bias in behavioral games conducted in Kenya.

⁴ Previous work on post conflict policies is often either correlational, or conducted at the country or region level. The empirical and experimental literature on conflict primarily focuses on the origins of conflict (see Blattman and Miguel, 2010, for a review) and its impact on growth (Rodrick, 1999), human capital (Miguel and Roland, 2011) and cooperation (see Bauer et al., 2016 for a review).

⁵ See Hogg, 2001 for a theoretical analysis and Devetag and Ortmann, 2007 for a survey of this literature.

⁶ See, among others, Brandts and Cooper, 2006; Brandts and Cooper, 2007; Brandts et al., 2007; Cartwright et al., 2013; Brandts et al., 2015; Brandts et al., 2016; Heursen et al. (2019).

We contribute to the literature on AA policies, where there is no previous evidence of how citizens react to AA under leaders of different religious or ethnic identity; the existing literature has been dominated by the analysis of gender quotas.⁷ Similarly, no previous paper has analyzed the effectiveness of intergroup contact in improving coordination, nor how this varies with leader identity (and conflict history). The previous literature on intergroup contact finds that such contact can change attitudes and prejudice towards the out-group (Paluck et al., 2018).⁸

Overall, our findings are relevant to research in several areas that are currently distinct from one another, including coordination, leader identity, quotas, segregation, in-group bias and conflict. The rest of the paper is structured as follows. Section 2 provides contextual information on religion and leadership in India. Section 3 describes the data collection, Section 4 delineates the experimental design and Section 5 discusses the regression specifications. Section 6 presents the empirical results, and Section 7 concludes.

2. Context: Religion and Leadership in India

India is a religiously diverse country, with Hindus constituting the majority religion (79.8% of the population in the 2011 census), Muslims forming the largest religious minority (14.2% of the population) and the rest being comprised of Christians, Sikhs, Jains, Buddhists and others. With 172 million Muslims in 2011, India has the third largest Muslim population in the world. India is a secular country, with the constitution enshrining the fundamental right to freely “profess, practice and propagate religion.” Close to 40% of Muslims live in urban areas relative to 29% of Hindus. The standard of living of Muslims

⁷ Many papers have examined how gender quotas influence policy outcomes, gender norms, women’s aspirations and political participation (see, among others, 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 almost exclusively focuses on whether gender quotas encourage women to take part in tournaments (Schotter and Weigelt, 1992; Balafoutas and Sutter, 2012; Niederle et al., 2013; Leibbrandt et al., forthcoming).

⁸ See also the literature surveys in Pettigrew and Tropp (2006), Ashraf and Bandiera (2017) and Bertrand and Duflo (2017). Only a few recent studies use random assignment of groups to examine intergroup contact in a developing country setting (Corno et al., 2018; Lowe, 2017; Rao, 2018; Scacco and Warren, 2018). Bhavnani et al. (2014) study the effects of intergroup contact on violence in Jerusalem but their data are observational rather than experimental. Barnhardt (2009) finds that contact between Hindus and Muslims in India reduces anti-Muslim prejudice.

is generally lower than that of Hindus: 43% of Muslims are illiterate relative to 36% for Hindus, 33% of Muslims are employed relative to 41% for Hindus and 25% of Muslims fall below the poverty line compared to 22% for all Indian citizens (NITI Aayog, 2016).

Violence between Hindus and Muslims has occurred frequently in India generating insecurity, displacement, segregation and loss of property and life. Previous research indicates that violence tends to disproportionately impact Muslims (Mitra and Ray, 2014). For instance, between 1985 and 1987, in the ten states with significant Muslim populations, Muslims experienced 60% of all deaths related to religious riots, 45% of all injuries and 73% of property damage (Wilkinson, 2004, p 30). Previous research has shown that Muslims (and Hindus) exhibit positive in-group bias in trusting behavior when they form a minority of the population, and that such in-group bias is increased by the degree of identification with the group (Gupta et al., 2018). In Indian elections (where Muslims form a minority of voters), Muslims are more likely to vote for Muslim candidates, but there is no evidence that Hindus discriminate against Muslim candidates (Heath et al., 2015).

Compared to their population share, Muslims are under-represented in leadership positions: Muslims comprised only 7.6% of state legislators over the period 1980-2010 (Bhalotra et al., 2014), less than 3% of national ministers (The Economist, 2016), and only 4% of the police force and the judiciary (Jaffrelot et al., 2019). This low representation has substantive consequences, since the religious identity of politicians has been shown to matter for policy outcomes such as health, education and sex-selective abortions (Bhalotra et al., 2014; Bhalotra et al., 2019). There are no quotas for Muslims in central government positions. However, some quotas for Muslims have been implemented in civil service and educational institutions in five states (Government of India, 2014).

3. Site Selection and Subject Recruitment

3.1. Site Selection and Balance Across Treatment Arms

Based on data on the incidence of religious violence over 1980-2010 (Kaysser et al., 2015), we selected two pairs of neighboring districts in Uttar Pradesh, India's most populous state. Each district pair was composed of one high conflict district and one low

conflict district (see map in Appendix A, Figure A1). One pair was in the western part of the state (Aligarh and Budaun) while the other was in the central part of the state (Allahabad and Pratapgarh).

The experiment was conducted in 44 sites in July 2017 with 1028 subjects. Within each district, we selected sites based on the following criteria: (i) they were officially listed as towns in the census of 2011,⁹ (ii) their population was below 50,000, (iii) they had a relatively high population share of Muslims, and (iv) selected towns were comparable in their demographics across high conflict and low conflict districts. The average town in our experimental sample had 40% Muslims in the population, compared to the overall state proportion of 19%.¹⁰ We focus on towns rather than villages, as Muslims are more likely to live in urban areas. This has the added benefit of comprising a more educated populace, reducing issues surrounding the understanding of the experimental tasks.

The 44 town level sessions consisted of 24 participants each,¹¹ resulting in a sample of 1028 subjects. The 44 sites were randomly assigned to three different treatment arms: 14 sites are retained as control, and 15 each are assigned to the intergroup contact treatment and to the affirmative action treatment (see Table A1 in Appendix A). The assignment was performed within each district (i.e. implicitly stratifying by high/low conflict status) and by further stratifying by Muslim population proportion and total population. We verify that census 2011 and respondent demographic characteristics are balanced across the different treatment arms (see Table 1, Panels A and B). In particular, we find that only 8 out of 56 comparisons in this table show statistically significant differences, a small proportion relative to the overall number of tests. Nevertheless, we control for all these characteristics

⁹ Towns in India are defined as (a) statutory towns i.e. all places with a municipality, corporation, cantonment board or notified town area committee or (b) all other places which satisfied the following criteria: (i) a minimum population of 5,000; (ii) at least 75 per cent of the male main working population engaged in non-agricultural pursuits; and (iii) a density of population of at least 400 persons per sq. km.

¹⁰ According to the census of 2011, Hindus and Muslims account for more than 99% of the population of Uttar Pradesh.

¹¹ There were a few exceptions due to recruitment difficulties: sessions consisted of 20 participants in three towns, and 16 participants in two towns. Our results control for town fixed effects and are robust to the exclusion of the towns with recruitment difficulties.

in our regressions to ensure that these small differences in participant characteristics do not bias our results; results without these controls are remarkably similar to those with controls.

3.2. Subject Recruitment

To recruit participants for the experiment, our research assistants visited both the Hindu and the Muslim sections of each town, and distributed flyers containing information about participant requirements (age and numeracy), remuneration, time and location of the experimental sessions. Flyers were also posted at prominent landmarks and distributed at shops, temples and mosques. Upon arrival, participants were screened for eligibility (over the age of 18 and could read numbers) and their names were recorded on a participant list that was not visible to other participants. Upon commencement of the experiment, an equal number of Hindus and Muslims were selected from the list. Once seated, participants were given a number tag representing their randomly assigned identification number.¹² Individuals were assembled in a room and randomly seated on mats, with four participants on each mat (see Appendix A, Figure A2). Participants were given large plastic folders to ensure their choices were not visible to others.

4. 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 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. We explain each component below.

¹² In India, it is relatively easy to identify Muslims from their names. After recording participant names, research assistants would implicitly allocate names to religion. Religion was then cross-checked once subjects completed the pre-experiment questionnaire. Subject ID numbers were never matched with participant names, and the list of participant names were destroyed at the end of the session. At no time were subjects informed of the names of their partners.

4.1. Pre-Experiment Survey and Puzzle Task

Prior to commencement of the incentivized tasks, subjects answered a brief survey about their personal characteristics that was designed to elicit the religion of the participant (see Section 4.2 for further details and Appendix B for complete experiment instructions). 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., 2018; Lowe, 2017). The time given for the puzzle assembly was twelve minutes and almost all participants were able to complete the task successfully in this time. If this task was chosen for payment, subjects were paid ₹ 400 for assembling the puzzle.

4.2. 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. In what follows we will explain the components of this task in more detail.

4.2.1. 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; they were not informed of the religion (or any other characteristics) of their firm members. They were told that the people sitting on their mat were *not* part of their firm, so that they did not attempt to observe the effort choices of their physical neighbors.

4.2.2. Effort Choices and Payoffs

The task is run across six periods. In each period employees decide how many hours (H) to devote to firm activities. Their choices vary between 0 and 20 in intervals of 5: $H_i \in \{0, 5, 10, 15, 20\}$. Employees' payoffs for each period are determined by equation (1) below, where i indexes individuals and j indexes groups. The payoff table is illustrated in Table A2 of Appendix A, where the units are Indian rupees (1USD \approx ₹68). Participants were shown the payoff table but not the payoff equation. Working is costly, each hour worked costing an employee ₹25. Thus, payoffs are decreasing in the employee's own hours of effort, but they are increasing in the minimum hours worked in the firm. Coordinating on any of the available effort levels is a Nash equilibrium.

$$(1) \pi_{ij} = 500 - 25H_{ij} + [\min_i (H_{ij}) * 40]$$

Given the payoff equation, it is only worthwhile for profit maximizing employees to raise their effort level if this will increase the minimum effort of the firm. For a profit maximizing employee to increase their effort by 1 unit, for example, from 0 to 5 hours, they must believe there is an 85.5% probability that each of the other three employees increase their effort.¹³ Given this, previous work has found that play often evolves towards the payoff-dominated equilibrium in which all players choose the lowest possible effort level (Brandts et al., 2006). Note that such weakest link production functions can explain many real world observations, such as workers being paid more in firms that use high value inputs (Kremer, 1993).

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). The absence of feedback about

¹³ To derive this probability, solve for p where $500 = 375(1 - p^3) + 575(p^3)$

individual effort levels makes it more difficult to escape coordination failure in the turnaround game (Brandts and Cooper, 2006).

4.2.3. 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). 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 (see questionnaire layout in Appendix B). The only characteristic that varies across leaders is their religion. However, to avoid making the research question salient, which might induce socially desirable response bias or experimenter demand effects, subjects are given the leader's characteristics (height, eye color, hair color, religion), and not just the leader's religion. Using a survey is a commonly used method to make identity salient (see Shih et al. 1999; Benjamin et al, 2010; Cohn et al., 2015).

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 (Table A3 of Appendix A). Leader characteristics other than religion, in particular, gender, age, and family income are balanced across Hindu and Muslim leaders, though Muslim leaders are less likely to have gone to college and more likely to pray several times a day, similar to the variation observed in 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.

4.2.4. 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. To measure the impact of quotas on behavioral reactions to leaders of different religions, we randomize a third of the groups within each district into an AA treatment. 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 then informed that their leader is in a reserved (unreserved) position. 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."¹⁴ It is important to reiterate that

¹⁴ Our estimates, which are intent-to-treat estimates, will therefore under-estimate the impact of religion-based reservation. Sample size considerations precluded the inclusion of a treatment with Hindu leaders being reserved and, in practice, quotas are usually for population minority groups.

here, as in the control arm, all groups contain two Hindu and two Muslim participants, and half of all groups within a session are assigned to Muslim leaders. 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.

4.2.5. 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. 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 Figure A3 in Appendix A). Subjects are encouraged to talk with their partner during the 12 minutes allowed for the puzzle. By comparing outcomes across the control arm and the contact treatment arm, we can infer the impact of intergroup contact on the ability of leaders to coordinate groups.

Importantly, 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.

4.3. Survey

After the coordination game, participants participated in a social norms elicitation task, which could also be selected for payment. We do not analyze this task in this paper. After the completion of the task, participants answered a short questionnaire measuring

demographics, beliefs about behavior in the experimental task, and views on politics and inter-religious violence.

4.4. *Experimental Procedures*

Effort choices in the coordination task were made individually using pen and paper and within a large binder (to hide choices). Participants were also instructed and monitored not to look at other players' choices. Each session was run with six research assistants including one experimenter. At the start of a task, instructions were read aloud by the experimenter to establish common knowledge. To determine whether subjects understood the instructions, each participant answered a set of control questions in private both before the first period and at the start of the fifth period (when the leader is introduced). Control questions mainly focused on understanding the payoff table. The experimenter and research assistants cross-checked the answers and started the experiment once satisfied that subjects understood the task.

To ensure the anonymity of the leader, in all sessions, the leader's proposed effort was elicited after the control questions were answered, but before collecting the control question answer sheets i.e. it was not possible to identify a leader by seeing who was writing a suggested effort, since everyone was answering control questions. At the end of the session, each subject was privately paid in cash for one of the tasks (excluding the pre-experiment survey). The average payoff was ₹ 610 including a ₹ 200 show up fee. This constitutes about 2.5 days' wage for a semi-skilled laborer.

5. **Empirical Strategy**

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:

$$(2) \text{MinGroupEffort}_{kjt} = \alpha + \beta \text{Leader}_{kjt} + X'_{kj} \gamma + \varepsilon_{kjt} ; t=1,2,\dots,6$$

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. X_{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. We run specification (2) separately for Hindu and Muslim leaders to examine whether leader identity matters for coordination outcomes. Recall that half of all groups within each town are randomly assigned to have Muslim leaders.

We then examine whether leader effectiveness varies across different policy environments by comparing coordination outcomes for the same leader identity (Hindu or Muslim) across the different treatment groups as follows:

$$(3) \text{MinGroupEffort}_{kjt} = \alpha + \beta_1 Leader_{kjt} + \beta_2 Leader_{kjt} * AA_j + \beta_3 Leader_{kjt} * Contact_j + X'_{kj} \gamma + \varepsilon_{kjt} ; t = 1, 2, \dots, 6$$

In equation (3), 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 (3) 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.

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 of the game. In this robustness check, since AA and contact are treatments at the town level, we cannot include town fixed effects.

6. Leader Identity and Coordination Outcomes

6.1. Leader Identity and Coordination in the Control Group

Consistent with the corporate turnaround game literature, we find that groups tend to 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 1).

A priori, it is not clear what to expect when we introduce Hindu and Muslim leaders. Previous research has identified that leaders improve coordination (see Brandts et al., 2015), however, this literature does not consider variation in leader identity. In heterogeneous societies, citizens may prefer their own social group over outside group members even when exercising their preferences imposes a cost to themselves or society (Bisin and Verdier, 2001). Marrying the literatures on leadership and identity, it is unclear ex-ante whether introducing a leader will improve coordination or reduce coordination because of strong in-group preferences.

We find that introducing Muslim leaders significantly improves minimum group effort in periods 5 and 6. In contrast, introducing Hindu leaders does not (Figure 1). 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 2, column 1). In contrast, the presence of Hindu leaders leads to a decline of 0.488 in minimum group effort, which is not statistically significantly different from zero (Table 2, column 2). The effect of Muslim leadership on minimum effort is lower than in the lab experiment of Brandts et al. (2015) who find that a randomly selected leader increases minimum effort (in homogeneous groups) from 3.33 to 11.25.¹⁵

The estimates are robust to controlling for the demographic and religious characteristics of group members, consistent with our randomized assignment of leader identity (columns

¹⁵ For comparison purposes the minimum effort in Brandts et al. (2015) is rescaled. The lower impact of leaders in our context 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 while those in Brandts et al. (2015) had the ability to send more detailed messages to participants. It could also be because homogeneous groups coordinate better.

3 and 4). The difference in coordination gains between Muslim and Hindu leaders is statistically significant (column 5). This holds even when we restrict the sample to periods 5 and 6, and control for leader proposals and for minimum effort in period 4 (Table 2, column 6). The results are also 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 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 (results available upon request).

6.2. Potential Mechanisms: Citizen Reactions to Leader Identity

In this section we investigate the mechanisms that may explain why Muslim leaders induce greater coordination in mixed groups than Hindu leaders. As discussed in the Introduction, leaders can influence outcomes through actions that reflect their personal preferences, or through citizen reactions to them. In general, leader actions or preferences are difficult to comprehensively measure, but in our laboratory-in-the-field-experiment, leader proposals are a measure of their preferences. Our data reveal that Muslim leaders propose 10.5 hours on average, compared to 9.4 hours for Hindu leaders (Table A4). This difference is not statistically significant (Table A5, column 1), and a Kolmogorov-Smirnov test shows that the distributions of proposals by leader identity are not statistically different either (p value 0.452; Figure A4). In any case, as shown above, the difference in minimum group effort between groups with Hindu vs Muslim leaders is robust to controlling for leader proposals (Table 2, column 6). We can therefore reject leader preferences as an explanation for the greater effectiveness of Muslim leaders, in favor of citizen reactions to leader identity.¹⁶

We investigate whether citizen reactions to leader identity are driven by in-group bias, with participants more likely to exert higher levels of effort when the leader is from their religion. To investigate this, we examine *individual* effort decisions, hypothesizing that

¹⁶ We also verify that participants' expressed views on trusting people of a different religion (measured in the post-game survey) do not vary by the identity of the leader their group was assigned i.e. experimenter demand effects are not driving our results.

Muslims will supply more effort in Muslim-led groups, and possibly less effort in Hindu-led groups. We expect smaller differences by leader identity in the individual effort choices of Hindus if, as the majority group, they have weaker in-group bias. We estimate the following specification:

$$(4) \text{ IndividualEffort}_{ikjt} = a + b\text{Leader}_{kjt} + X_{ikj}'g + w_{ikjt}; t = 1, 2, \dots, 6$$

where $\text{IndividualEffort}_{ikjt}$ is the effort choice of individual i in group k (of town j) and period t . As before, X_{ikj} includes town fixed effects, demographic controls and religious controls, and standard errors are clustered at the group level. Individual effort choices depend not only on their own reactions to leader identity, but also on their expectations of how other individuals in the group will react to the leader's proposal and the leader's identity. Thus, equation (4) is not strictly comparable to equation (2) where the outcome is group minimum effort.

We report specification (4) in Table 3 for each of the four possible pairs: Muslim vs Hindu employees supplying effort under Muslim vs Hindu leaders. We find that Muslim employees exhibit significantly higher effort in groups with a Muslim leader (compared to the earlier periods without any leader), but do not change their effort choices when confronted with a Hindu leader (Table 3, columns 1 and 2). This is consistent with Muslims exhibiting in-group (or co-ethnic) bias towards leaders who share their social identity, rather than a negative reaction to leaders of a different religion. Hindu employees similarly do not change their effort choices when faced with a Muslim leader (column 3). However, in contrast to Muslims, Hindu employees reduce their effort choices in groups with a leader from their own religion (column 4), a result that is reversed in the policy treatment arms and emerges from low conflict areas (we discuss this further in subsequent sections). Statistical tests confirm that the differences in individual effort are significantly different across the two types of leaders, for both Muslim and Hindu employees (p-values 0.088 and 0.001 respectively). The results remain similar when we exclude the leader's own effort

choices from this analysis and focus purely on employees' effort choices, and when we control for beliefs about other workers' hours (available upon request).

Previous research shows that in-group bias is more likely to be exhibited by population minority groups (Bisin and Verdier, 2001; Gupta et al, 2018). While Muslims are a minority group in the country as a whole, their population proportion varies dramatically from 12% to 89% across the towns in our sample. In fact, Muslims are a population majority in 16 out of 44 towns. We leverage this variation to examine whether the results are different in Muslim-majority versus Muslim-minority towns. Consistent with in-group bias being stronger in minority groups, we find that minimum group effort under Muslim leaders increases by a statistically significant 1.739 hours in towns where Muslim are in a minority, but by a statistically insignificant 0.208 hours ($=1.739-1.531$) in towns where Muslims are in a majority (Table 3, column 5). Although the difference is not statistically significant, the magnitude of the difference is compelling. There is no similarly large difference for Hindu-led groups (column 6).

Patterns of in-group bias satisfy another prior, that of being higher in areas with a history of inter-religious conflict. We see that Muslim employees exert somewhat greater effort in Muslim-led groups in high-conflict areas (Appendix A, Table A6, columns 1 and 5). Similarly, the significantly lower effort of Hindu employees under Hindu leaders occurs primarily in low conflict areas; in high conflict areas, Hindu employees do not reduce their effort levels significantly under a Hindu leader (Appendix A, Table A6, columns 4 and 8).

In the rest of this section, we analyze and reject three other potential explanations for higher minimum effort in Muslim-led groups. First, it is possible that citizens perceive Muslim leaders as more competent than Hindu leaders because they have achieved a leadership position despite being from a disadvantaged minority group. Our post-experiment survey asked participants "Do you think Muslim leaders are less capable relative to Hindu leaders?" We find that 27% of Hindus and 10% of Muslims agree with this statement, revealing that Muslims have more faith in Muslim leaders than Hindus. However, the *response* to Muslim leaders (minimum group effort) does not differ

significantly with the fraction of the group that believes Muslim leaders are less competent (Appendix A, Table A7, columns 1 and 2).

We next examine whether actual (field) exposure to Muslim leaders may have changed citizen views of their competence, as has been demonstrated for exposure to women leaders (Beaman et al., 2009). We exploited variation in the religious identity of mayors across the towns in our sample: 17 of the 44 towns elected Muslim mayors in the most recent urban local council elections of 2012. Contrary to the hypothesis that exposure to Muslim leaders generates positive responses, we find a lower minimum effort response to Muslim leaders in towns with Muslim mayors (Appendix A, Table A7, columns 3 and 4).

Third, we examine the beliefs that Hindu and Muslim citizens hold about their assigned leaders. In our post-experiment survey we asked participants how many hours they thought their leader actually worked on average. Hindu employees estimated Muslim leaders to work slightly more hours than Hindu leaders (10.5 vs 9.7), while Muslim employees estimated Hindu leaders to have worked more hours than Muslim leaders (10.5 vs 9.7).¹⁷ However, neither of these differences is statistically significant, and our main result that minimum effort is higher in Muslim-led groups is robust to controlling for participant beliefs about the hours worked by their leader (Appendix A, Table A7, columns 5 and 6). Notice that this specification also controls for the possibility that revealing the leader's religious identity may have changed employees' beliefs about the religious composition of the group even though nothing was explicitly stated about this, since the leader is one of the four in the group whose effort counts towards minimum effort. Our results are similarly robust to controlling for participants' estimates of the hours worked by fellow employees in their group (Appendix A, Table A7, columns 7 and 8).

To summarize, our investigations indicate that the ability of Muslim leaders to induce greater coordination towards Pareto-superior equilibria is likely driven by in-group preferences being enhanced when faced with a leader from the same religion, with such activation being stronger when the individual is part of a minority group in the local area.

¹⁷ Interestingly, 79% of Muslim leaders in the control group make effort choices that equal or exceed their own proposals, compared to 65% of Hindu leaders, a difference that is statistically significant (p-value=0.004).

We confirm the latter by exploiting the fact that, in our sample, Muslims are a majority in some towns: we only observe higher coordination in towns where they are a minority (Table 3).

6.3. Citizen Reactions to Leader Identity under Affirmative Action

In this subsection we analyze the impact of leader identity under affirmative action. The existing literature has found conflicting results on the impact of AA. Ip et al. (2018) find that an AA policy improves the outcomes of the recipient group if group members believe that quota recipients are indeed disadvantaged, but can lead to backlash if people believe that quota recipients are not suitable for leadership roles (because they are less skilled or not truly disadvantaged). Gangadharan et al. (2016) find that AA may prime the social identity of non-recipient groups, increasing their in-group solidarity. We know of no previous investigation of the impact of AA on coordination outcomes.

We examine the impact of the AA treatment on minimum group effort by estimating equation (3). We find that the AA treatment *decreases* the effectiveness of Muslim leaders relative to the control arm (Table 4, column 1).¹⁸ In sharp contrast, under Hindu leaders, we find a large and statistically significant improvement in minimum group effort in the AA treatment (column 2).¹⁹ These results are consistent with the hypothesis of increased backlash towards AA recipients and greater in-group solidarity among non-recipients.

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 A, Table A5, 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 4, columns 3 and 4). The concern is allayed,

¹⁸ Minimum group effort increases by an insignificant 0.227 hours under Muslim leaders in the AA group (1.067-0.840), much smaller than the 1.067 hour improvement obtained in the control arm.

¹⁹ Minimum group effort in Hindu-led groups improves by 2.391 hours compared to a decline of 0.488 hours in the control group.

as the coefficients on Leader*AA continue to be statistically significant and are now larger in magnitude.

Similar to the way we analyzed in-group bias for the control group, we investigate it in the AA arm using the following specification (similar to equation (4)) to model individual effort of Hindu and Muslim employees as a function of the religion of their group leader:

$$(5) \text{ IndividualEffort}_{ikjt} = f_0 + f_1 \text{Leader}_{kjt} + f_2 \text{Leader}_{kjt} * \text{AA}_j + f_3 \text{Leader}_{kjt} * \text{Contact}_j + X_{ikj}'g + z_{kjt}$$

Our estimates of the coefficient f_2 suggest that the main driver of the different results in the AA treatment arm relative to the control arm is the effort choice of Hindus in Hindu-led groups. In particular, Hindu employees chose significantly higher effort levels in Hindu-led groups under the AA treatment (Appendix A, Table A8, column 4). Muslim employees make similar effort choices in the two arms, as do Hindu employees in Muslim-led groups (Appendix A, Table A8, columns 1-3). Overall, our results suggest that the AA policy leads to intensified in-group bias among non-recipients (Hindus). We have verified that this “backlash” is not due to perceptions of Muslim leader competence: the minimum group and individual effort responses to Muslim or Hindu leaders under AA are not significantly different for groups in which more individuals assess Muslim leaders as less capable than Hindus (p-value = 0.567 for Muslim leaders and 0.645 for Hindu leaders).

6.4. Citizen Reactions to Leader Identity under Intergroup Contact

Previous research on intergroup contact suggests that it can reduce prejudice or increase “mutual understanding and regard” (Allport, 1954; Pettigrew and Tropp, 2006; Dahl et al., 2018; Paluck et al., 2018; Lett, 1945, p. 35). However, contact may instead activate in-group bias (Bagues et al., 2017), leading to greater suspicion, fear, resentment, disturbance, and open conflict (Baker, 1934, p. 120). Thus, given that activation of in-group bias appears to be a mechanism driving effectiveness of Muslim vs Hindu leaders, the question of how intergroup contact influences leader effectiveness is an open question.

We find that intergroup contact tends to improve minimum group effort under both Muslim and Hindu leaders compared to the control group, but we find the difference is larger for Hindu leaders (Table 4, columns 1 and 2).²⁰ In fact, this difference counterbalances the better performance of Muslim leaders in the control group such that minimum group effort in the contact treatment is almost the same across Muslim and Hindu leaders (see Figure 2).

As with the AA treatment, we find no significant differences in leader proposals under intergroup contact compared to the control group (Appendix A, Table A5, column 2). The coefficients of interest are larger in magnitude and statistically significant after controlling for leader proposals (Table 4, columns 3 and 4). In particular, once we control for leader proposals, Muslim-led groups also show a statistically significant improvement in coordination under intergroup contact.²¹

The mechanism driving the better performance of Hindu leaders under intergroup contact is similar to that in the AA treatment – namely that the policy activates in-group preferences among (majority group) Hindu employees towards Hindu leaders. Focusing on coefficients f_3 from estimating equation (5), we see that Hindu employees choose significantly greater effort under Hindu leaders in the intergroup contact arm (Appendix A, Table A8, column 4), but that intergroup contact does not significantly change individual effort choices in the other leader-employee combinations (columns 1-3).

Our results for both AA and intergroup contact suggest that commonly used policies can dramatically change citizen reactions to leader identity. In particular, in-group bias may be activated among the majority group, and undo the intended effects of the policy. For instance, our results suggest that intensified in-group bias among Hindu employees makes it difficult for Muslim leaders to achieve coordination under an AA policy.

²⁰ 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 4, 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).

²¹ The gender of an individual's puzzle partner may also be important. We re-estimate equation (4) for the contact group participants, interacting leader with female puzzle partner. We find that participants with female puzzle partners respond less to the introduction of a leader, but the difference is not statistically significant.

The results discussed so far average across areas with different histories of Hindu-Muslim conflict. In the next section, we examine whether past conflict influences in-group behavior and, thereby, citizen responses to leader identity.

6.5. Does Conflict History Matter?

We examine whether coordination outcomes under the two types of leaders in the treatment arms vary with whether the district had high or low exposure to conflict over the period 1980-2010. This is pertinent since both policies have the potential to correct for historical grievances. It is important to note that 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.²² 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 prior conflict, but rather the effectiveness of Muslim vs Hindu leaders in each of two groups of areas with different conflict histories.

What differences might we expect based on conflict history? If areas with a history of intergroup conflict are also areas where in-group favoritism and out-group discrimination are higher, then it is likely that the majority group will react more adversely to an AA policy. The impact of conflict on the effectiveness of intergroup contact is more ambiguous. On the one hand, conflict-affected areas may be precisely those where intergroup contact is most limited to begin with, and so the contact treatment may have a larger effect. On the other hand, a history of conflict may reduce the receptiveness of individuals of both groups to contact.

We find some significant differences in the effectiveness of the two interventions by the history of conflict in the district. Our earlier findings that Muslim leaders improve

²² Most of the characteristics documented in Table 2 are similar across towns in high conflict and low conflict districts, with the exception that female labor force participation and wage work are more common in the high conflict towns, possibly reflecting the fact that inter-religious violence in India is concentrated in more urbanized areas (results available on request).

minimum group effort in the control and contact treatments emerges entirely from low conflict areas (Table 5, column 1). A possible explanation is that low conflict areas tend to have greater levels of trust and labor market complementarities between the groups than high conflict areas (Jha, 2013). In high conflict areas, Muslim leaders have no impact on coordination in the control or contact arms, and an adverse effect on coordination in the AA arm (column 2).

The impact of Hindu leaders on coordination is less sensitive to the history of conflict in the district, with the exception that coordination improvements in the contact treatment arm are in fact higher in high conflict areas (Table 5, columns 3 and 4). Hindu leaders do not improve coordination in the control arm but they do significantly improve it under both the AA and the contact treatments relative to the control, in both high conflict and low conflict areas. All of these results are robust to controlling for leaders' proposals, and therefore primarily driven by citizen reactions to leader identity (results available upon request).

7. Conclusions

We conducted a lab-in-the-field experiment in India with a view to examining leader effectiveness in achieving economic coordination in diverse societies. We find that minority (Muslim) leaders improve coordination, but majority (Hindu) leaders do not. This is primarily the result of citizen reactions to leader identity, rather than differences in leader preferences. The mechanism driving citizen reactions appears to be in-group bias, which is stronger among the minority community. We confirm this by showing that this result is blunted in towns where Muslims constitute a local majority.

Additionally, we find that in-group bias appears to be activated for members of the majority group once policies designed to improve social integration of the minority group are introduced. This significantly changes the effectiveness of Muslim leaders, as compared to the control arm. In fact, the previous results are reversed under affirmative action that is perceived to assign leadership posts to Muslims, with coordination now being

higher under Hindu leaders. A policy of intergroup contact leads to higher coordination under both Hindu and Muslim leaders, but the gains are larger under Hindu leaders.

We find striking heterogeneity in the results by an indicator of past conflict between the two religious communities, across towns within a district. The main result is that social integration policies are 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 does not vary much across high conflict and low conflict areas.

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. Recent research suggests that leaders can actively shape citizen attitudes, and that social identity itself is amenable to policy initiatives (Miguel, 2004; Blouin and Mukand, 2019). While this is beyond the scope of our current paper, it can be a fruitful direction for future research.

References

- Adams, R. B., & Ferreira, D. (2009). Women in the boardroom and their impact on governance and performance. *Journal of Financial Economics*, 94(2): 291-309.
- Ahern, K.R., and A. Dittmar (2012). The changing of the boards: The impact on firm valuation of mandated female board representation. *Quarterly Journal of Economics* 127:1:137–197.
- Alesina, A., and La Ferrara, E. (2005). Ethnic Diversity and Economic Performance. *Journal of Economic Literature*, 43(3):762-800.
- Allport, G. (1954). *The Nature of Prejudice*. New York, NY: Basic Books.
- Ashraf, N. and Bandiera, O. (2017). Social Incentives in Organizations. Working Paper.

- Bagues, M., Sylos-Labini, M. and Zinovyeva, N. (2017). Does the Gender Composition of Scientific Committees Matter? *American Economic Review*, 107(4): 1207-1238.
- Baker, P. E. (1934). Negro–White adjustment. New York: Association Press.
- Balafoutas, L. and Sutter, M. (2012). Affirmative action policies promote women and do not harm efficiency in the lab. *Science*, 335: 579-582.
- Barnhardt, S. (2009). Near and dear? Evaluating the impact of neighbor diversity on interreligious attitudes. Working paper.
- Bauer, M., C. Blattman, J. Chytilová, J. Henrich, E. Miguel, and T. Mitts. (2016). Can War Foster Cooperation? *Journal of Economic Perspectives*, 30(3): 249–74.
- Beaman, L., Chattopadhyay R., Duflo E., Pande R. and Topalova P. (2009). Powerful women: does exposure reduce bias? *Quarterly Journal of Economics* vol. 124(4): 1497-1540.
- Beaman, L., Chattopadhyay R., Duflo E., Pande R. and Topalova P. (2012). Female leadership raises aspirations and educational attainment for girls: a policy experiment in India. *Science* vol. 335(6068): 582-586.
- Besley, T. and Coate, S. (1997), An Economic Model of Representative Democracy, *The Quarterly Journal of Economics*, 112(1): 85-114.
- Benjamin, D.J., Choi, J.J. and Fisher, G. (2016). Religious Identity and Economic Behavior. *Review of Economics and Statistics*, 98(4): 617–637.
- Berge, L., Bjorvatn, K., Galle, S., Miguel, E., Posner, D., Tungodden, B. and Zhang, K. (forthcoming). Ethnically Biased? Experimental Evidence from Kenya. *Journal of the European Economic Association*.
- Bertrand, M. and Duflo, E. (2017). Field Experiments on Discrimination. *Handbook of Economic Field Experiments*, 1:309–393.
- Bhalotra, S., Clots-Figueras, I. and Iyer, L. (2019). Religion and Abortion: The Role of Politician Identity. Working paper.
- Bhalotra, S., Cassan, G., Clots-Figueras, I. and Iyer, L. (2014). Religion, politician identity and development outcomes: evidence from India. *Journal of Economic Behavior and Organization*, 104: 4–17.

- Bhavnani, R. (2017). Do the Effects of Temporary Ethnic Group Quotas Persist? Evidence from India. *American Economic Journal: Applied Economics*, 9(3): 105-23.
- Bhavnani, R., Donnay, K., Miodownik, D., Mor, M. and Helbing, D. (2014), Group Segregation and Urban Violence. *American Journal of Political Science*, 58: 226-245.
- Bisin, A., and T. Verdier (2001). The Economics of Cultural Transmission and the Dynamics of Preferences, *Journal of Economic Theory*, 97: 298 – 319.
- Blattman, C., and Miguel, E.. (2010). Civil War. *Journal of Economic Literature*, 48 (1): 3-57.
- Blouin, A. and Mukand, S. W. (2019). Erasing Ethnicity? Propaganda, Nation Building, and Identity in Rwanda. *Journal of Political Economy*, forthcoming.
- Bowles, S., Foley, D. and Halliday, S. (forthcoming) *Microeconomics: Competition, Cooperation & Conflict*.
- Brandts, J., & Cooper, D. J. (2006). A change would do you good... an experimental study on how to overcome coordination failure in organizations. *American Economic Review*, 96(3): 669-693.
- Brandts, J., & Cooper, D. J. (2007). 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): 1223-1268.
- Brandts, J., Cooper, D. J., & Fatas, E. (2007). Leadership and overcoming coordination failure with asymmetric costs. *Experimental Economics*, 10(3): 269-284.
- Brandts, J., Cooper, D. J., & Weber, R. A. (2015). Legitimacy, communication, and leadership in the turnaround game. *Management Science*, 61(11): 2627-2645.
- Brandts, J., D. J. Cooper, E. Fatas, and S. Qi (2016). Stand by me - Experiments on help and commitment in coordination games. *Management Science*, 62(10): 2916–2936.
- Brollo, F. and Troiano, U. (2016). What happens when a woman wins an election? Evidence from close races in Brazil. *Journal of Development Economics*, 122: 28-45.
- Brooks, B., K. Hoff and P. Pandey (2015). Culture and the efficiency of coordination: Experiments with high- and low-caste men in rural India. Working paper.

- Burgess, R., Jedwab, R., Miguel, E., Morjaria, A. and Padró i Miquel, G. (2015). The value of democracy: evidence from road building in Kenya. *American Economic Review*, 105(6): 1817-1851.
- Cartwright, E., J. Gillet and M. Van Vugt. (2013). Leadership by Example in the Weak-Link Game. *Economic Inquiry*.
- Chattopadhyay, R. and Duflo, E. (2004). Women as policymakers : evidence from a randomized policy experiment in India. *Econometrica*, 72(5):1409–1443.
- Chin, A. and N. Prakash. (2011). The Redistributive Effects of Political Reservation for Minorities: Evidence from India. *Journal of Development Economics*, 96 (2): 265-277.
- Cohn, A., Marechal, M. and Noll, T. (2015). Bad Boys: How Criminal Identity Salience Affects Rule Violation. *Review of Economic Studies*, 82: 1289–1308.
- Coleman, J. S. (1987). Norms as social capital. In G. Radnitzky & P. Bernholz (Eds.), *Economic imperialism* (pp. 133-155). New York: Paragon
- Corno, L., E. La Ferrara and J. Burns. (2018). Interaction, stereotypes and performance: Evidence from South Africa. Working paper.
- Devetag, G. and Ortmann, A. (2007). When and why? A critical survey on coordination failure in the laboratory. *Experimental Economics*, 10: 331-344.
- Dahl, G., A. Kotsadam and D-O. Roth (2018). Does Integration Change Gender Attitudes? The Effect of Randomly Assigning Women to Traditionally Male Teams. NBER Working Paper 24351.
- Gächter, S., Nosenzo, D., Renner, E., and Sefton, M. (2012). Who makes a good leader? Cooperativeness, optimism, and leading-by-example. *Economic Inquiry* 50: 953-967.
- Gangadharan, L., T. Jain, P. Maitra, and J. Vecci (2016). Social Identity and Governance: The Behavioral Response to Female Leaders. *European Economic Review*, 90: 302–25.
- Government of India (2014). Status of Reservation of OBC in Various States. Ministry of Social Justice and Empowerment. Government of India. Available <http://pib.nic.in/newsite/PrintRelease.aspx?relid=108754>

- Gupta, G., Mahmud, M., Maitra, P., Mitra, S., and Neelim, A. (2018). Religion, Minority Status and Trust: Evidence from a Field Experiment. *Journal of Economic Behavior and Organization*, 146: 180 - 205.
- Güth, W., M. V. Levati, M. Sutter, and E. Van Der Heijden. (2007). Leading by Example with and without Exclusion Power in Voluntary Contribution Experiments. *Journal of Public Economics*, 91: 1023–42.
- Heath, O., Verniers, G. and Kumar, S. (2015). Do Muslim voters prefer Muslim candidates? Co-religiosity and voting behaviour in India. *Electoral Studies* 38: 10-18.
- Heursen, L., Ranehill, E. and Weber, R. (2019). Are women less effective leaders than men? Evidence from experiments using coordination games. Working paper.
- Hoff, K. (2000). Beyond Rosenstein-Rodan: the modern theory of coordination problems in development. *Proceedings of the Annual World Bank Conference on Development Economics*, 145-188.
- Hoff, K., and Stiglitz, J. (2001). Modern economic theory and development. In G. M. Meier & J. E. Stiglitz (Eds.), *Frontiers of development economics: The future in perspective*. New York: Oxford University Press.
- Hogg, M.A. (2001). A social identity theory of leadership. *Pers Soc Psychol Rev*, 5: 184–200
- Ip, E., Leibbrandt, A., and Vecchi, J., (2018). How Do Gender Quotas Affect Hierarchical Relationships? Complementary Evidence from a Representative Survey and Labor Market Experiments. *Management Science*, forthcoming.
- Iyer, L., A. Mani, P. Mishra, and P. Topalova (2012). The power of political voice: Women’s political representation and crime in India. *American Economic Journal: Applied Economics*, 4(4): 165–193.
- Jack, K. and Recalde, M. (2015). Leadership and the voluntary provision of public goods: Field evidence from Bolivia. *Journal of Public Economics* 122: 80-93.
- Jaffrelot, C., Verniers, G. and Das, S. (2019). The Dwindling Presence of Religious and Other Minorities from India’s Public Sphere. Work in progress.

- Jha, S. (2013). Trade, institutions, and ethnic tolerance: Evidence from South Asia. *American Political Science Review* 107(4): 806-832.
- Kaysser, N., Bhalotra S., Clots-Figueras I., and Iyer L. (2015). Hindu-Muslim violence in India 1950-2010: an update of the Varshney-Wilkinson dataset on Hindu-Muslim violence in India. Work in progress.
- Kremer, M. (1993). The O-Ring Theory of Economic Development. *Quarterly Journal of Economics*, 108(3): 551-575.
- Leibbrandt, A., Wang, L.C., Foo, C., (forthcoming). Gender Quotas, Competitions, and Peer Review: Experimental Evidence on the Backlash Against Women. *Management Science*.
- Lett, H. A. (1945). Techniques for achieving interracial cooperation. Proceedings of the Institute on Race Relations and Community Organization. Chicago: University of Chicago and the American Council on Race Relations
- Levy, D.M., Padgitt, K., Peart, S.J., Houser, D. and Xiao, E. (2011). Leadership, cheap talk and really cheap talk. *Journal of Economic Behavior and Organization*, 77(1): 40-52.
- Lowe, M. (2017). Types of contact: A field experiment on collaborative and adversarial caste integration. *Working Paper*.
- Matsa, D. A., & Miller, A. R. (2013). A female style in corporate leadership? Evidence from quotas. *American Economic Journal: Applied Economics*, 5(3): 136-69.
- Meyersson, E. (2014). Islamic rule and the empowerment of the poor and pious. *Econometrica*, 82(1): 229–69.
- Miguel, E. and Roland, G. (2011). The Long-Run Impact of Bombing Vietnam. *Journal of Development Economics*, 96: 1-15.
- Miguel, E. (2004). Tribe or Nation?: Nation Building and Public Goods in Kenya versus Tanzania. *World Politics* 56(3): 327–62.
- Mitra, A. and D. Ray. (2014). Implications of an Economic Theory of Conflict: Hindu-Muslim Violence in India. *Journal of Political Economy*, 122 (4): 719–65.
- Niederle, M., Segal, C. and Vesterlund, L. (2013). How costly is diversity? Affirmative action in light of gender differences in competitiveness. *Management Science*, 59(1):

1-16.

- NITI Aayog. (2016). Eliminating Poverty: Creating Jobs and Strengthening Social Programs. Occasional Paper 2. Government of India. Available at http://niti.gov.in/writereaddata/files/OccasionalPaper_No2_Poverty.pdf
- Osborne, M. and Slivinski, A. (1996). A Model of Political Competition with Citizen-Candidates, *The Quarterly Journal of Economics*, 111(1): 65-96
- Paluck, E. L., Green, S. and Green, D. P. (2018). The Contact Hypothesis Revisited. *Behavioural Public Policy*, 1-30.
- Pande, R. (2003). Can Mandated Political Representation Provide Disadvantaged Minorities Policy Influence? Theory and Evidence from India. *American Economic Review*, 93(4): 1132-1151.
- Pettigrew, T. F. and Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of personality and social psychology*, 90(5):751.
- Polania-Reyes, S. (2016). Coordination as unintended benefit: Lab-in-the-field evidence from a conditional cash transfer program. Working Paper.
- Rao, G. (2018). Familiarity Does Not Breed Contempt: Diversity, Discrimination and Generosity in Delhi Schools. Working Paper.
- Rodrik, D. (1999). Where Did All the Growth Go? External Shocks, Social Conflict, and Growth Collapses. *Journal of Economic Growth*, 4(4): 385–412.
- Scacco, A. and S. Warren (2018). Can Social Contact Reduce Prejudice and Discrimination? Evidence from a Field Experiment in Nigeria. *American Political Science Review* 1-24. doi:10.1017/S0003055418000151
- Schotter, A., and K. Weigelt (1992). Asymmetric Tournaments, Equal Opportunity Laws and Affirmative Action: Some Experimental Results. *Quarterly Journal of Economics*, 107(2): 511–539.
- Shih, M., Pittinsky, T. L., & Ambady, N. (1999). Stereotype Susceptibility: Identity Salience and Shifts in Quantitative Performance. *Psychological Science*, 10(1), 80–83.

The Economist (2016). India's Muslims: An uncertain community, October 29.

Available: <https://www.economist.com/asia/2016/10/29/an-uncertain-community>

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

Weber, R., Camerer, C., Rottenstreich, Y. and Knez, M. (2001). The Illusion of Leadership: Misattribution of Cause in Coordination Games. *Organization Science*. 12. 10.1287/orsc.12.5.582.10090.

Wilkinson, S. (2004). *Votes and Violence: Electoral Competition and Ethnic Riots in India*. Cambridge: Cambridge Univ. Press

Figure 1
Leader Identity and Coordination (Control Group)

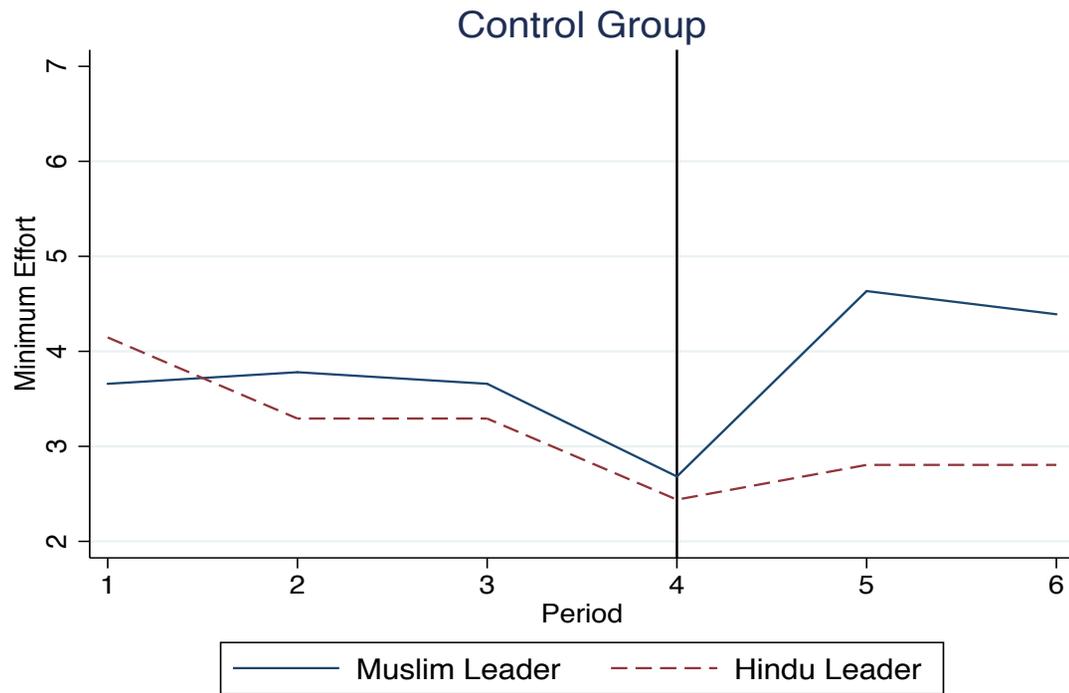


Figure 2
Leader Effectiveness Across Policy Environments

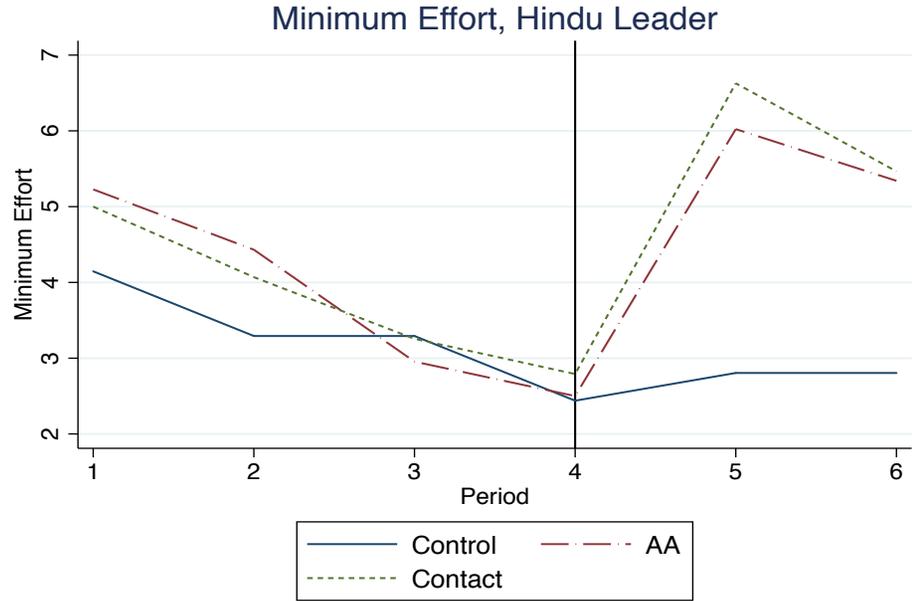
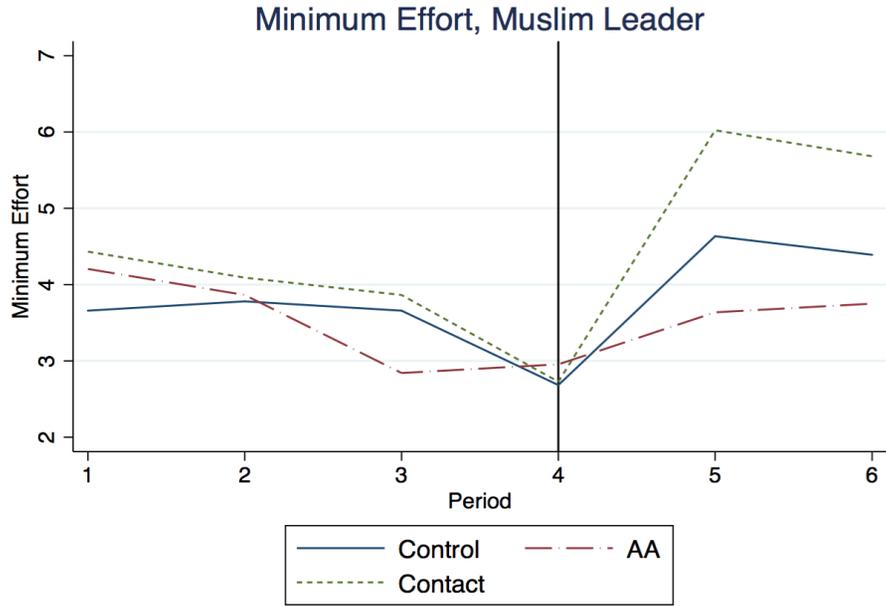


Table 1
Demographic Characteristics Across Towns in Different Treatment Groups

	Control	Contact	AA	p-value of difference (Contact v Control)	p-value of difference (Control v AA)
Panel A: Town characteristics (Census 2011)					
Total Population	13709	14646	14011	0.727	0.917
Fraction Muslim	0.43	0.40	0.38	0.803	0.631
Women per 1000 men	913	908	902	0.665	0.415
Girls per 1000 boys (age 0-6)	919	911	904	0.738	0.463
Fraction literate	0.58	0.57	0.58	0.808	0.978
Fraction women literate	0.50	0.49	0.50	0.880	1.000
Fraction SC	0.13	0.13	0.13	0.986	0.949
Fraction ST	0.00	0.00	0.00	0.488	0.323
Work participation, male	0.46	0.47	0.46	0.355	0.837
Work participation, female	0.11	0.14	0.14	0.167	0.262
Fraction working men in household enterprise	0.07	0.08	0.07	0.612	0.912
Fraction working men in non-agri and non-HH enterprise	0.65	0.60	0.64	0.555	0.894
Panel B: Experiment participant characteristics					
Male	0.67	0.60	0.60	0.065	0.059
Age	24.1	24.7	22.5	0.382	0.014
Scheduled Caste or Scheduled Tribe	0.16	0.13	0.16	0.357	0.847
Other Backward Caste	0.58	0.50	0.51	0.055	0.091
Only primary education	0.13	0.12	0.12	0.646	0.686
Completed grade 10	0.23	0.19	0.24	0.132	0.909
Completed grade 12	0.39	0.35	0.37	0.347	0.659
Completed college	0.25	0.34	0.27	0.007	0.497
Family monthly income <=Rs 5,000	0.19	0.21	0.16	0.568	0.335
Family monthly income Rs 5,000-10,000	0.31	0.30	0.29	0.875	0.574
Family monthly income Rs 10,000-15,000	0.24	0.19	0.22	0.147	0.507
Family monthly income Rs 15,000-30,000	0.17	0.23	0.23	0.069	0.063
Family monthly income > Rs 30,000	0.09	0.07	0.11	0.302	0.550
Pray several times a day	0.26	0.28	0.25	0.412	0.856
Pray once a day	0.61	0.58	0.60	0.519	0.811
Pray less than once a day	0.14	0.13	0.15	0.889	0.577

Notes: p-value of difference is calculated by regression on the treatment dummies and using robust standard errors.

Table 2
Leader Identity and Coordination (Control Group)

Dependent variable: Minimum Effort in the Group

	(1)	(2)	(3)	(4)	(5)	(6)
	Muslim Leaders	Hindu Leaders	Muslim Leaders	Hindu Leaders	All Leaders	All Leaders
Leader (Period>4)	1.067 (0.494)	-0.488 (0.381)	1.067 (0.508)	-0.488 (0.392)	-0.488 (0.379)	
Muslim Leader * (Period>4)					1.555 (0.620)	
Muslim Leader					-0.492 (0.616)	1.272 (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 parantheses, 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 3
Leader Identity and Individual Effort (Control Group)

	(1)	(2)	(3)	(4)	(5)	(6)
	<i>Individual Effort</i>				<i>Minimum Group Effort</i>	
	Muslim Leaders		Hindu Leaders		Muslim Leaders	Hindu Leaders
	Muslim Employees	Hindu Employees	Muslim Employees	Hindu Employees		
Leader (Period>4)	1.157 (0.434)	0.377 (0.531)	0.015 (0.517)	-1.156 (0.408)	1.739 (0.589)	-0.326 (0.498)
Leader*Muslim population majority					-1.531 (1.021)	-0.368 (0.800)
Observations	486	498	486	480	246	246
R-squared	0.309	0.278	0.242	0.280	0.481	0.436
Town FE	Yes	Yes	Yes	Yes	Yes	Yes
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes
Religious Controls	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parantheses, 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 4
Policy Environments and Leader Effectiveness

Dependent variable: Minimum Effort in the Group

	(1)	(2)	(3)	(4)
	Muslim Leaders	Hindu Leaders	Muslim Leaders	Hindu Leaders
Leader (Period>4) β_1	1.067 (0.495)	-0.488 (0.382)		
Leader (Period>4)*AA β_2	-0.840 (0.814)	2.391 (0.671)	-1.625 (0.837)	2.370 (0.627)
Leader (Period>4)*Contact β_3	1.007 (0.752)	2.755 (0.651)	1.738 (0.801)	2.85 (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 parantheses, 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 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.

Table 5
Does a History of Conflict Matter for Leader Effectiveness across Policy Environments?

Dependent variable: Minimum Effort in the Group

	(1)	(2)	(3)	(4)
	Muslim Leaders	Muslim Leaders	Hindu Leaders	Hindu Leaders
	Low Conflict Areas	High Conflict Areas	Low Conflict Areas	High Conflict Areas
Leader (Period>4)	1.181 (0.577)	0.978 (0.770)	-0.486 (0.555)	-0.489 (0.537)
Leader (Period>4)*AA	0.757 (1.194)	-2.176 (1.036)	2.236 (1.020)	2.520 (0.912)
Leader (Period>4)*Contact	2.014 (1.100)	0.320 (1.027)	1.663 (0.780)	3.470 (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 parantheses, 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.