

Caste, bureaucracy, and the limits to political affirmative action in India

Eduardo Campillo Betancourt^{*†}

December 30, 2022

[Click here for the most recent version](#) ([Google Drive](#)) ([Ungated](#))

Abstract

The aim of political affirmative action policies is to ensure that disadvantaged groups are represented in their governments and, in turn, that laws preferred by this group are more likely to be instituted. Often, however, they have not been found to be effective. I explore two reasons for this: 1) these policies target large, heterogeneous groups and ignore rigid boundaries within them, and 2) minority politicians might not have direct control over resource allocation. I focus on the context of India, where certain seats for state legislators are reserved for the historically discriminated lower castes (Dalits). Dalits belong to many heterogeneous castes and state legislators must influence local bureaucrats in order to affect the distribution of public goods. To overcome a lack of individual caste data, I exploit the link between names and caste membership and create a new dataset including the caste of workers involved in a public workfare program (NREGA). With this dataset and based on the fact that constituencies are reserved for low-caste legislators through a population cutoff rule, I use a regression discontinuity design to estimate the effect of having a low-caste state representative on the timing of payments to Dalit laborers within NREGA. I find that low-caste workers represented by a low-caste state legislator experience a 12% higher probability of receiving their payments late. The effect is constant across all individual castes and is concentrated in areas where politicians have lower bargaining power over the local bureaucrats. Hence, my findings point to the importance of considering vertical power structures when designing policies aimed at empowering under-represented minorities around the world.

^{*}Email: eduardocampillo2023@u.northwestern.edu

[†]I would like to thank Chris Udry, Lori Beaman, Seema Jayachandran, and Nancy Qian for their invaluable mentorship and support. I thank Kensuke Maeba, Lauren Clingan, José Luis Flor Toro, Eilidh Geddes, Nicole Holz, Udayan Vaidya, Paul Kim, and seminar participants in the Northwestern Development and HELP Seminars, and in the Development Lunch and Development Group Advising for their valuable comments and suggestions. Prashant Batule and Hrishikesh Morankar offered invaluable assistance with scrapping the NREGA website. This research was supported in part through the computational resources and staff contributions provided for the Quest high performance computing facility at Northwestern University, which is jointly supported by the Office of the Provost, the Office for Research, and Northwestern University Information Technology. All errors are my own.

1 Introduction

The aim of political affirmative action policies is to ensure that disadvantaged groups are represented in their governments and, in turn, that laws preferred by this group are more likely to be instituted. A prominent example of political quotas is the reservation of government seats for lower-caste (Dalit) Indians.¹ Confined to the bottom rung of society, Dalits have long struggled to gain access to most arenas of public life. In an attempt to ensure at least some degree of political representation for Dalits, following Independence from the British Empire, the Indian government reserved numerous government seats for lower-caste legislators in state assemblies (also known as MLAs). Despite the government's intentions, numerous empirical evaluations of these political quotas have documented underwhelming, and often null, effects on the welfare of Dalit citizens, calling into question the limitations of mandated representation for these disadvantaged population (Pande, 2003; Chin and Prakash, 2011; Jensenius, 2015).

In this paper, I explore two main limitations of these policies in helping under-represented groups. First, these programs target broad groups of people and ignore important differences within the disadvantaged beneficiaries. Although it is perhaps more practical to target populations at a higher level, this decision could lead to significant dispersion in the impacts of these policies. In particular, the null effect evidenced in the literature could mask significant caste heterogeneity in the benefits to Dalit constituents. Second, in settings where politicians do not have direct control over the allocation of public goods (such as India), securing seats for underprivileged candidates might not translate into better outcomes for the under-represented groups. Local bureaucrats often manage the logistics of public welfare programs and, for disadvantaged politicians, bargaining with these regional officers could be less effective.

India offers an ideal setting to explore both limitations to political affirmative action policies. Dalits belong to hundreds of individual castes. Each caste forms a highly endogamous and clearly defined group, so much that scholars of India recognize castes—and not caste categories, such as Dalits—as the main unit of social organization in the country (Munshi, 2019). The clear and consequential divisions within the targeted population motivate understanding the potentially disparate effects of caste quotas in the state governments. Moreover, political scientists have recently promoted caste inequality as an important driver of political polarization in the country and as a catalyst to the support for the populist, right-wing Bharatiya Janata Party (BJP) of Narendra Modi (Jaffrelot, 2022). In addition, MLAs control the allocation of public goods only through their influence over local bureaucrats. The stability and high remuneration that bureaucratic positions offer, as well as the intensive selection process for these posts, often selects more privileged Indians. Requiring Dalit MLAs to lobby more socioeconomically advantaged bureaucrats exacerbates the politicians' limitations to improve their constituents' welfare.

¹Even though most of the economics literature refers to Dalits by their official title (Scheduled Castes), I choose this term as it is their preferred demonym.

My empirical analysis proceeds in three steps. In the first part of this study, I estimate how being represented by a Dalit state legislator affects the provision of public goods to Dalit citizens within the world’s largest public workfare program: NREGA. While I am interested broadly on the provision of public goods and services, I focus particularly on the time it takes for NREGA laborers to receive payments following the completion of a project. I choose this measure because it is highly politically salient and it is directly controlled by the local bureaucracy, meaning that state legislators ought to be able to influence this outcome. Second, I evaluate whether the effect of having a low-caste MLA varies depending on the worker’s specific caste. This exercise will directly evidence which castes, if any, benefit the most from this policy and will shed light on whether the reservation of seats for Dalits in state legislatures has contributed to inequality across castes, within the Dalit category. The third and final section of my analysis studies how the bargaining power of MLAs over local bureaucrats mediates the politicians’ ability to aid their Dalit constituents. In order to affect the distribution of public goods and services, MLAs must be able to influence the actions of local bureaucrats. This lack of direct control over the allocation of resources could hinder the ability of Dalit MLAs, in particular, given the stark socioeconomic distance between these politicians and the (often) privileged bureaucrats.

As part of this analysis, I overcome two main empirical challenges: the endogenous selection of reserved seats for Dalit MLAs and a lack of data on individual castes. Constituencies are split into those where only Dalit candidates are allowed to run for election and the rest where the candidate pool is unrestricted using a cutoff rule based on the share of Dalit residents in these areas. Using this institutional feature, I estimate via a regression discontinuity design the effect of being represented by a Dalit MLA on the provision of timely payments to NREGA workers. Given this allocation rule, comparing worker in an average reserved and unrestricted constituencies would conflate the impact of these caste-based quotas with many other differences that correlate with having a higher Dalit share of the population. In particular, due to the legacy of discrimination against this minority group, areas with high concentrations of Dalits tend to have educational attainment and access to a broad range of public goods. In contrast, the regression discontinuity design compares only workers leaving in areas that were only narrowly designated as reserved versus those in constituencies that fell just short of the cutoff. I argue that these two sets of constituencies are broadly comparable and differ only in the identity of their elected state legislator.

Considered a sensitive social and political issue, individual caste information is almost entirely left out of large-scale data collection efforts in the country. Instead, comprehensive datasets often include only the broader caste category of respondents, e.g., whether the person is Dalit or not. Datasets that do contain caste information for respondents suffer from one of two problems: either they are not representative at levels below the district, or they are focused on specific villages or other small areas of the country.² In order to overcome the lack of detailed caste information, I make use

²Examples of these data are the Indian Human Development Surveys or the Rural Economic and Demographic Surveys. While both of these were constructed to be nationally representative and cover all states in the country,

of a common-knowledge proxy for caste: names. Surnames, in particular, are highly influenced by the caste to which a person or household belongs (Singh, 1996). Using the secure component of the 2006 Rural Economic and Demographic Survey (REDS), which includes a nationally representative sample of household heads, I estimate the joint distribution of castes and names. With these empirical probabilities in hand, I am able to predict the individual caste of NREGA workers.

Through an internal cross-validation exercise within the sample of names and known castes, I show that surnames allow me to estimate the correct caste of individuals with a likelihood of 73.6%. The out-of-sample performance of these predictions decreases when given names are also used. This is in line with the qualitative literature, which argues that given names are not used as signals of caste but of class and socioeconomic status (Singh, 1996). The procedure is able to correctly predict the caste of workers in the smaller castes with relatively high probability (60%).

Turning to the research question at hand, I estimate that having a Dalit MLA leads to an economically large but only marginally significant increase in the frequency of late payments to Dalit NREGA laborers. Dalit workers represented by a politician of their same caste category are 3.3 percentage points more likely to receive their payment past the stipulated deadline when compared to Dalit workers represented by a non-Dalit MLA. This represents a 12.4% increase relative to the comparison group mean. Second, I show that there is no widespread evidence of one Dalit caste being particularly favored at the expense of others. This contradicts the current right-wing narrative that singles out the largest Dalit group as having cornered the benefits of caste-based affirmative action. Third, this deleterious effect on Dalit workers is entirely concentrated in areas where the MLA has lower control over local bureaucrats. This result points to the importance of considering vertical power structures to address the ineffectiveness of political quotas for Dalits in India.

Caste and political reservations for low-caste politicians are two topics that have been widely studied in the economics literature. The focus of this work has been on either village council presidents or state representatives (the theme of this study). Early work on village-level reservations found positive effects of electing a Dalit village council president (Pradhan) on public good provision for low-caste residents (Besley et al., 2004, 2005; Bardhan et al., 2010). Dunning and Nilekani (2013), however, finds no effect using within-subdistrict variation (the level at which seats for Pradhans are reserved). Looking instead at reserved seats in state assemblies, Pande (2003), the seminal paper in this area, finds that having a greater share of Dalit MLAs leads to an increase in the number of public sector jobs allocated for Dalit workers but has no effect on spending for low-caste welfare programs. Echoing these results, Chin and Prakash (2011) conclude that being represented by a low-caste MLA has no effect on low-caste poverty. These two papers compare exogenous increases in the share of reserved MLA seats across states. Instead, using variation across reserved and general constituencies throughout India, Jensenius (2015) finds no evidence of improved welfare of Dalit constituents resulting from reservation. The effect I estimate on Dalit NREGA workers as a whole

they are concentrated in specific villages within select districts. Another example is district-level caste distributions in historical censuses.

is largely in line with the findings discussed here. My main contribution to this area of the literature is to consider how the effects of reservations vary by the caste of the recipients. By doing so, I aim to understand not just how caste-based reservations can affect the gap between Dalits and upper-caste Indians, but also how this policy impacts inequality within the lower-caste category. Recent work highlights the importance of looking not just across caste categories, but also within them (Joshi et al., 2022).

My project relates to a broader body of work on the relationship between caste (or ethnicity, more broadly) and public good provision. This work explores two main questions. The first is how caste (ethnic) heterogeneity relates to the political selection process and the performance of representatives (Anderson and Francois, 2017; Munshi and Rosenzweig, 2008, 2013). These papers point out the importance of how individual castes interact with village leaders and motivate the use of caste (as opposed to caste categories) as the relevant social unit of political mobilization. The second question is whether the existence of caste networks in a community enables clientelistic relationships between politicians and the governed (Anderson et al., 2015; Banerjee and Somanathan, 2007). Contrary to the first set of projects in this literature, these studies focus on aggregate measures of caste fractionalization instead of how the provision of public goods is impacted by the specific caste of individuals. My work relates to this broad body of work in that it explores how the caste of workers matters for understanding the benefits of being represented by a person of the same caste category. I contribute to this body of work by focusing on state legislators, instead of village leaders, and studying the relationship between their performance and the caste of their constituents. As mentioned above, state legislators play a central role in connecting individuals within their constituency to public welfare benefits. In addition, the set of incentives that drive their performance is drastically different from those of village leaders and this could have important consequences on how they direct the provision of public goods and services. Unlike village leaders, state representatives work within a political party, which can exert a great deal of control over their policy decisions. In my mechanisms section, I show how the political party of the legislator plays an important role in the impact of caste reservations, highlighting the importance of studying the interaction between caste-based political affirmative action (as it pertains to state legislators) and the caste of beneficiaries.

The remainder of the paper proceeds as follows. In Section 2 I present some background information on both the sociopolitical context and the government organizational structure vis-à-vis public good provision. Section 3 gives an overview of the data used in my analysis. Section 4 details the caste-prediction procedure and evaluates its performance. Section 5 expounds on my empirical strategy. Sections 6 and 7 discuss the results and mechanisms, respectively, and Section 8 concludes.

2 Social and political context

2.1 Caste categories and caste

Indian society, especially within the majority Hindu population, is stratified by caste.³ Castes are hereditary and endogamous social groups with origins in the Hindu religious texts (Moorjani et al., 2013). Caste plays a vital role in the functioning of labor and credit markets, with referral networks and risk sharing working efficiently within a caste group but often breaking down across castes even within the same caste category. Voting choices, in addition, are largely influenced by the caste identity of the electors and the candidates (Jaffrelot and Kumar, 2009).

Reifying historical hierarchies, the Indian government groups castes into three main categories: Dalits (officially known as Scheduled Castes, abbreviated as SCs), Other Backward Castes (OBCs), and Forward Castes (FCs). These categories are highly relevant for policy design in India, with most caste-based government programs targeting people, not based on their individual caste, but on whether or not they belong to one of these categories. The focus of this study is on Dalits, socially and economically marginalized groups located at the lowest rung of the caste ladder. They account for 20.7% of the population of Uttar Pradesh (Census of India, 2011). Each of these caste categories are composed of hundreds of individual castes. The literature on caste highlights that caste, not caste categories, are the relevant building blocks of social networks driving economic and political phenomena in the country (Munshi, 2019). Figure 1 represents the distribution of castes within the Dalit category. By a vast margin, the largest caste in the state of Uttar Pradesh is the Jatav caste (accounting for nearly six in every ten Dalits). This is precisely the group routinely accused by the right-wing of having monopolized the benefits from Dalit political representation, be it in the form of affirmative action programs or through the rise of “low-caste political parties” in recent decades.⁴ The next largest groups are the Pasis and Balmiki castes, each composing around 5% of the population. Together, the nine largest castes make up close to 90% of the Dalit population in the state. It is important to note that, despite all the castes being drastically less represented at the state level relative to the Jatavs, these social groups can be more concentrated in some areas within Uttar Pradesh. In fact, each of these castes comprises over half of the Dalit population in at least one village of this state.

2.2 Caste-based political affirmative action

Dalits often performed the least desirable occupations—such as manually cleaning public latrines or tanning leather—and contact with them was deemed “polluting” for people in higher castes,

³Hindus represent almost 80% of the total population both in the country as a whole and within the state of Uttar Pradesh, the geographic focus of this study. The second largest religious group are the Muslims (14%). The remaining share of Indians belong to the Christian, Sikh, or Buddhist tradition, among others (Census of India, 2011).

⁴As the title suggests, these are political parties that are overwhelmingly supported by Dalit voters and which explicitly favor the interests of members of these castes.

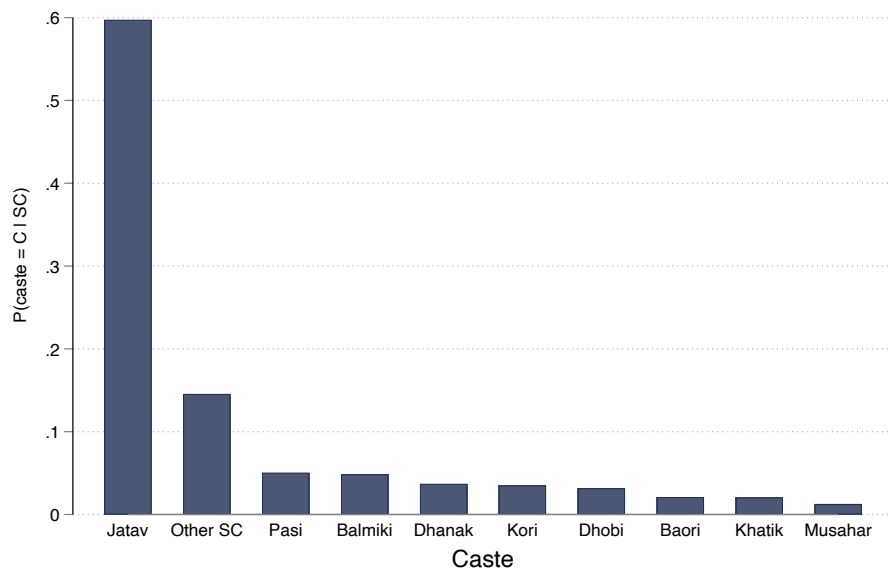


Figure 1: Caste distribution of Dalits within Uttar Pradesh

Note: This figure presents the distribution of Dalit individuals within each of the largest nine castes, plus a group incorporating the remaining smaller Dalit castes. The sample is restricted to observations within the state of Uttar Pradesh, the focus of this study.

Source: Rural Economic and Demographic Survey (2006)

thus leading many to call them “untouchable.” This long history of “untouchability” has resulted in significant and persistent socioeconomic disadvantages for Dalits. In an effort to correct these injustices, since the country’s independence the Indian government has implemented a series of affirmative action policies aimed specifically towards members of the Dalit castes. The focus of this paper is the study of political reservations for Dalit candidates in the state legislatures, which were created to ensure that Dalits were directly represented in these governing bodies. In Uttar Pradesh, the state assembly holds 403 seats, 85 of which are restricted to candidates belonging to a Dalit caste (reserved constituencies) whereas the rest can be contested by any person, regardless of their identity (general constituencies).⁵ The boundaries of these administrative units were last defined in 2008 based on the 2001 Population Census at the time. Each constituency is fully contained within a district.⁶ The resulting ACs were roughly equal in terms of their overall population, by design, but varied significantly in their share of Dalit residents. It is this variation that was used to determine the specific ACs that were to be reserved for Dalit candidates, a process that I discuss in more detail Section 5. Finally, in either type of constituency, state legislators (MLAs) are chosen by all voters, regardless of their caste identity vote, in their respective constituency’s elections. The winner is defined as the candidate who collects the largest share of votes in that constituency (first-past-the-post).

2.3 MLAs and public good provision

While their title would imply that their work is largely legislative, the most common role of these politicians is to help their constituents access to government welfare schemes (Jensenius, 2015). In order to affect the distribution of public benefits in their constituencies, MLAs interact with two main levels of public administration: Gram Panchayats (village executive governments) and blocks. Unlike ACs and panchayats, blocks are headed by career bureaucrats (Block Development Officers or BDOs). BDOs are not chosen democratically but rather crucially depend on higher-level politicians (such as MLAs) to get promoted or reallocated to more favorable postings (Iyer and Mani, 2012). It is precisely through this channel that MLAs can impact the allocation of public goods. Geographically, ACs often contain several blocks but some blocks span various constituencies: about half of the blocks in Uttar Pradesh are partially contained within more than one constituency. Intuitively, the influence of an MLA over a BDO should be higher in areas where the bureaucrat answers to only that politician. In areas where a block spans multiple ACs, the BDO can bargain

⁵There are two constituencies in this state that are reserved for candidates of what are called Scheduled Tribes. These are communities that exist outside the caste order and generally live in secluded areas of the country. For the purposes of this study, I will simply abstract from the distinction between a general constituency and those reserved for Scheduled Tribes candidates.

⁶It was the independent Delimitation Commission of India that carried out this process. This group was created by the central government with the explicit goal to create constituencies of a similar population size within each state. Iyer and Reddy (2013) finds that “the redistricting process to be politically neutral for the most part, though a few politicians who were advisory members for the redistricting process were able to avoid unfavorable redistricting outcomes for their specific constituencies.” Prior to 2008, the last delimitation was carried out in 1971.

with multiple MLAs and be less beholden to the will of any given one (Gulzar and Pasquale, 2017). This is a feature of the public administration in Uttar Pradesh that I will exploit in my mechanisms section.

There are two specific public goods that I explore in this paper: the National Rural Employment Guarantee Act (NREGA) and funding for school facilities. It has been documented in the literature that MLAs can impact both of these programs (Gulzar and Pasquale, 2017; Sarkar, 2019; Gulzar et al., 2020). NREGA is the largest public workfare program in the world. It guarantees 100 days of paid work each year to any household in India. Compensation in this program is set at the state minimum wage, which leads to participants being mostly low income workers. One of the most consequential roles of BDOs within MNREGA is to control the payments to workers. Timely payment of wages in NREGA is a large concern in India, particularly for Dalit workers. A payment is considered late by the program’s guidelines if it issued more than 15 days after the completion of work, a large delay when one considers the economic circumstances of the participants. Given the political salience of timely payments, it stands to reason that MLAs would choose to pressure BDOs within their constituency to speed up the remuneration process. Funding for schools is similarly controlled by the block bureaucrats. Importantly, whereas the influence of MLAs under NREGA ought to impact the poorer section of the electorate, increases in school funding stand to disproportionately benefit the more privileged segments of the population. Together, the two public goods can be useful in understanding how the incentives of MLAs to provide for their constituents vary depending on the socioeconomic status of the beneficiaries.

3 Data overview

3.1 Political data on constituencies and blocks

The first dataset that I use contains information at the block and constituency levels relevant to the political landscape of Uttar Pradesh. The most important variables here are the share of Dalit residents and the reservation status of each AC. The Dalit population shares are taken from the reports produced by the Delimitation Commission. Out of 403 constituencies, 85 are reserved for Dalit candidates and the remainder (318) are “general” constituencies without any restrictions on the candidate pool. This data contains information on the winning candidate’s party for the 2017 State Assembly elections, the legislative period relevant to my study. The latter two variables were obtained directly from the Election Commission’s website. In order to test for balance across both sides of the discontinuity, I include a set of socioeconomic indicators from 1971 and 2001 taken from the replication package for Jensenius (2015) and mapped to the post-2008 constituency boundaries using the constituency crosswalk developed by Raphael Susewind.⁷ Finally, from a mapping of blocks to constituencies, I create a variable denoting the number of ACs spanned by each block. I

⁷These crosswalk files are available on Prof. Susewind’s website.

use this measure in my mechanisms section where I compare workers whose block bureaucrat answers to a single MLA—the block is contained within a single AC—versus those whose bureaucrat answers to multiple—the block is spread across multiple ACs.

3.2 Worker-level outcome data from NREGA

I constructed the outcome variables for this study by scrapping the NREGA online portal and obtaining a random sample of the universe of workers who were active in either 2017 and 2018. To construct the random sample, within each block in Uttar Pradesh, I selected 10% of the Gram Panchayats—these are groups of one or more villages and the geographical unit that follows the block. The purpose of this random sampling was to ease the computational burden of web scrapping the data of all NREGA workers (5.7 million in 2020) while aiming to have the broadest possible geographic inclusion of constituencies in the state. Importantly, each worker can be identified by their full name and that of their father or husband (in the case of married women). In addition, this website provides the village, gram panchayat, and block of residence, as well as the workers’ caste category, age, and sex. Most of the results in the paper focus only on Dalit laborers totaling over 190,000 worker-year observations.

The two outcome variables used in this study relate to the timing of payments made to workers. After a worker completes their assigned employment period, they are entitled to receive their wages within 15 days at the latest. This definition of a late payment comes directly from the NREGA Operating Manual (Ministry of Rural Development, 2013). For each worker, the portal includes information of the amount paid and the number of days until the bureaucrat processed the payment for each period of employment. From these variables, I construct two measure of how late payments are. The first is a share of late payments and the second is the average number of days until payment, both spanning all work done in the year. To construct the share of late payments, I begin by creating an indicator variable for each wage amount equal to one if it was paid after the required 15 days, and zero otherwise. I then average across all payments made to that worker over the year, using as weights the total amount paid in each occasion. I calculate the average time to payment in a similar manner, but averaging over the total number of days until the wages are processed instead of the indicator variable I describe above. On average, I estimate that workers wait 14 days after the completion of their work to receive their wages. This translates to almost 27% of wages being paid late, as per the program’s definition. These statistics have only worsened over the years, with a study estimating that 71% of wages were paid after the required deadline during the first half of 2021.

3.3 Representative sample of names and castes

The final dataset that I use in my analysis is the secure component of the 2006 Rural Economic and Demographic Survey (REDS). The main feature of this dataset is that it contains a nationally-

representative sample of household heads and it includes their names as well as their sub-caste. With this sample in hand, I can construct a reliable estimate of the joint distribution of names and castes that allows me to predict the caste of NREGA workers based on their name. I make two restrictions to the sample of this dataset. First, I restrict the sample to those households that reside in Uttar Pradesh—the focus of this study—or one of the seven other North India states that compose the Hindi Belt: Bihar, Chhattisgarh, Haryana, Jharkhand, Madhya Pradesh, Rajasthan, and Uttarakhand. Instead of using only Uttar Pradesh, I include these seven other states that all have Hindi as their primary language. Coupled with the fact that the qualitative literature suggests that naming conventions are common to the Northern India region, this allows me to innocuously expand the sample of households that I use for prediction, meaning a wider set of possible names associated with each caste Singh (1996). The ten most predominant Dalit castes in Uttar Pradesh, the focus of this study, spread into other states in the Hindi Belt, highlighting the relevance of including these adjacent geographical entities. Within these eight states, I further restrict the sample to individuals belonging to the Scheduled Caste (Dalit) category. By restricting both the NREGA and REDS sets of observations to Dalits, I avoid some of the more egregious caste prediction errors, such as categorizing an “upper caste” or Muslim person as belonging to a Dalit caste. Both of the former groups have a socioeconomic status and relationship to public good provision that is vastly different to that of Dalits. The final sample consists of about 40,000 household heads. I know each of their full names, that of their father or husband, and their sex and age. The REDS dataset contains much more information on these respondents, however, since these variables have no counterpart in the NREGA dataset, I could not use it to predict the workers’ castes.

4 Predicting caste using names

The choice of a name is often informed by a family’s cultural and social background. Within the Indian context, it has long been common knowledge that there is a strong link between a person’s caste and their name. In Northern India (where Uttar Pradesh is located), for instance, names “consist of a personal name and a surname which is generally based on the caste name” (Singh, 1996). Personal names, in contrast, often stem from religious texts or serve to indicate social status, rather than caste belonging (Jayaraman, 2005; Pal, 2019). Taking this convention as given, I exploit the availability of name and caste data for a representative sample of households from the Rural Economic and Demographic Survey (REDS) 2006 wave to estimate the joint probability distribution of these two variables in Uttar Pradesh. I focus on the nine most prominent Scheduled Castes in the state—representing 90% of this population—and group the remaining Dalit castes under “Other SC.”⁸ Operationally, I calculate the share of people in the REDS data with a specific

⁸The ten castes are, in alphabetic order: Balmiki, Baori, Dhanak, Dhobi, Jatav, Khatik, Kori, Musahar, Pasi, and “Other SC.”

name that belong to each of these ten castes. I define as “name” a name-part with unique Hindi pronunciation. By name-part, I mean each component of a name, that is, a person whose name is “Rajesh Singh Kumar” has three name-parts: “Rajesh”, “Singh”, and “Kumar.” Therefore, I calculate $P(\text{caste}|\text{Rajesh})$, $P(\text{caste}|\text{Singh})$, and $P(\text{caste}|\text{Kumar})$, not $P(\text{caste}|\text{Rajesh Singh Kumar})$.⁹ Lastly, due to the vast number of Roman spelling variants of Hindi names in my data, I consider any two name-parts with the same Hindi pronunciation as being the same name-part. I argue that this leads to more meaningful caste predictions, seeing as it stands to reason that names are most commonly used orally and not in written form among the workers of this study. Hence, it is different pronunciations that would distinguish one person’s name from another and not homophonous differences in spelling.

4.1 Internal validation of caste predictions

Using the fact that in the REDS dataset I can observe a person’s name and their true caste, I run a series of cross-validation exercises to estimate the accuracy and precision of my predictions. This process involves splitting the dataset randomly into two subsets: training and testing. I use the training subset to estimate $P(\text{caste}|\text{name-part})$. Then I use these probabilities to predict the probability that people in the testing dataset subset belongs to each of the ten Dalit castes—pretending I do not know the *true* caste of each observation. I do so by assigning each of the individual’s name-parts the predicted probability for each caste, conditional on having that name part. Finally, I aggregate these predictions across all name-parts associated with an individual using a formula based on Bayes’ rule.¹⁰ This results in ten continuous measures of the likelihood that an individual (i) belongs to each caste (C), conditional on all their name-parts:

$$P(\text{caste}_i = C | \text{names}_i)$$

Lastly, I compare the survey-recorded caste of individuals (the true caste) with my name-based predictions. I measure the performance of this procedure, separately for each caste (C), using two metrics:

$$\text{Efficiency} : \frac{1}{N} \sum_{i=1}^N P(\text{caste}_i = C | \text{names}_i, \text{true caste}_i = C) \quad (1)$$

⁹Ideally, I could include all combination of names in my estimation to create even better predictions. However, given the restricted sample size of the REDS dataset, including all different permutations of the name-parts would greatly hinder my statistical power to make predictions.

¹⁰Intuitively, this formula serves as a geometric average of all predictions, weighted by the relative likelihood of a person of each caste having that name versus the share of people across all Dalit castes with that name. Both in the simulation of these predicted castes, as well as in my empirical results, I show robustness to using three alternative aggregation methods: a simple average, an average weighted by the over-representation of each name-part within that caste, and the maximum operator. Appendix Section A contains a more detailed explanation of this procedure as well as an illustrating example.

$$\text{Inaccuracy} : \frac{1}{N} \sum_{i=1}^N P(\text{caste}_i = C | \text{names}_i, \text{true caste}_i \neq C) \quad (2)$$

Expression (1) can be interpreted as the efficiency of the predictions, that is, the probability that this procedure assigns to the correct caste. The metric in equation (2) is a measure of the Type I error (or inaccuracy) rate, i.e., the probability that I assign to a given incorrect caste. Ideally, we should see that $P(\text{caste}_i = C | \text{names}_i, \text{true caste}_i = C) = 1$ and $P(\text{caste}_i = C | \text{names}_i, \text{true caste}_i \neq C) = 0$. I vary two main parameters: whether or not I use all name-parts or only surnames, and the share of the original dataset assigned to train the algorithm.¹¹¹² The first variation is meant to test the out-of-sample predictive power of all names versus surnames. The second variation is an attempt to simulate the results when the training dataset is much smaller than the data where the researcher will predict the caste of individuals. Under each combination of parameters, I repeat the process a total of 100 times and report the average efficiency and inaccuracy rates across all simulations.

4.1.1 Simulation results

The first result from these simulations is that the out-of-sample efficiency of caste predictions is greater when using only last names in the prediction algorithm. The first two vertical bars in Figure 2 present the average probability that the algorithm assigns to the true caste of a person based on either all their name-parts (shown in green) or only their last name (shown in blue).¹³ From these two columns, we can see that the out-of-sample efficiency almost doubles when using only last names, as opposed to including first and middle names. While it is true that in-sample, the predictive power of the algorithm should be weakly better by using more information (i.e., more names), this issue arises as using first and middle names (which are more indicative of class than caste) introduce noise in the predictions and leads them to perform poorly outside of the training dataset. This finding lends credibility to the algorithm, as the qualitative literature argues that it is last names that are culturally linked to caste membership and not first or middle names (Singh, 1996). Complementing this finding, the right-most bars present the same comparison (using all name-parts or last names only), but predicting instead whether a person is Hindu or Muslim as opposed to their caste. Here, we see that there is little distinction in the procedure’s efficiency when predicting a person’s religion across the two sets of name types. The religious origin of names is more clearly tied to language and sacred texts than caste-based names. For this reason, both given names—often based on the names of deities or prophets—and last names are indicative of a person’s religion, leading to similar out-of-sample predictions when using either of these sets of names.¹⁴

Even conditional on each caste, the efficiency of the predictions is high. In light blue, Figure 3 presents the results shown in Figure 2 disaggregated by caste, with the predictions constructed using

¹¹I use either 50%, 25%, or 2.5% of the full dataset to train the algorithm and the respective complement to test its predictions.

¹²I also vary the method used to aggregate across all the individual’s name-parts and report those results in

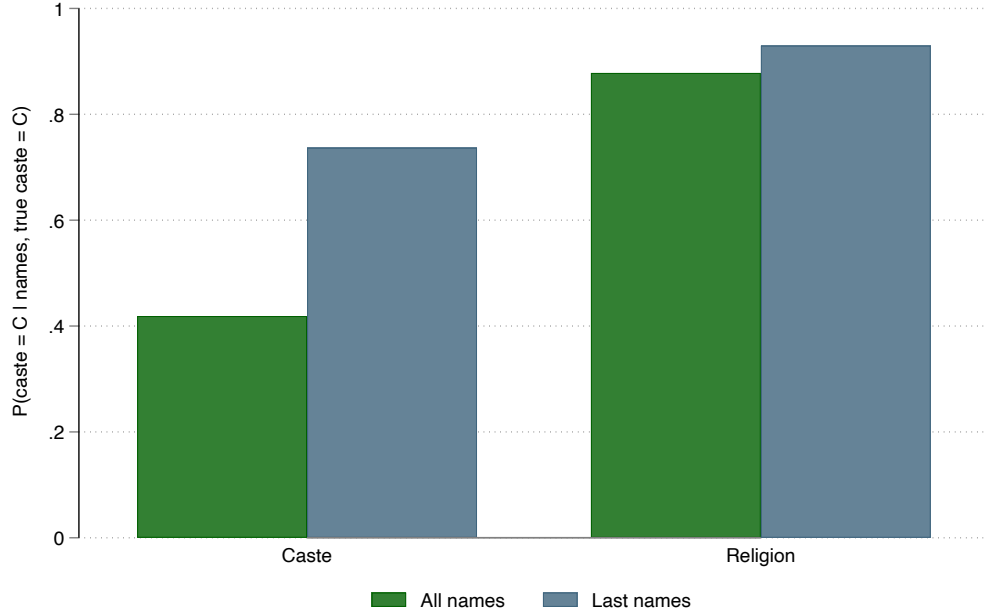


Figure 2: Efficiency of caste and religion predictions using all names vs last names only

only last names. The horizontal axis lists each caste ordered by their relative size in the population (represented in dark blue). Jatavs, representing almost 60% of Dalits in Uttar Pradesh, are most efficiently identified by the caste prediction. Within the group of people known to be Jatav in the testing subset, the algorithm estimates, on average, a higher than 80% likelihood that they belong to that same caste. Overall, the efficiency rate drops as the size of the caste group decreases. Even for the smaller groups, however, the average probability of correctly predicting their caste is well above 50%. The only exception is the “Other SC” group. This is not surprising, since this group is composed by an amalgam of the smallest sub-castes that do not fall within one of the other nine groups in the figure. The Type I error rate exhibits the opposite pattern. In Figure 4, we see that the algorithm over-predicts the likelihood that people belong to the larger caste groups when that is not their correct caste.

Finally, the simulations illustrate that the efficiency of predictions decreases as the size of the training dataset becomes smaller, relative to the size of the testing dataset. Using only 2.5% of the entire dataset, however, I am able to predict whether or not the individuals in the testing subset are Jatav or belong to any other Dalit caste. Figure 5 illustrates both of these results. Even though the efficiency of the predictions is lower when using only the smallest training set, it remains above

Appendix Section B.

¹³These results are averages taken across all castes.

¹⁴I include the complementary results for the Type I error rate of these prediction in Figure 11 in Appendix Section B.

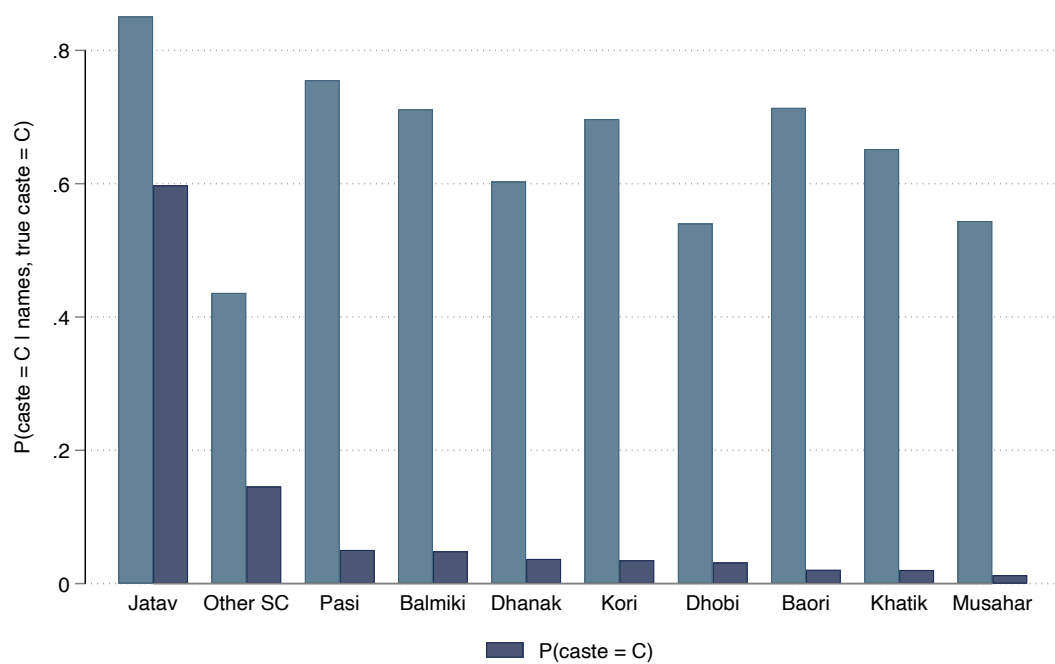


Figure 3: Efficiency of predictions for each caste

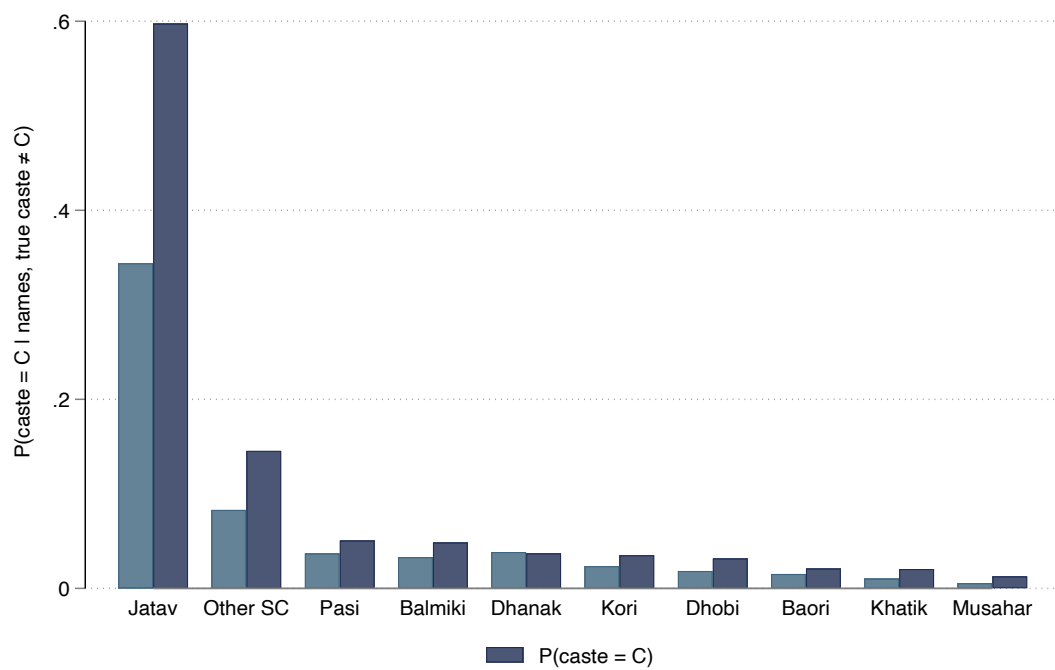


Figure 4: Inaccuracy of predictions for each caste

70% for both Jatavs and non-Jatavs. Given the right-wing discourse of Jatavs being the group most advantaged by political reservations, this result is particularly important as I am able to efficiently categorize workers into one of these two categories.

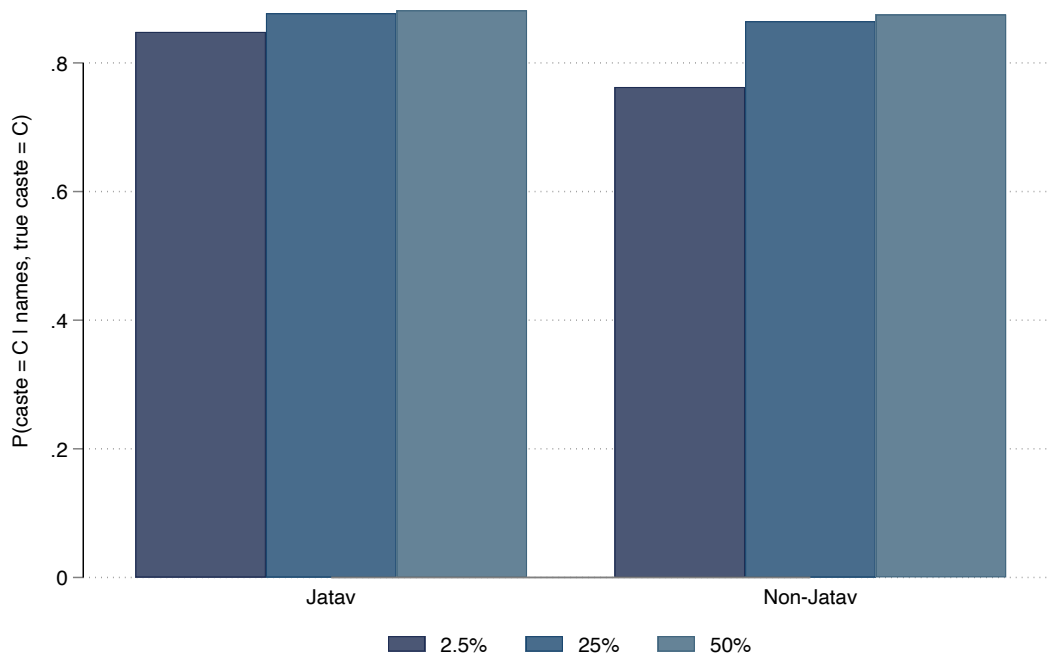


Figure 5: Efficiency of predictions for Jatavs and non-Jatavs, varying the size of the training dataset

To place these results in context, I compare them to the efficiency and accuracy rates within the literature that seeks to match people across historical censuses. This literature serves as a useful benchmark, since it is subject to similar sources of measurement error, such as spelling mistakes or strategic name choices.¹⁵ Within this parallel body of work, Abramitzky et al. (2021) summarizes a series of procedures and lists estimates of their efficiency and accuracy. The average efficiency of my caste predictions is 73.6%. In comparison, the maximum efficiency of the methods surveyed in Abramitzky et al. (2021) lies around 60%–70%. The accuracy of my predictions, that is, one minus the Type I error rate that I present in my results, is also reasonable in relation to the census linking methods. As Figure 4 shows, the accuracy of my caste estimates is 95% on average, ranging from slightly under 70% (for the Jatav caste) to over 95% for most of the smaller caste groups.¹⁶

¹⁵The latter occurs in the U.S. context, for instance, when black people during the Great Migration chose to identify as white in an attempt to face less discrimination (Dahis et al., 2020).

¹⁶Note that Figure 4 reports the *inaccuracy* of my estimates, i.e., 1 – the accuracy rate.

4.2 Making predictions in the NREGA dataset

Having calculated the empirical probabilities for each caste–name pair, I then predict the caste of workers in the NREGA data. As mentioned above, the data provides the name of each individual, the name of their father (or husband in the case of married women), and the unique ID of their household. I use this information to create a single prediction for each caste at the household level. This is done to maximize the number of name signals for each worker. Since caste is hereditary and highly endogamous, I argue that this step is done without loss of generality, while it does improve my ability to predict the caste of many laborers in the data. The steps that I follow to create a unique prediction for each caste and NREGA laborer is outlined below:

1. Restricting attention to only surnames, I assign to each worker’s name-part the respective predicted probability for each caste conditional on that surname. Names that are in the NREGA data but not in the REDS data are left as missing.¹⁷
2. Next, I aggregate the conditional caste probabilities across all surnames associated with all individuals in a household. For this step, I use a formula based on Bayes’ rule that begins with a caste prior equal to the share of all Dalits in the REDS sample that belong to each caste. ¹⁸This procedure can be interpreted as a geometric average of the conditional caste probabilities for each name-part where each signal is weighted by the ratio of the name’s likelihood within the specific caste group versus the likelihood in the overall population: $P(\text{caste} = C, \text{name})/P(\text{name})$.
3. The result, up to this point, is a continuous measure of $P(\text{caste} = C | \text{surname}^1, \dots, \text{surname}^n)$ for each of the ten castes (C). These measures can be interpreted as the likelihood that a bureaucrat with only the name and sex of an NREGA worker assigns to them belonging to a caste. In order to be able to estimate the effect of the reservation policy on the sub-sample of individuals belonging to a caste, I finish this procedure by creating an indicator of whether a worker belongs to a given caste, as shown below:

$$1\{\text{caste} = C\} = 1\{P(\text{caste} = C | \text{surname}^1, \dots, \text{surname}^n) > P(\text{caste} = C)\}$$

This indicator variable is equal to one if the probability of them becoming to a caste conditional on their surname is greater than the likelihood of sampling someone from that caste at random from the state’s Dalit population. The intuition behind this metric is to assign people to a

¹⁷Later on, when the caste probabilities are aggregated across name-parts and individuals in a household, I ignore these missing values. If all name-parts of all individuals in the household are missing, then I do not create any caste prediction for this individual.

¹⁸The posterior from this initial operation becomes the prior for incorporating the signal from the second name-part. I follow this procedure iteratively until I have incorporated all name-parts for that household.

caste for whom the algorithm assigns a greater-than-random chance of being members of that particular caste.

5 Empirical strategy

In 2008, the Indian government assigned, within each district of Uttar Pradesh, certain constituencies to be contested by only Dalit candidates and the remainder to be contested by an unrestricted pool of politicians. In principle, Dalit candidates can also participate and win election in the unrestricted (or “general”) constituencies. In practice, there are no general constituencies won by a Dalit candidate in Uttar Pradesh during the period of my study. By district, the government made two determinations: how many and which ACs were to be reserved. The number of reserved ACs in a district was chosen so that the proportion of reserved constituencies equaled (or was as close as possible to) the share of inhabitants belonging to a Dalit caste. As a note, I use the abbreviation “SC” to refer to Dalits in this table and in all other algebraic expressions in order to economize on space. These two terms are completely interchangeable. In Table 1, I present the example of Unnao district in Uttar Pradesh. The district has a total of six ACs and 30% of its total population is comprised by Dalits. Therefore, two ACs were reserved for Dalit candidates in Unnao. Next, the specific constituencies restricted for Dalit representatives were chosen to be those with the highest proportional representation of Dalit inhabitants in the district.

Table 1: Reservation of ACs in Unnao district

AC Name	SC%	Reserved
Mohan	38.8%	SC
Safipur	36.6%	SC
Purwa	31.8%	Gen
Bhagwantnagar	29.7%	Gen
Bangermau	26.6%	Gen
Unnao	22.0%	Gen

Note: This table lists each of the constituencies in the sample district of Unnao in the state of Uttar Pradesh. The constituencies are ranked in descending order by the proportion of residents who are Dalit. As a reminder, Dalits are officially termed Scheduled Castes by the Indian government, hence the SC acronym. Just over 30% of the district’s population is composed of Dalits, which results in two out of the total six ACs being reserved.

Given the long legacy of discrimination against and ostracism of Dalits in India, it is not difficult to understand that, on average, individuals belonging to any of these castes tend to be poorer, less educated, and have lower access to public goods, on average (Munshi, 2019). For this reason, a simple OLS regression comparing any reserved and general constituencies (e.g., Mohan vs. Bangermau in Table 1) would not yield a consistent estimate of the effect of being represented by a Dalit

state legislator. In the example above, however, Safipur was narrowly reserved for Dalit candidates and Purwa fell narrowly short of getting the same status. This motivates the use of a regression discontinuity design. Restricting our attention to a small neighborhood of the cutoff yields two groups of constituencies that are, on average, equal, save for the fact that one set is assigned a Dalit representative and the remainder is assigned a non-Dalit representative. By choosing the constituencies with the highest representation of Dalit inhabitants, the government effectively created a population cutoff within each district.¹⁹ In the case of Unnao district, I define its cutoff (c_d) as the average of 36.6% and 31.8%, the shares of Dalits in Safipur and Purwa, respectively.²⁰ Based on this cutoff, ACs fall into a reserved (or general) status if they are located above (or below) the cutoff.²¹ In my analysis, I exploit this quasi-random assignment of ACs to be represented solely by Dalit MLAs in order to exogenously vary the caste category of the state representative. This allows me to study the effect on Dalit constituents of having an MLA of their same caste category and whether this form of political representation changes the distribution of public goods across castes. Lastly, in order to pool observations across districts—each having its own cutoff—I subtract from the share of Dalits in each constituency the respective district’s cutoff, as shown in equation (3). Hence, this re-centered measure ($\widetilde{SC}\%_{ad}$) is positive whenever an AC falls above the cutoff and negative otherwise. This creates a running variable where the cutoff for every district is now equal to zero, which simplifies the analysis and its exposition.

$$\widetilde{SC}\%_{ad} := SC\%_{ad} - c_d \quad (3)$$

Consider Y as the outcome of interest: the delay of payments to an individual Dalit worker. The potential outcomes $\{Y(0), Y(1)\}$ represent the outcome that is realized if the unit lies in an AC is “general” or reserved, respectively. I estimate the effect (at the cutoff) of having a Dalit MLA on the quality of public service delivery (Y). This parameter can be written as:

$$\beta := E \left[Y(1) - Y(0) | \widetilde{SC}\% = 0, SC \right] = \lim_{\widetilde{SC}\% \downarrow 0} E[Y|SC] - \lim_{\widetilde{SC}\% \uparrow 0} E[Y|SC]$$

Where the second equality follows from the quasi-random assignment of constituencies to reserved or general status around the cutoff. Given that there are no “general” constituencies won by

¹⁹Jensenius (2015) explores a matching design that is similar in spirit to estimate the average treatment effect of being in a reserved constituency, but focusing on the period of 1970 to 2001 (before the re-drawing of constituency boundaries). In a working paper, Sarkar (2019) also employs this implicit population cutoff to explore the impacts of reservations. Neither of these studies, however, studies how the effects of being represented by a Dalit state legislator varies with respect to the caste of the constituents.

²⁰The cutoff could be defined to be any number between these two shares. My results are robust to choosing different bandwidths for the RDD estimator, which provides evidence that the choice of a specific cutoff value does not significantly impact my analysis. In general, for a district d where r constituencies are chosen to be reserved, consider $SC\%_d^{(i)}$ the AC with the i^{th} highest share of SC residents. Then, the generic formula for the district-specific cutoff is given by $c_d = \frac{1}{2} (SC\%_d^r + SC\%_d^{r+1})$.

²¹Figure 15 in the Appendix shows that the assignment of constituencies complied with this rule perfectly.

Dalit candidates in my data, β can be estimated as the average difference between the provision of timely payments to NREGA workers in constituencies represented by a low-caste candidate versus those in which the MLA is unrestricted. Lastly, notice that the expression above is conditional on the workers belonging to a Dalit (SC) caste, as I am only interested in estimating the effect on Dalit workers. I estimate β using a linear regression of the form:

$$Y_{ivt} = \beta^{SC} \text{Dalit MLA}_{ad} + f(\widetilde{SC\%})_{ad} + \mathbf{X}'_{ad}\Gamma + \phi_d + \lambda_t + \varepsilon_{ivt} \quad (4)$$

Here Dalit MLA_{ad} is an indicator equal to 1 if observation i is located in an AC a that is reserved for only Dalit candidates and 0 otherwise. The function f controls flexibly for $\widetilde{SC\%}$ at either side of the cutoff. Conditional on f and the district fixed effects ϕ_d , β estimates the difference between the average outcome for observations in constituencies just above the cutoff versus those in ACs just below it, within each district. The district fixed effects are necessary in this instance as the Delimitation Commission quasi-randomly assigned ACs to either side of the cutoff within each district. The year fixed effects (λ_t) and the constituency-level controls (\mathbf{X}_{ad}) are meant to increase the precision of β 's estimate by absorbing extraneous variation in the outcome variables. I estimate this regression using only observations in ACs where the share of Dalits in the population falls within 2.5 percentage points of the district's cutoff. This bandwidth is constant across specifications to make the results comparable. I specify f to be linear function with flexible slope at either side of the cutoff and weigh the observations using a triangular kernel.²² All results, unless otherwise stated, are robust to using the optimal bandwidth and the non-parametric estimate of f (Calonico et al., 2014). Finally, since the identity of the state legislator varies at the AC level, I cluster all standard errors by constituency.

The main identifying assumption in the RD design is that the potential outcomes $\{Y(0), Y(1)\}$ are continuous around the cutoff $\widetilde{SC} = 0$. Intuitively, this requires that there be no manipulation of constituencies to either side of the cutoff and that this deterministic rule is not used in granting access to another policy. As Iyer and Reddy (2013) finds, the re-districting process was centrally controlled and there is little evidence of constituencies being manipulated to be either reserved or general. In addition, the central government independently corroborated the statistics reported by the State leaders using the 2001 Census data. Furthermore, the purpose of this re-districting was solely to define new administrative boundaries and reserve certain constituencies for disadvantaged groups, not to allocate any other public service or goods using this same decision rule. While the continuity of the potential outcomes cannot be tested directly— $Y(0)$ and $Y(1)$ are fundamentally unobservable jointly for any worker or constituency—I present some evidence consistent with this assumption. First, Figure 6 shows that the density of \widetilde{SC} is continuous around zero. A sudden jump in the probability of falling below or above the threshold would indicate that certain constituencies

²²This is the default kernel used in the `rdrobust` Stata package by Calonico et al. (2014). Results are robust to using either a uniform or Epanechnikov kernel.

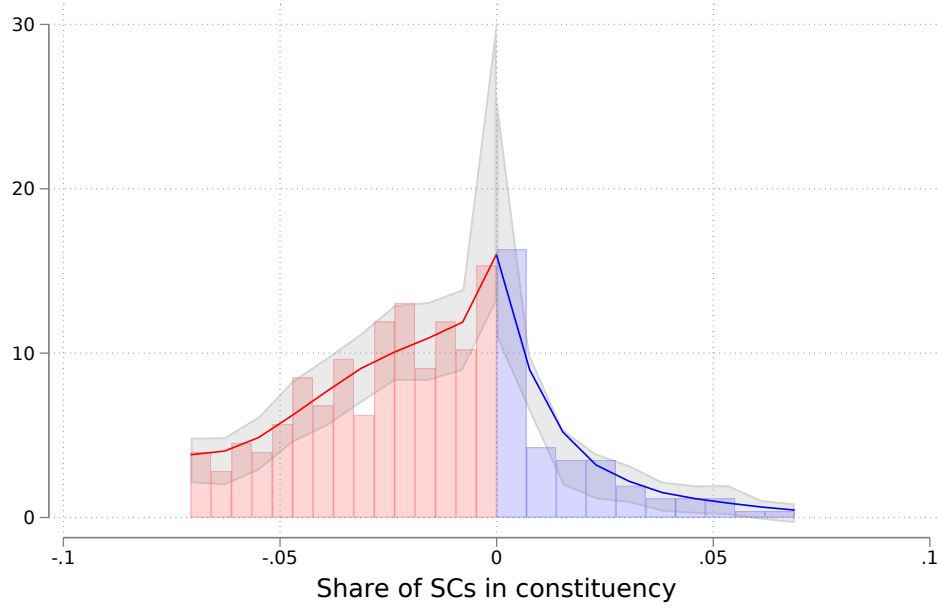


Figure 6: Probability density function for the normalized share of SCs in each constituency

Table 2: Balance across baseline covariates

Variable	N	General	Reserved	Difference	Std. Error
<i>Panel A: Socioeconomic indicators</i>					
Total population (1971) (1000s)	372	208.21	208.30	0.10	(1.37)
Dalit share of population (1971)	372	20.67	19.05	-1.62**	(0.81)
Literacy rate (1971)	372	20.61	20.10	-0.51	(0.96)
Dalit literacy rate (1971)	372	9.87	9.77	-0.10	(0.61)
Literacy rate (2001)	372	44.99	44.42	-0.57	(1.18)
Dalit literacy rate (2001)	372	36.68	36.29	-0.39	(1.06)
Main workers % (2001)	372	23.63	23.80	0.18	(0.33)
Dalit main workers % (2001)	372	22.53	22.82	0.28	(0.43)
Agricultural labor (2001)	372	8.48	7.52	-0.96	(0.61)
Dalit agricultural labor % (2001)	372	42.45	42.73	0.28	(2.48)
<i>Panel B: Public good and services provision</i>					
Electricity access % (2001)	372	82.47	82.94	0.47	(2.15)
Dalit electricity access % (2001)	372	83.16	82.79	-0.37	(2.19)
Educaction access % (2001)	372	88.76	86.92	-1.83*	(1.01)
Dalit educaction access % (2001)	372	89.83	88.11	-1.72*	(0.99)
Medical access % (2001)	372	44.92	43.96	-0.96	(2.30)
Dalit medical access % (2001)	372	43.80	41.75	-2.04	(2.46)
Communication access % (2001)	372	55.48	51.62	-3.86	(2.67)
Dalit communication access % (2001)	372	54.18	49.55	-4.64*	(2.74)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

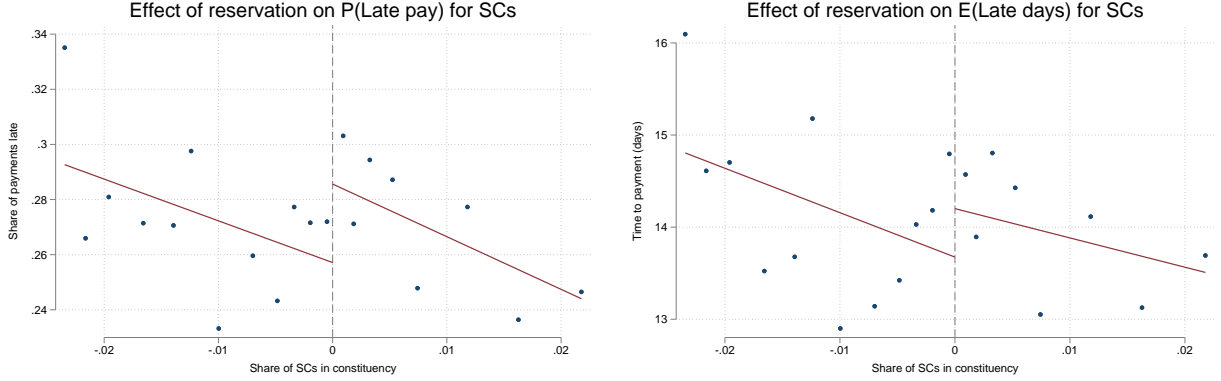
might have been artificially placed on either side of the cutoff. Figure 6 does exhibit a sudden drop in the mass of constituencies slightly over the cutoff. Nevertheless, this is mechanically the result of there being proportionately far fewer reserved versus general constituencies within any single district. In particular, it is often the case that only one or two ACs are reserved in a given district, which results in most of the mass above zero being very close to the cutoff. Second, Table 2 presents the results of several individual regressions (at the AC level) testing whether constituencies in a small neighborhood of the threshold are balanced across a series of baseline covariates. These variables present a complete picture of the socioeconomic condition (Panel A) of each constituency and the provision of public goods and services (Panel B), both for Dalits and the general population. These estimates show that there is no systematic difference in the socioeconomic status of constituencies or the provision of public goods on either side of the cutoff. Potentially, the only concerning result that is both economically and statistically significant is the share of Dalits in each constituency as measured in 1971. A priori, we could be concerned that Dalits in some reserved constituencies after 2008 might be better off, relative to Dalits in “general” constituencies due to being *previously* assigned a Dalit state legislator. This, however, would result in a positive difference across the two types of constituencies, which is the opposite of what I find. Similarly, if anything, it seems that reserved constituencies experienced a slight under-access to education (for both Dalits and the general population) and communication services (particularly for Dalits). These results would work in the opposite direction of my hypothesized effect of Dalit representation on public good provision for people of the same caste category. Furthermore, none of these statistically significant differences are robust to using slightly different bandwidths. In general, Table 2 supports the fact that, at baseline, reserved and “general” constituencies are comparable, on average.

6 Results

6.1 Timely payment to NREGA workers

Whereas Dalits make up 20.7% of the total population, they comprise over a third of active workers in NREGA. Given the over-representation of Dalits in this public workfare program, we might expect that Dalit MLAs use their influence over block bureaucrats to advantage workers of their same caste category. Moreover, the prompt payment of wages is a highly salient issue among NREGA participants, especially for Dalit workers. For constituents to receive their wages on time, would bode well for MLAs as this could improve the public’s perception of their efficacy in aiding their constituents. In this section, I begin by comparing Dalit workers in constituencies that are narrowly reserved for a Dalit candidate versus those that live in a constituency with an unrestricted candidate pool. As stated above, the winning candidate in the “general” constituencies is never a person who belongs to a Dalit caste, so this comparison can be interpreted as how being represented by an MLA of their same caste category affects Dalit workers. A priori, we might expect to find

that Dalit MLAs either improve (lower) payment times for workers of their own caste group or that they do not affect this outcome whatsoever. If candidates can perfectly commit to enact the policies preferred by their parties, the caste identity of the representative should not affect their behavior in office, as discussed in Pande (2003). However, absent this commitment, we would expect a weak increase in the public benefits targeted to Dalit voters. This would imply that $\beta \leq 0$ in equation (4). Meaning that Dalit workers in constituencies headed by a Dalit MLA receive their payments sooner or just as quickly relative to Dalit workers in a constituency represented by a non-Dalit MLA.



(a) RDD effect on the share of payments received late (b) RDD effect on the average time to payment

Figure 7: Discontinuities in payment delays at the reservation cutoff

Note: This figure uses binned scatterplots to visually represent the discontinuity in both the share of late payments and the average delay of payments (measured in days) at the population cutoff that determines whether a constituency is assigned a Dalit (SC) state legislator or a non-Dalit state legislator; observations to the right of the cutoff are in reserved constituencies and those to the left of the cutoff are in “general” constituencies.

I obtain the parameter values in Table 3 by estimating the regression in equation (4). This regression uses only Dalit workers in ACs that lie within 2.5 percentage points of their districts eligibility threshold. This comparison holds constant the caste category of the workers while varying the caste category of the MLA. The standard errors are clustered by constituency to reflect the fact that the identity of the MLA is assigned at this same level. The regressions include year and district fixed effects, as well as a set of constituency-level variables controlling for various measures of public good provision and socioeconomic status at baseline to improve the precision of my estimates. In this table, we see that Dalit representatives cause an economically large increase in the delay of payments to workers of their same caste group. In the first column, I estimate that NREGA participants under an SC representative are 3.3 percentage points more likely to receive late payments (defined as over 15 days after completion of the work). This represents a 12.4% increase relative to the average of probability of being paid late for workers in “general” ACs. While this point estimate is only marginally significant, using a local linear approach to control for the share of Dalits in each constituency—represented by $f(\widehat{SC})\%$ in equation (4)—results in a far more precise estimate of the

Table 3: Effects of Dalit representation on the payments to Dalit workers

	P(Late pay)	E(Late days)
Dalit MLA	0.033* (0.018)	0.692 (0.568)
Observations	195256	195256
Avg. Y	0.268	14.05
Std.Dev. Y	0.356	12.74
P-value	0.0670	0.223
Bandwidth	0.025	0.025
Yr., Dist. FEs	Yes	Yes
Sample	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses, p-values are also constructed using the clustered standard errors. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).

same magnitude, thus lending more credibility to this finding.²³ Taking this result at face value, it would imply that around 64,000 Dalit workers are receiving their payments after the required deadline. For people who select into working for NREGA (meaning this might be their best work opportunity), having to wait even an additional day without receiving their due compensation could have highly adverse effects on their well-being.

The second column suggests that this increase in the time to payment is concentrated along the 15 day margin that defines payments as being late and not necessarily reflects an overall shift in the distribution of the time it takes for the workers to get paid. This is evidenced both in the lack of statistical significance (with a p-value of more than 20%) and in the magnitude of the coefficient itself, indicating only a 4.9% increase in the number of the time to payment—relative to the 12% effect shown above for the likelihood of the payment being delayed. The results in Table 3 for the share of late payments and average time to payment are presented graphically in Figures 7a and 7b, respectively. These figures demonstrate that the patterns in the previous table are brought to bear in the raw data, as well, and that they do not seem to be driven by outlying sets of observations.

This deleterious effect of Dalit representation on workers of that same caste category does not conform to the prediction of the model laid out in Pande (2003)—which predicted a weak improvement for Dalit workers. That being said, it is important to remember that the ability of MLAs to affect policy outcomes hinges on their control over the block bureaucrats. A possible

²³I use the procedure (and Stata package) introduced by Calonico et al. (2014) to construct these alternative estimates.

explanation for this finding is that BDOs believe that Dalit MLAs would be less able to influence their career trajectory, relative to non-Dalit MLAs. MLAs belonging to a lower caste might have lower bargaining power within the political apparatus and be less effective at securing the promotions or re-allocations desired by the BDO they are trying to persuade. More generally, the average Indian bureaucrat tends to be positively selected in terms of income and caste, which creates a power imbalance vis-à-vis the lower caste politicians. Previous work has shown that ethnic heterogeneity within an organizational structure can lead to lower productivity (see for instance, Hjort (2014) in the context of Kenyan firms). The results in Section 7.1, support the argument that the balance of power between MLAs and BDOs is an important mediator in the representative’s ability to influence policy outcomes.

6.2 Heterogeneity by caste in the effects of Dalit reservations

Having shown a weakly positive increase in the wait times that Dalit workers face before they are paid, a natural question is whether this detrimental impact is borne uniformly by all Dalit workers or whether the average treatment effects are masking important heterogeneity across individuals. Specifically, I investigate how these effects vary across individual castes within the Dalit category. Motivated by the media and right-wing political discourse in Uttar Pradesh, I first test how the effect of being in a reserved constituency differs for members of the largest Dalit caste (Jatavs) relative to non-Jatav Dalits. To test this prediction, I modify the regression in equation (4) by including an interaction term between the indicator for a constituency being reserved for Dalit legislators and an dummy variable equal to one if I predict the worker to be Jatav, using the procedure outlined in Section 4.²⁴ The resulting equation takes the following form:

$$Y_{ivt} = \beta_1 \text{Dalit MLA}_{ad} + \beta_2 \text{Dalit MLA}_{ad} \times \text{Jatav}_{ivt} + \delta \text{Jatav}_{ivt} + f(\widetilde{SC})_{ad} + \mathbf{X}'_{ad}\Gamma + \phi_d \lambda_t + \varepsilon_{ivt} \quad (5)$$

Following the same sample restrictions I laid out for the regression in equation (4), I estimate this expression using only Dalit workers. The coefficient β_1 in the equation above estimates the effect on non-Jatav Dalits of having an MLA of their same caste group. β_2 , in turn, estimates the difference in this effect for Jatav Dalits. If the assertion that Jatavs have been most advantaged by caste-based reservations is true in the context of punctual payments to NREGA workers, we expect that $\beta_1 > 0$ and $\beta_2 < 0$. The first sign restriction would indicate an increase in payment times for non-Jatav Dalits, as we found for all Dalits in Table 3. In contrast, β_2 being strictly negative would imply that these deleterious effects are attenuated for members of the largest Dalit caste in the state—and even reversed if $|\beta_2| > |\beta_1|$.

²⁴Specifically, the indicator variable for a person being Jatav takes a value equal to one if the predicted probability of them belonging to this caste is greater than the average share of Jatavs in the population. The intuition behind this definition is that people for whom I predict a greater-than-random likelihood of belonging to this caste are defined as Jatav.

Beyond exploring heterogeneity by caste, this next set of results reflects whether being represented by a Dalit state legislator has an effect on inequality across caste groups. Work by political scientists in India has identified caste inequality within Dalits as an important driver of support for populist parties, such as Narendra Modi’s BJP (Jaffrelot, 2022). The overwhelming increase in the support for the BJP has had significant policy implications in India. Given these ramifications, it is important to understand how caste-based political reservations influence inequality across castes.

Table 4: Heterogeneity in the effect of reservations for Jatavs vs. non-Jatavs

	P(Late pay)	E(Late days)
Dalit MLA	0.0353* (0.0186)	0.790 (0.643)
Dalit MLA \times Jatav	-0.00763 (0.0102)	-0.666** (0.331)
Observations	101275	101275
Avg. Y	0.265	13.94
Bandwidth	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes
Sample	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).

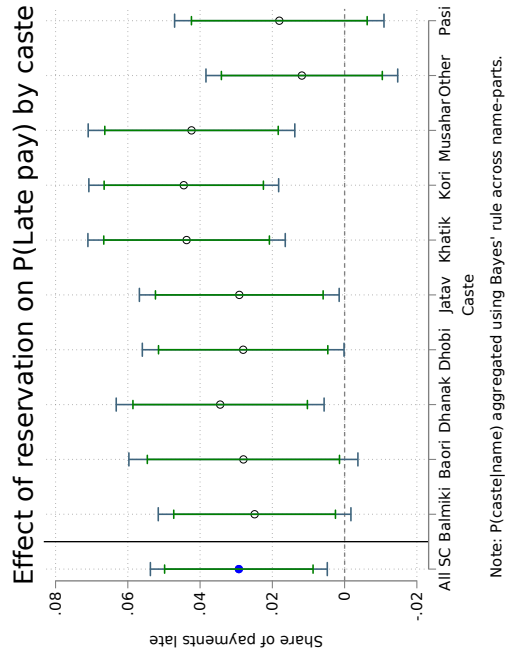
Table 4 shows the estimates from equation (5) for both the probability of receiving a payment late (column one) and the average days until the workers are paid (column two).²⁵ In column one, I estimate that both non-Jatavs experience a 3.5 percentage point increase in the likelihood of receiving a late payment due to having a Dalit representative. Based on the coefficient in the second row of this same column, I cannot reject that the effect is statistically different for Jatav workers. The interaction term in column two does indicate that the adverse effect of having a Dalit MLA on the average time delay of payments is statistically smaller for Jatavs. Column two suggests a negligible increase in the number of days until payment—only 0.124 days or 0.88% of the mean in the unreserved group—for Jatav workers. In contrast, the estimate in the first row suggest a 0.79 day increase in payment time for non-Jatav Dalits. Nevertheless, we cannot reject that the effect is statistically different from zero for non-Jatavs.

²⁵ A significant proportion (48%) of workers do not report their surname in the NREGA data. Since I use only the last names for individuals in the household to predict the worker’s caste, I am unable to create a prediction for this subset of laborers. This lead to an important reduction in the sample size relative to the results shown in Table 3

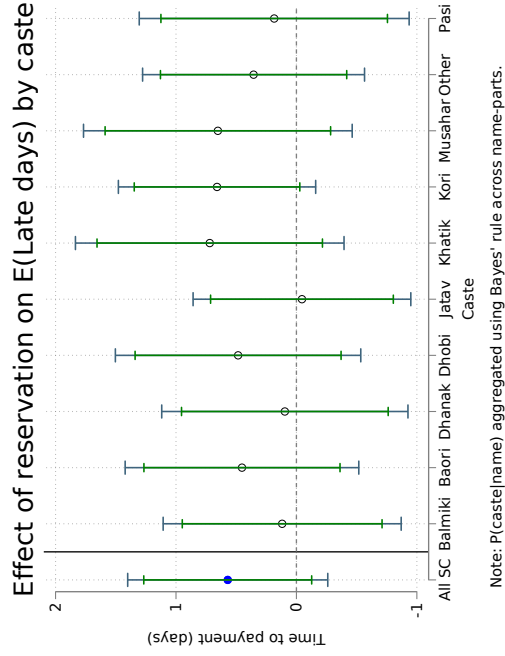
Under the premise that there is substantial heterogeneity in the effect of Dalit reservations by caste, one explanation for why the baseline coefficient in the first row is so noisily estimated is the fact that it is averaging the effect across all non-Jatav lower castes. Next, I study how the effect of having a Dalit state legislator varies beyond the dichotomous comparison of Jatavs versus the rest of Dalit workers. To do so, I plot the estimates for β in equation (4) restricting the sample to only people from each of the ten castes that I am able to predict. Figures 8a and 8b display the results of this exercise. The horizontal axes denote the caste of the workers. Each point represents the average treatment effect—at the threshold—of having a Dalit MLA on either measure of late payments for NREGA workers of each respective caste. In Figure 8a, the ATE for each individual caste lies within the 95% confidence interval of the estimated effect for all Dalit workers (shown with a blue filled dot). This implies that, statistically, there is no systematic heterogeneity in how each caste is affected by having a low-caste MLA. Notably, the majority of these caste-specific estimates are statistically different from zero at the 95% level—and all but the estimates for Pasis and the “catch-all” category (“Other”) are different from zero with 90% confidence. Hence, the lack of heterogeneity across castes is unlikely to be generated by imprecisely estimated effects for each caste, rather could be indicative of reservations having no impact on within-Dalit caste inequality. Figure 8b, echoes the results shown in Table 4. Taking the magnitude of the estimated ATEs, Jatavs are least affected by the average increase in time to payment. However, the estimates for each caste (including Jatavs) fall within the 95% confidence interval of the ATE for all Dalit workers. This reiterates the result found in the previous figure and supports the lack of significant heterogeneity in the effect of Dalit reservations across Dalit castes.

7 Mechanisms

The results above suggest that Dalit MLAs are unable to improve the delivery of NREGA payments to constituents of their same caste category and might even perform worse in this regard, relative to non-Dalit state legislators. This effect appears uniform across each Dalit caste. In this section, I explore two mechanisms that mediate these results. The first mechanism pertains to the balance of power between MLAs and BDOs. It asks whether Dalit MLAs perform better in areas where they have higher bargaining power relative to the bureaucrats. The second channel explores the incentives of the MLA as dictated by their political party. In this last part of the analysis, I ask whether Dalit MLAs belonging to parties which explicitly support the interests of low caste voters perform differently than those belonging to other political institutions.



(a) Effects by caste on the share of late payments



(b) Effects by caste on the average time to payment

Figure 8: Caste-specific average treatment effects of having a Dalit state legislator

Note: Each figure represents the average treatment effect, at the reservation cutoff, of having a Dalit state legislator on a measure of late payments to NREGA workers. Figure 8a uses the share of late payments as the outcome while Figure 8b uses the average number of days until payment as the dependent variable. The horizontal axes denote the caste composition of the sample, ordered alphabetically: either all Dalit (SC) workers (the leftmost point represented by a filled circle) or each of the individual caste groups (represented by empty circles). The dark blue bars mark the width of the 95% confidence interval around each ATE, while the green bars mark the 90% confidence intervals. The caste predictions were aggregated across name-parts and individuals to the household level using the procedure based on Baye's rule.

7.1 MLA bargaining power in their relationship with BDOs

The control that MLAs have over the timing of payments for NREGA workers relies crucially on the influence that they have over block officers (BDOs) within their constituency. As described in Section 2, BDOs can be incentivized to heed the MLAs' wishes as the latter play an important role in the promotion and re-allocation of the bureaucrats. The administrative geography in India, provides some useful variation in the degree of control that MLAs have over BDOs. While blocks are generally smaller than ACs, it is often the case that blocks span more than one constituency. When this happens, each MLA must compete with other legislators as they try to influence the BDO. Intuitively, this reduces the bargaining power that each of the MLAs has over the bureaucrat whose block spans their constituencies. Gulzar and Pasquale (2017) uses this feature of the relationship between politicians and bureaucrats to vary the balance of power among state legislators and BDOs.

In the analysis below, I will compare the effect of having a Dalit MLA separately for low-caste workers in blocks that are contained within a single constituency and those in blocks that span more than one constituency. Given the discussion above, I would expect that Dalit MLAs are better able to pressure BDOs into paying their constituents more promptly in areas where they are the sole state representative responsible for that bureaucrat.

Table 5: Low-caste MLA performance and their relationship with BDOs

	P(Late pay)		E(Late days)	
	Single MLA	Mult. MLAs	Single MLA	Mult. MLAs
Dalit MLA	0.00511 (0.0242)	0.0562*** (0.0217)	0.576 (0.765)	0.855 (0.746)
Observations	92677	102579	92677	102579
Avg.Y	0.259	0.277	13.50	14.59
Std.Dev.Y	0.352	0.360	12.10	13.30
Bandwidth	0.0250	0.0250	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes	Yes	Yes
Sample	Dalit Only	Dalit Only	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).

Table 5 shows the results of this exercise. Columns one and three estimate the effect of having a Dalit MLA on Dalit workers overseen by BDOs who answer to a single state legislator (higher bargaining power for the MLAs). Columns two and four estimate this same effect but for Dalit workers living in blocks that span multiple constituencies (lower bargaining power for the MLAs).

Comparing columns one and two, I show that the increase in the share of late payments—as reflected in Table 3—occurs only within blocks that span multiple constituencies. That is, it is only in blocks where Dalit MLAs have lower bargaining power in relation to the block bureaucrat that low-caste workers are more likely to receive their payments late, relative to Dalit workers represented by a non-Dalit MLA. This distinction is statistically significant, evidenced by the fact that each estimate lies outside the confidence interval of its counterpart. Columns three and four echo the same qualitative pattern. Nevertheless, I am not able to statistically distinguish the treatment effect across the two samples.

7.2 Caste ideology of the MLA’s party

Even if a Dalit MLA could, in principle, coerce the BDOs, the representative could belong to a political party that does not explicitly favor the interests of low-caste voters. Even within some parties that are not explicitly opposed to redistribution towards Dalit workers, politicians might believe that by placing a higher weight on Dalit welfare they might alienate non-Dalit voters—who almost always represent the majority of the electorate. It stands to reason that the ideology of the MLA’s party should play an important role in determining whether or not Dalit representatives place greater pressure on BDOs to aid Dalit workers in their constituencies.

In the 2017 State Assembly elections, three sets of parties gained control of the chamber. The BJP, a right-wing, Hindu nationalist party, won almost 80% of the seats reserved for Dalit candidates and about 75% of the overall assembly. The Samajwadi Party (SP) won the second largest contingent of reserved seats (seven out of 85). The SP, a longtime opponent of the “low-caste” parties in the state, has a base that is mainly composed of the Other Backward Classes—a set of voters that are economically underprivileged but not historically marginalized by the caste system. Lastly, it was the “low-caste” parties who won the majority’s share of the remaining seats (both reserved and general).²⁶ Together, these parties secured seven out of 85 reserved seats and 32 of the 403 total assembly positions. These three groups of parties differ greatly in terms of their support for government benefits towards the low-caste population. While the BJP has historically been opposed to caste-based redistribution, it secured such a large share of the reserved seats by garnering the support of a large swath of Dalit voters. Hence, it is not clear whether they would be particularly in favor of or against their Dalit MLAs deviating resources towards Dalit workers through NREGA. Within the constituencies won by the SP—being a staunch, historical rival of the BSP and other “low-caste” parties—I would expect to see no improvement in the payment wait times of Dalit workers. Lastly, the representatives with the biggest incentives to favor Dalit workers are those belonging to the “low-caste” parties. Moreover, within low-caste voters, Jatavs have most con-

²⁶These parties are the BSP, ADAL, and SBSP. Typically, the BSP is the largest low-caste party in the state and has gone as far as to control of the Uttar Pradesh State Assembly in previous cycles. During the 2017 elections, however, the BSP was unable to secure a large amount of seats (only 19 seats total) despite of gaining 22% of the total votes in the state.

sistently supported these (Jaffrelot, 2022). Therefore, if any Dalit MLAs was to disproportionately advantage Jatav workers, they ought to belong to a “low-caste” party.

Table 6: Dalit MLA performance and party ideology

	E(Late days)		
	SP	BJP	Non-BJP/SP
Dalit MLA	13.08*** (0.0206)	-0.255 (0.768)	-3.158*** (0.918)
Observations	17370	151748	35114
Avg. Y	13.03	14.08	13.22
Bandwidth	0.0250	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes	Yes
Sample	Dalit Only	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).

Table 6 presents the results of estimating the regression in equation (4) separately for constituencies ruled by either the SP, BJP, or a party other than these two (either a low-caste party or an independent candidate). The results in this table largely agree with the predictions made above. In column one, we see that the main opposition to the “low-caste” parties in the state is largely responsible for the increase in payment times across reserved and general constituencies. This could be driven by SP representatives attempting to instead advantage OBC voters, who compose the majority or the party’s vote base. Within ACs won by the BJP, there seems to be no significant difference induced by the caste group of the MLA. While the BJP has been routinely opposed to affirmative action policies for Dalits, an important reason that the party was able to secure such a large number of seats in Uttar Pradesh’s assembly was their increased support among the low-caste voters. This could explain why we do not see positive or negative effects of reservation in areas ruled by the BJP, as the two forces described above might cancel each other out. Lastly, in column three we see that in constituencies won by “low-caste” parties MLAs belonging to a scheduled caste lead to a large and significant decrease in the payment times for low-caste workers. The estimated effect represents a reduction of almost 24% in the average number of days until NREGA laborers receive their wages.

Table 7 presents the parameters estimated via the regression in equation (5). The first row represents the effect of having an Dalit MLA on non-Jatav, low-caste workers within the sample of constituencies won by each of the three sets of parties. These estimates are essentially the same as

the effects for the entire Dalit NREGA workforce. The most important result in this table appears in the second row, which presents the difference in the effect of having an Dalit MLA for Jatav workers, relative to the rest of the low-caste participants in NREGA. Dalit MLAs in ACs governed by either the SP or BJP do not demonstrate a differential advantage towards Jatavs relative to non-Dalit MLAs. In constituencies won by the “low-caste” parties, however, we can see a large (albeit marginally significant) *further* decrease in the time it takes for Jatavs to receive their NREGA wages. Jatavs in reserved constituencies are paid 5.59 days sooner than Jatavs in ACs represented by a non-Dalit state legislator. This represents a nearly 30% increase in the benefits from political affirmative action received by non-Jatav workers under a low-caste legislator, who on average receive NREGA wages 4.3 days earlier than non-Jatavs in general constituencies. Hence, while there does not seem to be conclusive evidence of a widespread favoritism of Jatavs, the gap between the largest scheduled caste in the state and the remaining Dalits is widened significantly by the election of low-caste MLAs in ACs won by the “low-caste” parties.

Table 7: Caste heterogeneity in the performance of Dalit MLAs by party ideology

	E(Late days)		
	SP	BJP	Non-BJP/SP
Dalit MLA	12.63*** (0.404)	-0.0554 (0.771)	-4.319*** (1.061)
Dalit MLA \times Jatav	-0.940 (1.027)	-0.182 (0.379)	-1.270+ (0.717)
Observations	9637	77611	19314
Avg. Y	13.01	14.00	13.05
Bandwidth	0.0250	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes	Yes
Sample	Dalit Only	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).

8 Discussion and conclusion

In this paper, I explore how state MLAs influence, differentially by caste, the timing of payment to NREGA workers. In order to answer this question, I first create a new dataset where I predict the caste of workers in this public workfare program using the cultural link between castes and names. Based on simulations, I show that surnames allow me to estimate the correct caste of individuals with 73.6% probability, on average. Notwithstanding a loss in efficiency stemming from having a small training dataset (relative to the size of the dataset where I predict the caste of workers), this procedure allows me to reliably estimate whether a person belongs to the Jatav caste or another Dalit caste group, the most important dichotomy for the purposes of this study. Using this new dataset, I estimate that having a Dalit MLA leads to a 3.3 percentage points increase in the frequency of late payments to low-caste NREGA laborers. This represents a 12.4% increase relative to the comparison group mean. Despite the large economic significance, this result is relatively imprecisely estimated. The effect is concentrated in areas where the block officer is overseen by multiple MLAs, meaning that each of them has lower control over local bureaucrats. In contradiction with the right-wing narrative of Jatavs having cornered the benefits of political affirmative action, I show that there is no widespread evidence of one caste being particularly favored over others. Restricting attention to constituencies won by parties expressly aligned with the low-caste voters, Dalit workers as a whole receive their payments 24% days sooner relative to their low-caste counterparts in the comparison group. Only within constituencies won by parties expressly aligned with the low-caste voters, I observe that Dalit workers as a whole receive their payments 24% days sooner relative to their low-caste counterparts in the comparison group and that this effect is even more pronounced when comparing Jatav workers to the remainder of the Dalit NREGA beneficiaries. While this subsample produces results that align with the BJP's divisive rhetoric, non-Jatav Dalits, too, receive a large and significant benefit from being represented by a member of their same caste category—receiving their payments 33% sooner than low-caste workers in the control group. Therefore, even though some degree of favoritism towards the largest Dalit caste in the state was to be expected, it has not come at the expense of workers from other underprivileged castes.

A potential limitation of this study is that its findings might not extend to other caste-based affirmative action policies or even other outcomes within the control of Dalit MLAs. Besides being the feature of NREGA that is most directly under the control of MLAs, I chose to focus on the timing of payments because of its political salience. Workers often cite late payments as one of the main problems in the implementation of NREGA schemes, making this policy lever one of great relevance to MLAs as they aim to win the support of their constituents (Nandy, 2021). As liaisons between their constituents and the government bureaucrats who manage the bulk of public welfare programs, Dalit state legislators occupy an important place in the political hierarchy. Their position could be instrumental in helping other low-caste Indians gain access to much-needed public aid. Due to the attention that voters pay to the timely payment of NREGA wages and the importance of the

MLAs’ political position, it is not unreasonable that any potential favoritism towards Jatavs would become evident in this study’s setting.

In future iterations of this project, I will look to include outcomes that affect a broader segment of the population. One such outcome is school funding. Just as with the payment of NREGA wages, MLAs can influence the allocation of education funds to villages and target those that have a higher concentration of specific Dalit castes. The appeal of studying this additional public good’s allocation is that it is informative about whether Dalit MLAs choose to advantage people from a particular caste group when the beneficiaries belong to a broader socioeconomic category. Children attending schools in India come from a wide range of backgrounds, as opposed to NREGA beneficiaries who are mostly poor and live in rural settings. Beyond extensions to this project, an interesting question arises from the empirically proven connection between castes and names: if individual recognize name choice as an important signal of their group identity, does this choice respond to social and economic forces? For instance, we might expect that Dalit families choose names that are less caste-salient in areas where the rates of caste-based violence are higher. In contrast, we might see low-caste families adopting “upper-caste” surnames in places where employment opportunities are improving. Using data that I have begun collecting from the universe of official name changes in Maharashtra, I hope to answer these questions in future work.

The implications and motivation for this study could extend to affirmative action programs elsewhere. Policies such as caste-based reservations that intend to equalize opportunities for disadvantaged individuals are often aimed at broad and heterogeneous groups of people. Scholarships for Hispanic students in the United States, for instance, target a set of the population that is composed of many different socioeconomic and political backgrounds. The welfare gaps between Hispanics and the white majority in the US mirror those between Dalits and the rest of Indians. Likewise, we might think it valuable to consider how affirmative actions aimed at the Hispanic population affect distinct groups within this broader ethnic category.

Given the wide-reaching and central role that caste plays in the everyday lives of Indians, it is imperative that we consider how this dimension of identity interacts with public policies. Political reservations for Dalits, in particular, are aimed at closing the gap between low-caste Indians and their better-off counterparts. Just as we have sought to understand the determinants and mitigating factors of inequality across these broad caste categories, it stands to reason that we explore how different Dalit caste groups might be advancing at different rates. Recent work has highlighted the degree of within-caste-category inequality (Joshi et al., 2022). Without exploring the differing impacts that public policies such as the one I study have across castes, we would be ill-equipped to design welfare programs aimed at remedying these cleavages.

References

- Ran Abramitzky, Leah Boustan, Katherine Eriksson, James Feigenbaum, and Santiago Pérez. Automated linking of historical data. *Journal of Economic Literature*, 59(3):865–918, September 2021. doi: 10.1257/jel.20201599. URL <https://www.aeaweb.org/articles?id=10.1257/jel.20201599>.
- Siwan Anderson and Patrick Francois. Reservations and the Politics of Fear. *Working Paper*, page 59, May 2017.
- Siwan Anderson, Patrick Francois, and Ashok Kotwal. Clientelism in Indian Villages. *American Economic Review*, 105(6):1780–1816, June 2015. ISSN 0002-8282. doi: 10.1257/aer.20130623. URL <https://pubs.aeaweb.org/doi/10.1257/aer.20130623>.
- Abhijit Banerjee and Rohini Somanathan. The political economy of public goods: Some evidence from India. *Journal of Development Economics*, 82(2):287–314, March 2007. ISSN 03043878. doi: 10.1016/j.jdeveco.2006.04.005. URL <https://linkinghub.elsevier.com/retrieve/pii/S0304387806000587>.
- Pranab K. Bardhan, Dilip Mookherjee, and Monica Parra Torrado. Impact of Political Reservations in West Bengal Local Governments on Anti-Poverty Targeting. *Journal of Globalization and Development*, 1(1), January 2010. ISSN 1948-1837. doi: 10.2202/1948-1837.1025. URL <https://www.degruyter.com/document/doi/10.2202/1948-1837.1025/html>.
- Timothy Besley, Rohini Pande, Lupin Rahman, and Vijayendra Rao. The Politics of Public Good Provision: Evidence from Indian Local Governments. *Journal of the European Economic Association*, 2(2-3):416–426, May 2004. ISSN 1542-4766, 1542-4774. doi: 10.1162/154247604323068104. URL <https://academic.oup.com/jeea/jeea/article/2194944/The>.
- Timothy Besley, Vijayendra Rao, and Rohini Pande. Participatory Democracy in Action: Survey Evidence from South India. *Journal of the European Economic Association*, 3(2):648–657, 2005.
- Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014. doi: <https://doi.org/10.3982/ECTA11757>. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA11757>.
- Census of India. All India Census Data 2011, 2011.
- Aimee Chin and Nishith Prakash. The redistributive effects of political reservation for minorities: Evidence from India. *Journal of Development Economics*, 96(2):265–277, November 2011. ISSN 03043878. doi: 10.1016/j.jdeveco.2010.10.004. URL <https://linkinghub.elsevier.com/retrieve/pii/S030438781000115X>.

- Ricardo Dahis, Emily Nix, and Nancy Qian. Choosing racial identity in the united states, 1880-1940. Working Paper 26465, National Bureau of Economic Research, November 2020. URL <http://www.nber.org/papers/w26465>.
- Thad Dunning and Janhavi Nilekani. Ethnic Quotas and Political Mobilization: Caste, Parties, and Distribution in Indian Village Councils. *American Political Science Review*, 107(1):35–56, February 2013. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055412000573. URL https://www.cambridge.org/core/product/identifier/S0003055412000573/type/journal_article.
- Saad Gulzar and Benjamin J. Pasquale. Politicians, bureaucrats, and development: Evidence from india. *American Political Science Review*, 111(1):162–183, 2017. doi: 10.1017/S0003055416000502.
- Saad Gulzar, Nicholas Haas, and Benjamin Pasquale. Does Political Affirmative Action Work, and for Whom? Theory and Evidence on India’s Scheduled Areas. *American Political Science Review*, 114(4):1230–1246, November 2020. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055420000532. URL https://www.cambridge.org/core/product/identifier/S0003055420000532/type/journal_article.
- Jonas Hjort. Ethnic Divisions and Production in Firms *. *The Quarterly Journal of Economics*, 129(4):1899–1946, 10 2014. ISSN 0033-5533. doi: 10.1093/qje/qju028. URL <https://doi.org/10.1093/qje/qju028>.
- Lakshmi Iyer and Anandi Mani. Traveling Agents: Political Change and Bureaucratic Turnover in India. *The Review of Economics and Statistics*, 94(3):723–739, August 2012. ISSN 0034-6535. doi: 10.1162/REST_a_00183. URL https://doi.org/10.1162/REST_a_00183. _eprint: https://direct.mit.edu/rest/article-pdf/94/3/723/1916990/rest_a_00183.pdf.
- Lakshmi Iyer and Maya Reddy. Redrawing the lines: Did political incumbents influence electoral redistricting in the world’s largest democracy? Working Paper 14-051, Harvard Business School, 2013.
- C. Jaffrelot and S. Kumar. *Rise of the Plebeians?: The Changing Face of Indian Legislative Assemblies*. Exploring the political in South Asia. Routledge, 2009. ISBN 9780415460927. URL <https://books.google.com/books?id=MLu6E1AjLmgC>.
- Christophe Jaffrelot. Why do poor voters choose a pro-rich party in india?, Feb 2022.
- Raja Jayaraman. Personal identity in a globalized world: Cultural roots of hindu personal names and surnames. *The Journal of Popular Culture*, 38(3):476–490, 2005. doi: <https://doi.org/10.1111/j.0022-3840.2005.00124.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.0022-3840.2005.00124.x>.

- Francesca Refsum Jensenius. Development from Representation? A Study of Quotas for the Scheduled Castes in India. *American Economic Journal: Applied Economics*, 7(3):196–220, July 2015. ISSN 1945-7782, 1945-7790. doi: 10.1257/app.20140201. URL <https://pubs.aeaweb.org/doi/10.1257/app.20140201>.
- Shareen Joshi, Nishtha Kochhar, and Vijayendra Rao. Fractal inequality in rural India: class, caste and jati in Bihar. *Oxford Open Economics*, 1, 02 2022. ISSN 2752-5074. doi: 10.1093/ooec/odab004. URL <https://doi.org/10.1093/ooec/odab004>. odab004.
- Ministry of Rural Development. 4th edition, 2013.
- Priya Moorjani, Kumarasamy Thangaraj, Nick Patterson, Mark Lipson, Po-Ru Loh, Periyasamy Govindaraj, Bonnie Berger, David Reich, and Lalji Singh. Genetic evidence for recent population mixture in india. *The American Journal of Human Genetics*, 93(3):422–438, Sep 2013. doi: 10.1016/j.ajhg.2013.07.006.
- Kaivan Munshi. Caste and the Indian Economy. *Journal of Economic Literature*, 57(4):781–834, December 2019. ISSN 0022-0515. doi: 10.1257/jel.20171307. URL <https://pubs.aeaweb.org/doi/10.1257/jel.20171307>.
- Kaivan Munshi and Mark Rosenzweig. The Efficacy of Parochial Politics: Caste, Commitment, and Competence in Indian Local Governments. *NBER Working Paper Series*, September 2008. doi: 10.3386/w14335. URL <http://www.nber.org/papers/w14335.pdf>.
- Kaivan Munshi and Mark Rosenzweig. Networks, Commitment, and Competence: Caste in Indian Local Politics. *NBER Working Paper Series*, July 2013. doi: 10.3386/w19197. URL <http://www.nber.org/papers/w19197.pdf>.
- Debmalya Nandy. Poor remuneration and late payments: The nrega payments trap. *The Wire*, Sep 2021. URL <https://thewire.in/labour/poorremuneration-and-late-payments>.
- Sajith Pal. What’s in a first name? your social status. *The Times of India*, Jul 2019. URL <https://timesofindia.indiatimes.com/blogs/open-patent/whats-in-a-first-name-your-social-status/>.
- Rohini Pande. Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India. *American Economic Review*, 93(4):1132–1151, August 2003. ISSN 0002-8282. doi: 10.1257/000282803769206232. URL <https://pubs.aeaweb.org/doi/10.1257/000282803769206232>.
- Sourav Sarkar. *Essays on political economy and development*. PhD thesis, MIT, Department of Economics, 2019.

K.S. Singh. *Communities, Segments, Synonyms, Surnames and Titles*. National series. Anthropological Survey of India, 1996. ISBN 9780195633573. URL <https://books.google.com/books?id=bfAMAQAAMAAJ>.

A Aggregating caste predictions within households

aggregates This section includes a detailed example of how I aggregate the predictions for each name-part across individuals in a household. The three aggregation methods that I use, besides the one motivated by Bayes' rule, are a simple average, an average weighted by the relative over-representation of a name within the given caste, and the maximum over name-parts and household members.

Consider a household consisting of two three individuals: a mother "Aastha Jatav," a father "Rajesh Chamar," and a son "Ameet Chamar." I am interested in predicting the probability that all members of this household belong to the Jatav caste. Since I only use last names in my caste predictions, I need to aggregate into a single measure the probability that each last name has of belonging to a Jatav person: $P(\text{caste} = \text{Jatav} | \text{name} = \text{Jatav})$, $P(\text{caste} = \text{Jatav} | \text{name} = \text{Chamar})$, and $P(\text{caste} = \text{Jatav} | \text{name} = \text{Chamar})$. Notice that the last name "Chamar" is repeated twice. In the current version of this procedure, I have not corrected for this to allow for repeated names within a household to serve as stronger signals of caste belonging. While I recognize the potential drawbacks of this decision, such as larger households mechanically appearing to be more caste-salient without necessarily being so, this is a prediction exercise and the internal simulations suggest that the estimated caste probabilities rendered by this algorithm are of high quality.

Using the Bayes' rule method, these probabilities would be aggregated as follows. First, I start with a prior set equal to the average share of Dalits from the specific caste (Jatavs in this case) in the Hindi Belt region: $P(\text{caste} = \text{Jatav})$. In what follows, I abbreviate $\text{caste} = \text{Jatav}$ as Jatav . The first posterior will be given by this expression:

$$P(\text{Jatav} | \text{name}_1 = \text{Jatav}) = \frac{P(\text{name} = \text{Jatav} | \text{Jatav})P(\text{Jatav})}{P(\text{name} = \text{Jatav})} \quad (6)$$

The second posterior will use the same expression, replacing the prior with $P(\text{Jatav} | \text{name}_1 = \text{Jatav})$:

$$P(\text{Jatav} | \text{name}_1 = \text{Jatav}, \text{name}_2 = \text{Chamar}) = \frac{P(\text{name} = \text{Chamar} | \text{Jatav})P(\text{Jatav} | \text{name}_1 = \text{Jatav})}{P(\text{name} = \text{Chamar})} \quad (7)$$

This procedure continues until all last names have been incorporated. In this case, the final expression is:

$$\begin{aligned} P(\text{Jatav} | \text{name}_1 = \text{Jatav}, \text{name}_2 = \text{Chamar}, \text{name}_3 = \text{Chamar}) = \\ \frac{P(\text{name} = \text{Chamar} | \text{Jatav})P(\text{Jatav} | \text{name}_1 = \text{Jatav}, \text{name}_2 = \text{Chamar})}{P(\text{name} = \text{Chamar})} \end{aligned} \quad (8)$$

This formula can be interpreted as generating a Bayesian posterior from all last names, assuming

that each name-part is an independent signal. This is certainly not true, names (especially last names) are highly correlated within households. However, due to sample size restrictions within the REDS dataset, I am not able to calculate the multiple permutations of these names across individuals and within households that would be necessary to relax this assumption. Finally, I normalize the sum of probabilities within a household, across castes to be equal to one. This last step is necessary precisely because the independence assumption is not entirely correct. That being said, as I mentioned above for the case of repeated last names, since this is a prediction exercise, the theoretical correctness of the procedure is less important relative to generating the most correct predictions possible.

Finally, the alternative mechanisms that I use to aggregate are taking a simple average across the last names, that is:

$$P(Jatav|name_1 = Jatav, name_2 = Chamar, name_3 = Chamar) = \frac{1}{3} [P(Jatav|name = Jatav) + P(Jatav|name = Chamar) + P(Jatav|name = Chamar)] \quad (9)$$

Using a weighted average, similar as the expression in equation (9) where each signal is weighted by the relative over-representation of a last name within a specific caste, e.g.: $\frac{P(name=Chamar|Jatav)}{P(name=Chamar)}$

Finally, I consider the maximum operator across all predictions as my fourth aggregation method. The results in Appendices B and C show that all aggregation methods perform relatively well and generate the same qualitative patterns.

B Additional caste prediction simulation results

B.1 Efficiency results

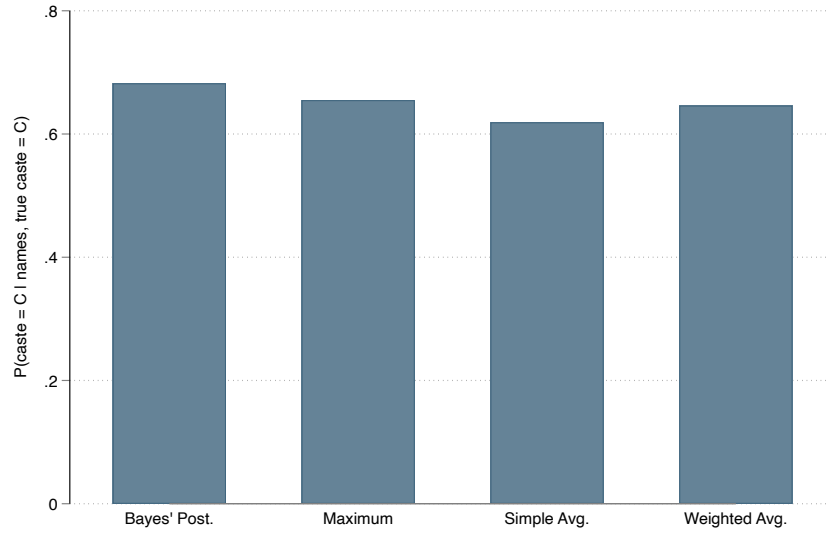


Figure 9: Efficiency of predictions for each aggregation method

Note: This figure represents the efficiency (one minus the Type II error rate) across all castes and across all sizes of the training dataset, using only last names. The horizontal axis represents the method used to aggregate the predictions across name-parts and individuals.

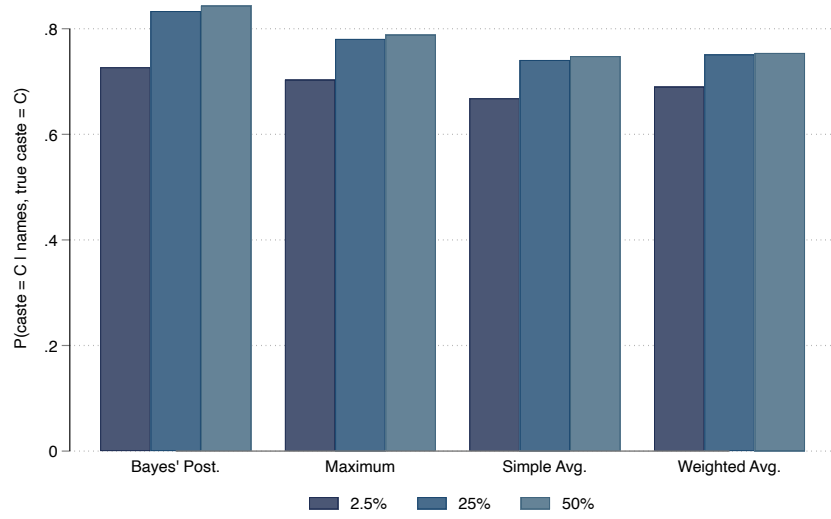


Figure 10: Efficiency of predictions for each aggregation type, varying the size of the training dataset

Note: This figure represents the efficiency (one minus the Type II error rate) across all castes, using only last names. The horizontal axis represents the method used to aggregate the predictions across name-parts and individuals. The differently-colored bars represent the size of the training dataset, relative to the size of the entire dataset. The darkest color represents the smallest training dataset and the lightest represents the largest training dataset.

B.2 Type I error results

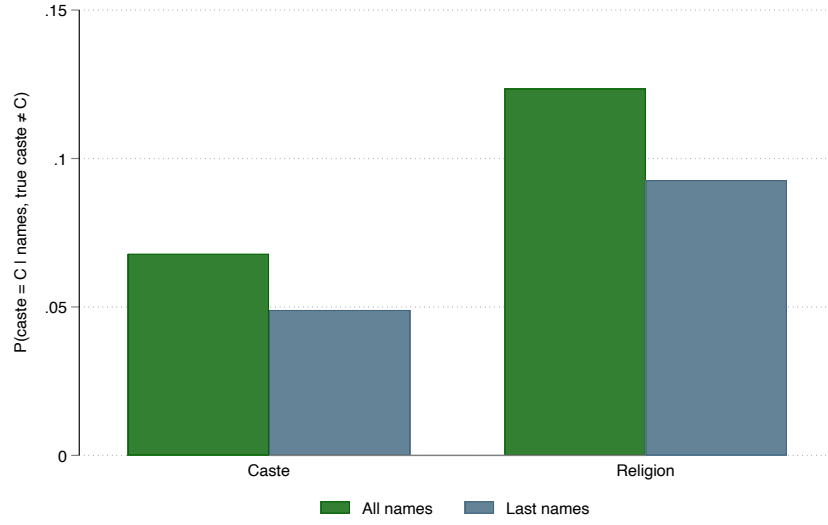


Figure 11: Type I error of caste and religion predictions using all names vs last names only

Note: This figure represents the inaccuracy (the Type I error rate) across all sizes of the training dataset and aggregation methods, and across all castes (for the first set of bars) or both religions (Hinduism and Islam, for the second set of bars). The horizontal axis represents the whether the predictions are aimed at caste or religion. The green bars represent the Type I error rate when using all name-parts and the blue represent the error rate when using only last names.

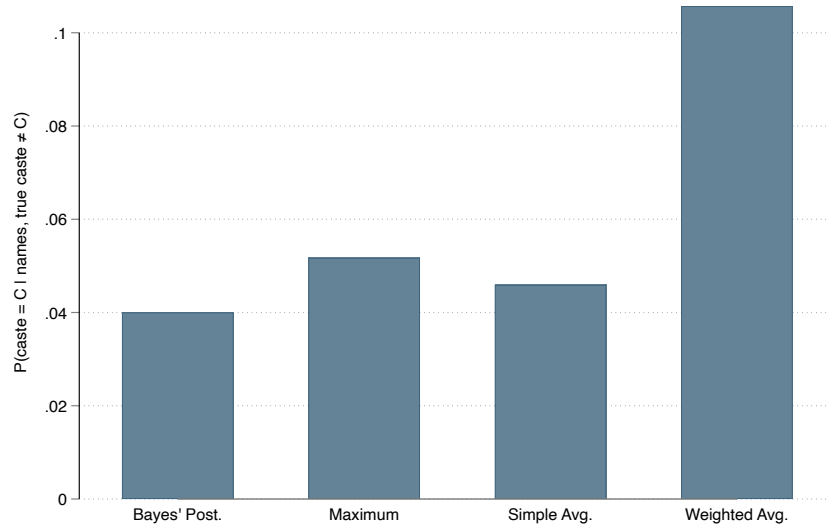


Figure 12: Type I error of predictions for each aggregation method

Note: This figure represents the inaccuracy (the Type I error rate) across all castes and across all sizes of the training dataset, using only last names. The horizontal axis represents the method used to aggregate the predictions across name-parts and individuals.

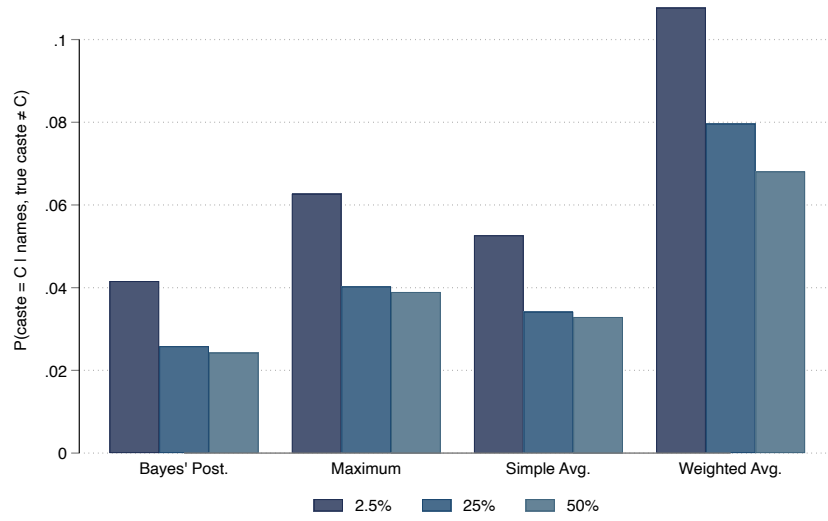


Figure 13: Type I error of predictions for each aggregation, varying the size of the training dataset

Note: This figure represents the inaccuracy (the Type I error rate) across all castes, using only last names. The horizontal axis represents the method used to aggregate the predictions across name-parts and individuals. The differently-colored bars represent the size of the training dataset, relative to the size of the entire dataset. The darkest color represents the smallest training dataset and the lightest represents the largest training dataset.

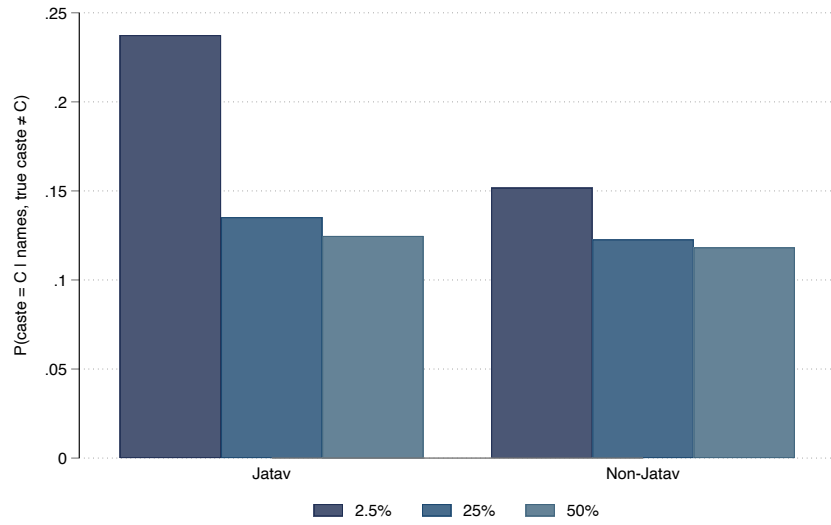


Figure 14: Type I error of predictions for Jatavs vs. non-Jatavs, varying the size of the training dataset

Note: This figure represents the inaccuracy (the Type I error rate) of caste predictions, using only last names. The horizontal axis represents the caste being predicted, separated into Jatav (the largest caste group among Dalits in Uttar Pradesh) and non-Jatav Dalits. The differently-colored bars represent the size of the training dataset, relative to the size of the entire dataset. The darkest color represents the smallest training dataset and the lightest represents the largest training dataset.

C Additional empirical results

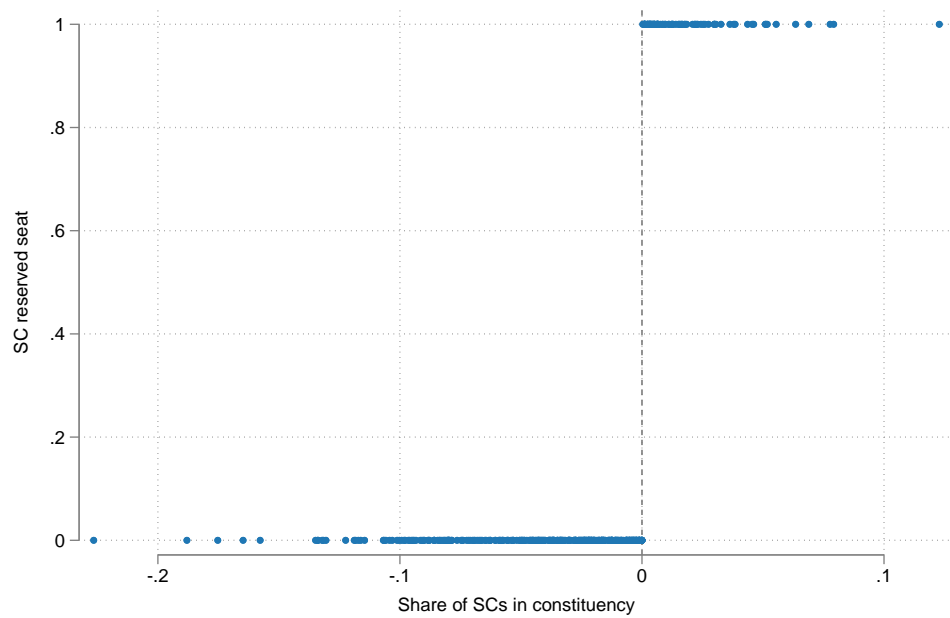


Figure 15: Perfect compliance with the assignment of reserved and general constituencies

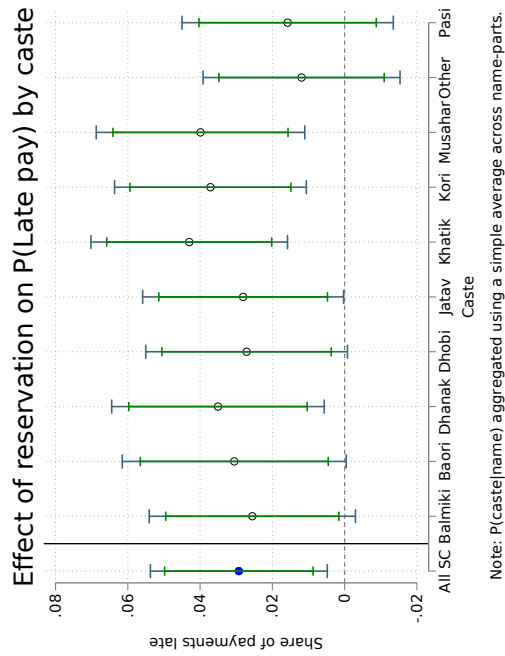
Note: This figure is a scatter plot of all constituencies based on the share of Dalit residents, relative to the district reservation cutoff (horizontal axis). The vertical axis represents whether these constituencies are reserved or not.

Table 8: Heterogeneity in the effect of reservations for Jatavs vs. non-Jatavs

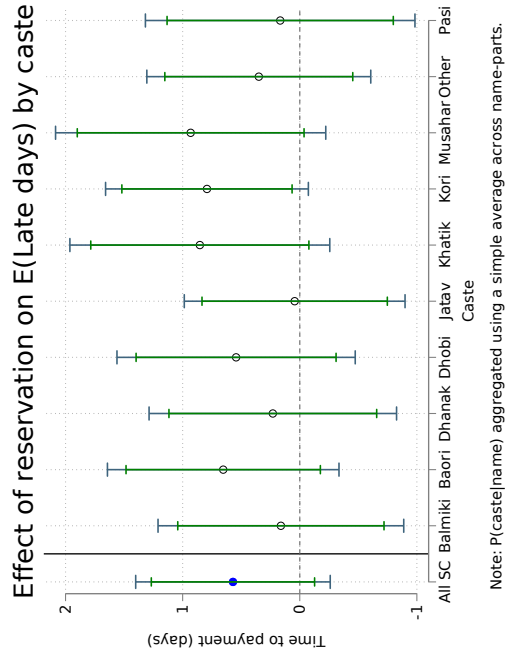
	P(Late pay)	E(Late days)
Dalit MLA	0.0309 (0.0260)	0.869 (0.927)
Dalit MLA×P(Jatav)	0.000957 (0.0294)	-0.674 (1.001)
Observations	101275	101275
Avg. Y	0.265	13.94
S.D. P(Jatav)	0.172	0.172
Bandwidth	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes
Sample	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).



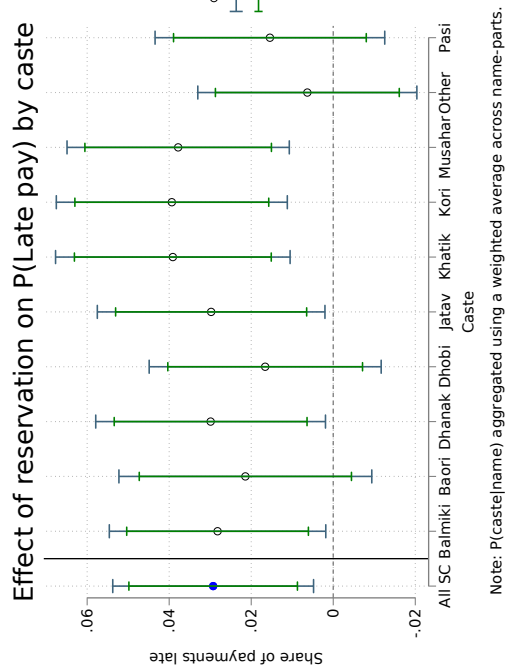
(a) Effects by caste on the share of late payments



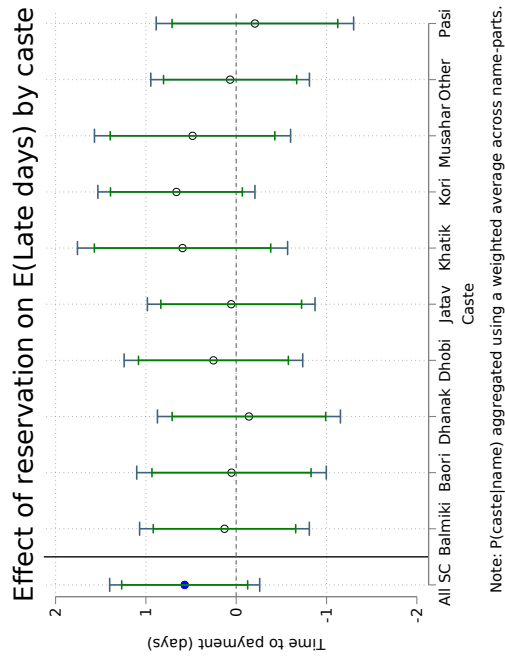
(b) Effects by caste on the average time to payment

Figure 16: Caste-specific average treatment effects of having a Dalit state legislator

Note: Each figure represents the average treatment effect, at the reservation cutoff, of having a Dalit state legislator on a measure of late payments to NREGA workers. Figure 16a uses the share of late payments as the outcome while Figure 16b uses the average number of days until payment as the dependent variable. The horizontal axes denote the caste composition of the sample, ordered alphabetically: either all Dalit (SC) workers (the leftmost point represented by a filled circle) or each of the individual caste groups (represented by empty circles). The dark blue bars mark the width of the 95% confidence interval around each ATE, while the green bars mark the 90% confidence intervals. The caste predictions were aggregated across name-parts and individuals to the household level using a simple average.



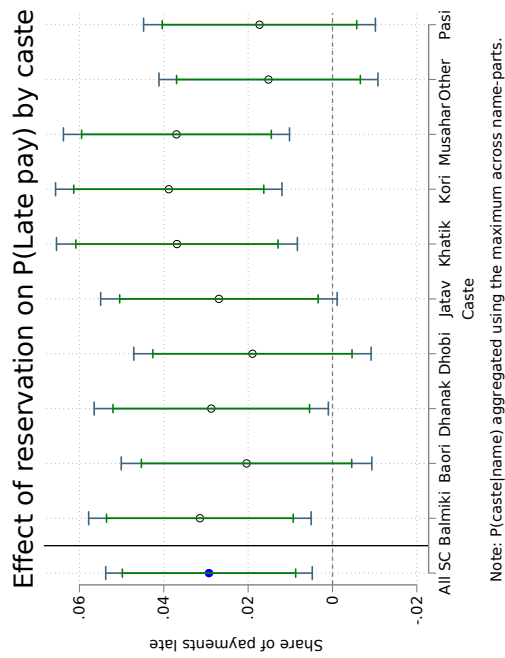
(a) Effects by caste on the share of late payments



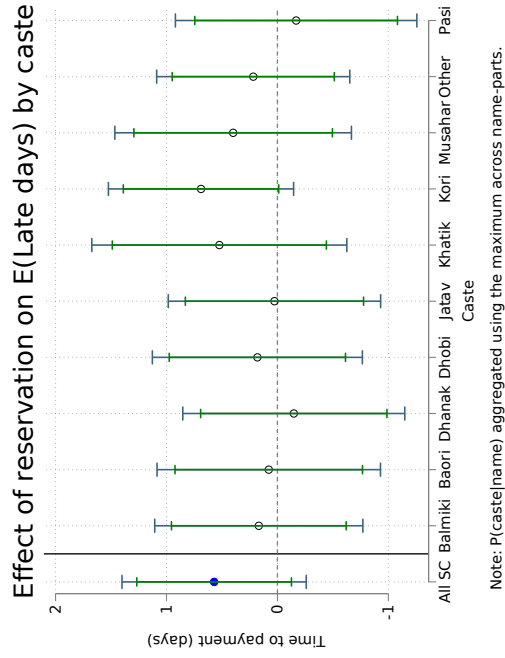
(b) Effects by caste on the average time to payment

Figure 17: Caste-specific average treatment effects of having a Dalit state legislator

Note: Each figure represents the average treatment effect, at the reservation cutoff, of having a Dalit state legislator on a measure of late payments to NREGA workers. Figure 17a uses the share of late payments as the outcome while Figure 17b uses the average number of days until payment as the dependent variable. The horizontal axes denote the caste composition of the sample, ordered alphabetically: either all Dalit (SC) workers (the leftmost point represented by a filled circle) or each of the individual caste groups (represented by empty circles). The dark blue bars mark the width of the 95% confidence interval around each ATE, while the green bars mark the 90% confidence intervals. The caste predictions were aggregated across name-parts and individuals to the household level using a weighted average.



(a) Effects by caste on the share of late payments



(b) Effects by caste on the average time to payment

Figure 18: Caste-specific average treatment effects of having a Dalit state legislator

Note: Each figure represents the average treatment effect, at the reservation cutoff, of having a Dalit state legislator on a measure of late payments to NREGA workers. Figure 18a uses the share of late payments as the outcome while Figure 18b uses the average number of days until payment as the dependent variable. The horizontal axes denote the caste composition of the sample, ordered alphabetically: either all Dalit (SC) workers (the leftmost point represented by a filled circle) or each of the individual caste groups (represented by empty circles). The dark blue bars mark the width of the 95% confidence interval around each ATE, while the green bars mark the 90% confidence intervals. The caste predictions were aggregated across name-parts and individuals to the household level using a the maximum operator.

Table 9: Caste heterogeneity of MLA performance and their relationship with BDOs

	P(Late pay)		E(Late days)	
	Single MLA	Mult. MLAs	Single MLA	Mult. MLAs
Dalit MLA	0.0172 (0.0252)	0.0544** (0.0226)	0.563 (0.902)	1.082 (0.770)
Dalit MLA \times Jatav	-0.0106 (0.0127)	0.0106 (0.0132)	-0.343 (0.462)	-0.378 (0.454)
Observations	44733	56542	44733	56542
Avg.Y	0.258	0.271	13.52	14.29
Bandwidth	0.0250	0.0250	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes	Yes	Yes
Sample	Dalit Only	Dalit Only	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).

Table 10: Dalit MLA performance and party ideology

	P(Late pay)		
	SP	BJP	Non-BJP/SP
Dalit MLA	0.285*** (0.000600)	0.00358 (0.0215)	-0.335*** (0.0282)
Observations	17370	151748	35114
Avg.Y	0.225	0.269	0.240
Bandwidth	0.0250	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes	Yes
Sample	Dalit Only	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).

Table 11: Caste heterogeneity in the performance of Dalit MLAs by party ideology

	P(Late pay)		
	SP	BJP	Non-BJP/SP
Dalit MLA	0.226*** (0.0149)	0.00847 (0.0211)	-0.405*** (0.0365)
Dalit MLA \times Jatav	0.0163 (0.0366)	-0.000587 (0.0120)	-0.00886 (0.0211)
Observations	9637	77611	19314
Avg.Y	0.219	0.268	0.230
Bandwidth	0.0250	0.0250	0.0250
Yr., Dist. FEs	Yes	Yes	Yes
Sample	Dalit Only	Dalit Only	Dalit Only

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors clustered at the constituency level are shown in parentheses. Each regression includes a series of control variables at the constituency level. These encompass several socioeconomic (e.g., literacy of both Dalit and all residents and share of agricultural workers), demographic (e.g., total population, share of Dalit workers prior to 2001), and public good provision measures (e.g., access to education, electricity, health, and communication facilities) at baseline (1971 or 2001).