Representation in Politics and Labor Markets

Sugat Chaturvedi

Thesis submitted to the Indian Statistical Institute
in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

Representation in Politics and Labor Markets

Sugat Chaturvedi

October 2021

Thesis Supervisor: Prof. Tridip Ray

Thesis submitted to the Indian Statistical Institute in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Acknowledgements

I have now traversed the long road to this doctoral thesis. One that is enriching, full of intriguing questions as well as bestowing many opportunities to experiment and learn. Yet, it is hard to imagine completing this journey without constant support and encouragement from many people. I would first like to thank my advisor Prof. Tridip Ray for his invaluable support and guidance as well as the freedom to pursue interesting research directions. I am especially grateful to Dr. Sabyasachi Das for always pushing me to do better, for his constant guidance and invaluable feedback throughout the duration of my PhD. I have enjoyed many discussions with him and learnt a great deal in the process. I also wish to thank Dr. Kanika Mahajan for being an incredible mentor and pushing me towards new research directions. I am grateful to Dr. Zahra Siddique as well for her guidance. I also wish to thank Prof. Sourav Bhattacharya whose perspective and comments on my work have motivated me to think more deeply about research.

I would also like to thank all the faculty members of the Economic and Planning Unit, Indian Statistical Institute, Delhi. I would especially like to thank Prof. Farzana Afridi who introduced me to empirical research as well as labor economics and gave many useful comments on my work. I am also grateful to Professors Monisankar Bishnu, Abhiroop Mukhopadhyay, Debasis Mishra and Chetan Ghate for their feedback and encouragement at various stages of my research. I am thankful to my seniors Dhritiman Gupta and Prachi Singh with whom I have been fortunate to have many enriching discussions.

I would like to convey my deepest thanks to all my family members. They have taught me many important skills in life that have also enabled me to smoothly navigate through my PhD. I have also relied on their brilliant insights from time to time. First, I wish to thank my mother Dr. Usha Kanta Chaturvedi who has always ensured that I stay nourished and healthy. I have learnt the importance of hard work and being open minded towards new viewpoints from her. My father Prof. Mithilesh Chaturvedi has motivated me to maintain a broad vision and has inspired me to persevere in the pursuit of knowledge. My sister Rochana Chaturvedi has been a wonderful coauthor and has had an instrumental role in nudging me towards the exciting fields of machine learning and natural language processing. Her encouragement has pushed me to excel. I am also grateful to my sister Shesha Chaturvedi who has often provided me a welcome break from research through her quick-witted humor but has also often shared a new angle to think about social issues during many of our discussions.

Finally, I wish to thank Indian Statistical Institute for providing a conducive environment for research. The generous fellowship provided by the institute is gratefully acknowledged.

To my parents

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

Representation in politics and the workplace is not only intrinsically important to achieve a more equal and inclusive society, but also plays an instrumental role in the economic development process. It ensures that a diverse set of perspectives are put forward and a varied set of voices are heard. It is well known that political leaders allocate more resources to their co-ethnics (Franck and Rainer, 2012; Hodler and Raschky, 2014; De Luca et al., 2018). Thus, reducing inequality in access to power can directly decrease between-group inequality. Members of different groups may also have contrasting personal preferences, and therefore, their representation can change policy even in the absence of ethnic favoritism. For example, Bhalotra et al. (2014) show that political representation of Muslims improves health and education outcomes in their constituencies in India but do not find any evidence of religious favoritism. Meyersson (2014) similarly find that Islamic party rule in Turkish municipalities had a persistent positive effect on female education.

In particular, women's representation in politics may change public policy to align it with women's preferences (Chattopadhyay and Duflo, 2004; Bhalotra and Clots-Figueras, 2014). It may also reduce gender stereotypes and improve perceptions of women's effectiveness as leaders (Beaman et al., 2009). In addition, female leadership may also entail role model effects. Beaman et al. (2012) find that female leadership in village councils in India reduced the gender gap in career aspirations and education achievement of adolescent girls. Improved access to the labor market may also have similar effects on women's well being. Jensen (2012) show that awareness of job opportunities not only increased the schooling of women, but also affected their marriage and fertility decisions. Against this backdrop, this thesis deals with the question of representation of marginalized groups in politics and the labor markets. It studies how political institutions and cultural biases shape economic outcomes.

It considers three open questions:

- Why gender quotas in politics may have an effect on public policy?
- What gender stereotypes do employers hold and how they affect application decisions of candidates?
- How do political institutions affect representation of minorities?

The sections that follow provide an overview of the research questions, methods used to answer them and the main results from the subsequent chapters.

1.2 The Importance of Being Earnest: What Explains the Gender Quota Effect in Politics?

There is now an extensive literature documenting the effects of electoral gender quotas on policy after the pioneering work by Chattopadhyay and Duflo (2004). These studies are motivated by the observation that preference for public goods is gendered, and therefore, representation of women should change policy as elected women would prioritize a different set of public goods than men. However, policy may be affected both by the preferences of the elected leaders (Alesina, 1988) as well as that of the voters (Downs, 1957).

Chapter 2 examines why gender quotas in local elections in India might matter for public good provision. Female voters may express the demand for their preferred policy more strongly in the presence of female leadership. This could happen because registering demand to elected authorities is a costly process, especially for women, who face a significantly higher cost of political participation. Women may feel that female leaders are more likely to accede to their demands and may systematically choose to publicly express their demands more in their presence (Iyer and Mani, 2019). We test whether the effect of gender quotas on policy occurs due to supply (female leaders prioritizing different policies) or demand (female voters expressing their demands more vigorously in the presence of a female leader).

To disentangle the two mechanisms, we compile novel data on *household level* allocation of a publicly provided good—toilets—for the entire rural population of Uttar Pradesh comprising over 25 million households. We use the religious and gender identities of the council presidents and household heads as proxies for preference for toilets to disentangle demand effects from supply. We manually classify the religion of Sarpanches as Hindu or Muslim using their names which are culturally distinct. Since classifying religion of over 25 million households manually is not feasible, we identify the religion of household heads based on their names using a character sequence based machine learning model developed by Chaturvedi and Chaturvedi (2020) that can predict the religion based on names with very high accuracy.

Open defecation is widely practiced in rural India due to severe lack of sanitation facilities. Consequently, the construction of toilets received a lot of emphasis from the policy makers making it a politically salient public good. In 2014, the new government at the center launched the Swachh Bharat Mission (SBM) or "Clean India Mission" after emphasizing toilet provision in their policy platform during election campaigns. Under this, a subsidy of ₹12,000 was provided to each eligible household, paid as reimbursement in two equal installments.[1] It is well known that women have a stronger preference for toilets than men due to notions of shame and dignity, and also because of the possibility of harassment (Coffey et al., 2014; Khanna and Das, 2016; Stopnitzky, 2017). Furthermore, the gender gap in preference is higher for Muslims than Hindus due to cultural bias of Hindus against in-home toilets (Coffey et al., 2014, 2017).[2] We show using pan-India survey data sets that women in female headed households enjoy greater autonomy, bargaining power and participation in the political processes. Therefore, female headed households are more likely to express greater demand for toilets especially among Muslim households. This implies that if demand is important, then the gender quota effect, among either Hindu or Muslim Sarpanches, is likely to be greater for households headed by women.

We examine the effect of female quotas in the 2015 elections for village council presidents (or *Sarpanches*) for the universe of 59,000 village councils or gram panchayats (GPs) in U.P. on the provision of household toilets in the subsequent year 2016–17. At least a third of Sarpanch positions are reserved for women based on the population composition of the GPs. We use a population threshold-based fuzzy regression discontinuity design that exploits the discontinuities created by the rule used for the allocation of gender reservation for the position of village council president in the state. We separately estimate the impact of gender quota among Hindu and Muslim Sarpanches.

Among Hindu Sarpanches, gender quota has no average effect on toilet construction. On the other hand, gender quota among Muslim Sarpanches increases toilet construction by 15 percentage points over a sample average of 13%. In other words, the probability of allocation more than doubles. This greater effect could either result from a larger gender gap in preferences among Muslim Sarpanches or a greater demand from women in the Muslim households which are relatively numerous in GPs where Muslim leaders win.

[1]The eligible households include all Below Poverty Line (BPL) households and certain categories of Above Poverty Line (APL) households such as Scheduled Castes/Scheduled Tribes, small and marginal farmers, landless laborers and women headed households.

[2]Another explanatory factor might be greater restrictions on the mobility of Muslim women.

To identify the demand channel, we focus on the gender of the household head. We find that gender quota among both Hindu and Muslim Sarpanches leads to a large increase in probability of toilet allocation among Muslim female headed households. The Hindu female headed households also experience a differential increase in the allocation probability. The effect is, however, smaller in magnitude and statistically imprecise. The fact that Hindu female leaders allocate disproportionately more toilets to Muslim female headed households can only be explained by their response to demand expressed demand by Muslim female headed households. We also show that the result is stronger in GPs with high Muslim population share of Muslims that are likely to have high collective demand for toilets.

We investigate the supply mechanism by focusing on GPs having a close election between a Hindu and a Muslim candidate both in GPs with and without female reservation. Within this sample the religion of the Sarpanch is, in effect, randomly assigned. Hence, the GPs with Hindu and Muslim Sarpanches in this sample will, on average have the same demographic composition and voter preferences. Consequently, any differential treatment effect in the GPs with Muslim Sarpanches (relative to Hindu ones) would be evidence of the supply mechanism. We, however, find that the interaction effect of the religion and gender of the Sarpanch is null implying that the supply mechanism is not present in our context. We, therefore, conclude that the overall effect of gender quota and its heterogeneity across villages is explained exclusively by demand.

Our results reconcile the mixed evidence on the effects of gender quota in elections. The importance of demand further suggests that gender quotas may be more effective in presence of more engaged female voters. From a policy perspective, the results imply that empowering women voters and facilitating their participation in local political processes can lead to significant improvement in their substantive representation.

1.3 Words Matter: Gender, Jobs and Applicant Behavior

In 2017, only 24% of urban Indian women above the age of 15 were working or looking for work compared to 82% of men. Moreover, women who were employed earned 38% lower wages than men on average.[3] The extant literature on gender disparities in employment, wages and occupational segregation finds that this can be due to supply-side factors or discrimination by employers and can manifest during the hiring process as well as while

[3]The gender wage gap is computed using individual-level wage data in the Periodic Labor Force Survey (2017–18), controlling for their age, education and occupation. For young urban workers aged 18-32, the gender wage gap is lower at 27%.

workers advance through their careers. While the final matching of candidates with jobs plays an important role, these disparities can have their genesis in the job search and application behavior of the candidates.

Previous studies for China, Indonesia, and Mexico (Kuhn and Shen, 2013; Kuhn et al., 2020; Ningrum et al., 2020; Helleseter et al., 2020) have examined discrimination at the hiring stage by studying gender preferences exhibited in online job ads. In India too it is not unusual for employers to express explicit gender preferences in the job ads they post online (Chowdhury et al., 2018).[4] Chapter 3 examines explicit employer preferences for hiring men vs women using job ads posted between July 2018 and February 2020 on an online job portal in India that caters primarily to young urban job seekers. Our final data comprises 1,57,890 job ads that specify job location within a single state and for which we can obtain a detailed occupational classification using an unsupervised machine learning algorithm. We also use data on all applications made to these ads by 1.06 million active job seekers to study the impact of explicit preferences on search behavior.

The portal does not have a separate field that allows employers to directly state the preferred gender for a job to the job seekers. However, employers can state this in the accompanying job description. We construct variables indicating an employer's gender preference by carrying out a text search within the job title and description for each job ad.[5] Approximately 4.5% of the job ads in our sample have an explicit female preference, 3.8% have an explicit male preference, and the rest have no explicit gender preference. Job titles such as 'telecaller' and 'office executive' occur frequently in female-targeted jobs while titles such as 'delivery boy' and 'sales executive' are common among male-targeted jobs. These confirm that explicit gender preferences are related to existing occupational gender stereotypes.

We find that jobs that specify higher education/experience/wage are less likely to exhibit an explicit gender preference even within the same occupation. This indicates that jobs with

[4] Such explicit requests are now banned in some countries. For instance, the U.S. banned gendered job ads in 1973. Austria banned these in 2004, while China banned these more recently in 2014. In China, discriminatory job ads violate the Chinese Labor Law. In 2016, after a spate of legal cases, China directly imposed fines on job portals and employers posting such ads. In India, the Equal Remuneration Act 1976 implements the provisions of Article 39 of the Indian constitution and prohibits sex-based discrimination in the payment of salary for the same work (or work of similar nature) as well as in recruitment, promotion, training, and transfer. However, these provisions within the Indian legal framework are ineffective in preventing employers from posting job ads that explicitly request a male or a female.

[5] We search the text for the following words which indicate an explicit female preference: 'female', 'females', 'woman', 'women', 'girl', 'girls', 'lady' or 'ladies'. Similarly, we undertake a search for the following words which indicate an explicit male preference: 'male', 'males', 'man', 'men', 'guy', 'guys', 'boy', 'boys', 'gent' and 'gents'. Some job ads include words related to both genders in the job title or job description. We categorize such job ads as having no explicit gender preference, together with ads that did not include words related to either gender.

lower skill requirements in the Indian labor market are more likely to indicate a gender preference, consistent with the previous evidence from other countries. Within the set of gender-targeted ads, jobs having an explicit male preference are less likely to require higher education but are more likely to require more experience (> 2 years) and advertise higher wages than the jobs that explicitly prefer women.

Gender segregation at the application stage will occur if candidates comply with the gender requests. Using the applications data, we find that the fraction of female applicants to a vacancy increases by around 16 percentage points when the employer indicates an explicit female preference and reduces by 9 percentage points in case of an explicit male preference, even within the same occupation. These translate into an increase of 50% and a decrease of 28% in the share of female applicants to female and male-targeted job ads respectively, which are quite substantial. We also find that an employer's explicit preference for women significantly reduces the total number of applications to a job ad.

Even in the absence of an explicit gender request, the text contained in a job ad may send an implicit signal to a candidate regarding whether the employer posting the ad prefers a female or a male candidate for the job. To test this, we calculate how predictive the text in a job ad is of the explicit gender requests using a machine learning (ML) approach.[6] We find that for job ads that are not explicitly targeted at any gender, an increase in predicted femaleness (increasingly female job description i.e. use of words in the description which reflect job traits more likely to be associated with females) from zero to one (i.e., zero to a perfect female association) leads to a reduction in the advertised wages by 29% within the same occupation and location.

We then examine how the predicted gender mix of applicants varies with the associated femaleness and maleness within the same occupation. Two striking findings emerge. One, as femaleness associated with a job description increases in the gender non-targeted ads, the female application share increases. Second, even in the presence of an explicit request for women, the proportion of female applicants increases with femaleness.

To further delve into the source of such variation, we systematically uncover words reflecting gender stereotypes related to skills, job flexibility, and personality traits using recent research in explainable artificial intelligence. For instance, female skills related words include 'makeup', 'Tally', 'emailing', 'proofreading'—skills associated with stereotypical female tasks whereas male-related skills include knowledge of tools like 'OOP/Corel/API', 'driving/repairing', 'negotiation' and 'pitching'. Personality and appearance traits like 'pleasant'

[6] Kuhn et al. (2020) use the Naïve Bayes (NB) classifier, while we use the logistic regression (LR) classifier with balanced class weights and TF-IDF vectors. In our context, using the LR classifier almost doubles the accuracy.

and 'complexion' occur when requesting for women candidates while 'multitask', 'determination', 'vigilant' occur when requesting for men. Job flexibility is lower in jobs requesting men —since words like 'rotational/night/evening shifts', 'travel', 'relocate' for work are more common in ads requesting men. We also find that female skills and male job flexibility related words increase and decrease female applicant shares respectively. Moreover, jobs requiring skills associated with women entail a wage penalty, whereas those associated with men engender a significant wage premium.

Collectively, these findings indicate that gendered wording of job descriptions matters for female applicant shares and contributes to the gender wage gap in the Indian labor market. We are further able to detect sources of gender disparity by examining which stereotypes matter. Our results have implications for whether banning gender targeting in job ads can be an effective tool for narrowing gender segregation and eliminating the gender pay gap in the labor market. Given that female applicant shares are higher (lower) in gender non-targeted ads (male-targeted ads), restricting employers from stating their gender preference can reduce gender segregation at the application stage. However, it is unlikely to eliminate it as long as stereotypical gender associations exist in the labor market. Achieving greater gender parity in the labor market requires a broader transformation of traditional gender roles that affect applicant behavior and employers' preferences.

1.4 Group Size and Political Representation Under Alternate Electoral Systems

Representation of minorities in national governments ensures a more inclusive public policy and improves the allocation of public resources towards them. In contrast, their systematic exclusion is antithetical to democratic principles and may lead to conflict and instability in democracies (Cederman et al., 2010). Our data show that only about a third of minorities are politically represented on an average in national cabinets of post-World War II democracies. Many minority groups such as the Tamils in Sri Lanka, Kurds in Turkey, Palestinians in Lebanon, Bengali Hindus in Bangladesh, and Roma in several European countries—are still subject to targeted discrimination by their respective governments. In contrast, the majority group is represented in 94% of cases.

Chapter 4 deals with the question of political representation of minorities in democracies and asks why political parties represent some minorities but not others. Specifically, we examine how the group size of minorities affects their representation in national government and consequently, allocation of public resources under majoritarian (MR) and proportional

(PR) electoral systems. We propose a novel theoretical framework that models spatial distribution of *multiple* minority groups in a probabilistic voting setup. It predicts that a minority's population share has *no effect* on its representation and per capita resource allocation under PR, but has an inverted-U shaped relation under MR. We verify this by compiling an ethnicity level panel data set comprising over 400 minority groups across 87 countries during 1946–2013. Our main measure of political representation is an indicator that measures whether a group has any representation in the national government.[7]

MR elections are typically contested over single-member districts. The candidate or party with the highest number of votes or an absolute majority in an electoral constituency wins. MR is currently followed for legislative elections in countries such as India, Nigeria, and the United Kingdom. MR can result in a significant disparity in the vote share of a party across the country and their seat share in the parliament. Such disparity does not arise in PR in which seats are allocated to parties in proportion to their vote share in multi-member constituencies. Examples of countries that currently have PR are Argentina, Belgium, South Africa, and Israel. To illustrate, in the 2014 Indian general election conducted using the first-past-the-post system, the BJP polled 31.3% of national votes but won 51.9% seats. In contrast, the INC got 19.5% of national votes but could only obtain 8.1% of seats in the parliament. A more striking case is of BSP, which did not get a single seat in the Lok Sabha, despite polling 4.2% votes at the national level.

We face two challenges in isolating the effect of group size on political representation. First, a group's access to power may depend on a lot of country-level observable and unobservable factors. These include the number and size composition of all the groups including the majority, their political alliances, electoral strategies of political parties, voters' attitudes towards specific groups and political, economic or social contingencies that may affect representation in complex and unpredictable ways. Therefore, we compare minorities within a country in a given year. This allows us to remove all country-year specific factors statistically.

Second, the electoral system of a country is not random. Political actors in positions of power may strategically choose electoral systems that maximize their chances of winning (Trebbi et al., 2008). This means that the electoral system at the time of democratization, and even changes in it later may depend on existing power distribution across groups.

To address this non-randomness, we exploit the fact that the electoral system of the

[7] The Ethnic Power Relations core data set 2014 (Vogt et al., 2015) defines ethnicity as "any subjectively experienced sense of commonality based on a belief in common ancestry and shared culture." It provides a measure of political representation in the national government called "power rank" for every ethnic group in a country for every year. This measure can belong to one of six categories in descending order of power—monopoly, dominant, senior partner, junior partner, powerless, and discriminated by the state. We consider a group as represented if it is neither powerless nor discriminated by the state.

former colonial rulers systematically predicts the electoral system of the colonies but does not directly affect ethnic politics post-independence. Most of the countries that were once British and French colonies adopted MR. In contrast, those that had been colonized by Belgium, Netherlands, Portugal, and Spain adopted PR. Therefore, we use the electoral system of the primary colonial ruler in the colony's independence year as a natural experiment that quasi-randomly determines the electoral system of a colony.[8] We also use a second identification strategy where we compare the same ethnic group across countries within a continent.

We find that under PR, the group size of minorities does not affect their representation in the national government. In contrast, it has an inverted-U shaped effect under MR. Therefore, under MR, there is an "optimal" minority size above which representation begins to fall. We estimate the optimal population share to be 22%.

The explanation for the result we see in PR is straight forward. In PR, parties maximize votes share across the entire country. Consider two minorities with different sizes. Though, offering higher representation (and hence, per capita transfers) to the larger group gets a party more total votes, it is cheaper for a party to attract a higher share of voters from the smaller group. Since the cost of obtaining a vote from an additional person is the same irrespective of which minority the person belongs to, therefore, representations equalize across groups of different sizes.

The result under MR is driven by groups that are concentrated within a geographic region in a country. In MR, parties want to win constituencies and have to consider settlement patterns of groups across constituencies. Consistent with insights from the settlement scaling theory (Bettencourt, 2013), we show that the area occupied by a minority group has a concave relationship with its population share. This implies that, on average, larger minorities have higher population density than smaller ones. For minorities which are not dispersed throughout the country, there is an "optimal" density that maximizes their presence across districts. If a minority is too small, they become less important everywhere. If they become too large, their importance remains clustered around a few districts only. Thus, mid-sized groups enjoy the highest level of access to the national government in MR.

We use nightlight intensity as a proxy for public resource allocation towards groups in a country's geographically well-demarcated region. Electricity in most countries is publicly provided and is an essential public good for any region within a country. Michalopoulos and Papaioannou (2013a) use micro-data from Afrobarometer surveys to confirm that the

[8] Some countries, such as Indonesia and Brazil, became dictatorships after independence and remained so for many decades before democratizing. Eighteen countries democratized over 30 years after independence. Of them, 10 have PR, though only 2 were colonized by countries with a PR system. In such cases, the colonial ruler's electoral system matters much less for a country. Therefore, we omit countries that democratized more than 30 years after independence for our main causal estimates.

measure is indeed a good proxy for various public goods such as "access to electrification, presence of a sewage system, access to piped water, and education" within settlement areas of ethnic groups. In general, nightlights are also shown to be a good proxy for economic activity at the sub-national level (Henderson et al., 2012).

The average nightlight luminosity in an ethnic group's settlement area is measured by overlaying GIS maps of the settlement areas of ethnic groups using the GeoEPR dataset (Wucherpfennig et al., 2011) with DMSP-OLS Nighttime Lights Time Series data. The use of nightlight luminosity imposes two restrictions in the data—it is available only from 1992 onwards. Therefore, it can be used only for groups with a well-demarcated and contiguous settlement area as specified by the EPR dataset. The result for political inclusion is replicated with nightlight as the outcome variable. The estimated population share with peak nightlight intensity is also 22%—same as what we estimated for political inclusion. We find similar patterns when comparing the same group across different countries in a continent, as well as when we measure road construction in an ethnic group's settlement area as a measure of public resource allocation using geo-spatial data from the Global Roads Inventory Project (Meijer et al., 2018).

The possibility of a change in electoral systems is already being examined for many countries. We provide evidence that the PR system may reduce between-group inequality in political representation of minorities and in their material well-being. Our findings are consistent with other comparative studies that show PR to be associated with lower ethnicization of politics (Huber, 2012) and a lower likelihood of electoral violence (Fjelde and Höglund, 2016).

Chapter 2

The Importance of Being Earnest: What Explains the Gender Quota Effect in Politics?

2.1 Introduction

There is now an extensive literature that documents whether gender of a political leader affects policy in India (Chattopadhyay and Duflo, 2004; Bardhan et al., 2005; Ban and Rao, 2008; Bardhan et al., 2010; Clots-Figueras, 2011, 2012; Gajwani and Zhang, 2014; Deininger et al., 2015) as well as in other countries.[1] Several of these papers use affirmative action policies for women, in the form of gender quotas in political positions, to identify the causal impact. These studies are primarily motivated by the observation that preferences for public goods are gendered, i.e., women and men prefer different public goods. Hence, gender quotas should change policy, as elected women would prioritize allocation of a different set of public goods than their male counterparts.

However, the theoretical models of political economy tell us that policy is shaped by the preference of the elected leader (Alesina, 1988) as well as the preferences (Downs, 1957) of and communication of these preferences (Banerjee and Somanathan, 2001) by the voters. Importantly, female voters may express the demand for their preferred policy more vigorously in the presence of female leadership (Chattopadhyay and Duflo, 2004; Beaman et al., 2007). This can happen because registering demand to elected authorities is a costly process, especially for women who face a significantly higher cost of political participation and engagement. Consequently, women may feel that their demands may be heard more

[1] We discuss this literature later in this section.

in the presence of female leaders (Iyer and Mani, 2019). Hence, gender quotas can lead to changes in policies either due to differential prioritization of policies by the female leader (i.e., supply side factors) or due to differential *demand* (from the female voters) faced by her.[2] Understanding the source of the effect is important because it has implications for policies to improve the substantive representation of women. If demand matters, then policies that empower female voters and facilitate their political participation can also contribute towards this goal. It would also imply that gender quotas may be more effective in the presence of more engaged female voters.

The two mechanisms, however, are hard to disentangle as they co-move. When a female leader is elected, both the demand and supply mechanisms come into force. In this chapter, we overcome this identification challenge by compiling data on allocation of a publicly provided good that is targeted at the *household* level. We identify salient population groups that systematically differ in their gender gap in preference for the good and examine whether female quota leads to differential allocation of the same good across households belonging to these groups. This allows us to isolate demand from supply. Heterogeneous effect of gender quota across households is by itself not an evidence of demand. The female leader could be differentially focused on certain types of households, for electoral or other reasons, which may cause heterogeneous allocation. As discussed below, our analysis is robust to this criticism.

We examine the effect of female quotas in the 2015 village council president (or, *Sarpanch*) elections for the universe of village governments or *Gram Panchayats* in Uttar Pradesh (UP), the largest state in India (and comparable to the fifth most populous country in the world) on the provision of household toilets in the subsequent year, 2016–17. We use household toilets as our target public good due to a combination of factors. First, it is well known that women have a greater preference for toilets than men (Coffey et al., 2014; Khanna and Das, 2016; Stopnitzky, 2017). Hence, we expect the gender quota to have an effect on its allocation. Moreover, open defecation is widely practiced in rural India due to severe lack of sanitation facilities. Consequently the construction of toilets has recently received a lot of emphasis from the policymakers, making it a politically salient public good. Finally, since toilets are allocated to individual households, unlike many other publicly provided goods, we can identify the demand and supply mechanisms.

We use the allocation of household toilets under the the *Swachh Bharat Mission* (SBM)

[2] Gender quotas can also lead to different policies because female politicians could be less experienced (Gajwani and Zhang, 2014) or due to male backlash (Gangadharan et al., 2016). However, such differences tend to be temporary in nature, as women politicians have been shown to catch up very rapidly in such contexts (Afridi et al., 2017). We also show in our context that women politicians coming through gender quotas do not have differential ability in governance.

or 'Clean India Mission', launched by the Government of India in 2014 as our outcome variable. The algorithm used to assign female quotas in the 2015 Sarpanch elections allows us to employ a fuzzy regression discontinuity design (RDD) to estimate the causal effect of the quota.[3] We compile a household level dataset on allocation of toilets for the entire rural population of UP (over 25 million households) and match it to the election results for the universe of about $59,000$ GPs of the state. We find that, on average, gender quotas in 2015 lead to 15–28% increase (depending on bandwidth) in the probability of allocation of toilets to eligible households in 2016–17. The average effect, however, is imprecisely estimated. We then argue and show that the average noisy effect masks significant heterogeneity across the Gram Panchayats (GPs).

To uncover the heterogeneity, we first notice that the gender gap in the preference for toilets is significantly higher among Muslims than Hindus. This is due to religious considerations of purity and pollution that Hindus associate with in-home toilets (Coffey et al., 2014, 2017). Such purity concerns are absent among Muslims, leading to a greater gender gap in preference for household toilets.[4] Using survey data on usage of and preference for toilets in rural North India (including UP), we find that the gender gap in preference is indeed significantly greater among Muslim households compared to Hindu households. We divide the GPs by the Sarpanch's religious identity, i.e., Hindu and Muslim, and separately estimate the impact of gender quota within each sample. While the samples of GPs with either Hindu or Muslim Sarpanches are clearly endogenous, our identification method works in each of the sample separately, producing causal effects of gender quotas in the two samples. We discuss this in detail in Section 2.5.4.

We argue that if the preferences of either voters (demand) or Sarpanches (supply) at all matter for toilet provision, then we should find a larger effect of the quota among Muslim Sarpanches compared to Hindu Sarpanches. This greater effect could either result from the larger gender gap in preferences among Muslim Sarpanches or a greater demand from women in the Muslim households, who are relatively numerous in GPs where Muslim leaders win. Consistent with this, we find that among Hindu Sarpanches, gender quota has no average effect on toilet construction. On the other hand, gender quota within Muslim Sarpanches increases toilet construction by 15 percentage points on a sample average of 13%, i.e., allocation probability more than doubles.

To identify the demand mechanism, we examine heterogeneity in the treatment effect

[3] We discuss the quota assignment algorithm in Section 2.2.2 and the associated identification strategy in Section 2.5.1.

[4] Another possible consideration among Muslim households might be greater restrictions on the mobility of women compared to Hindus, and consequently, restrictions on Muslim women to defecate in the open.

across households headed by men and women.[5] For this, we use data on the gender identity of the household head, available for a subset of households that were eligible for the scheme. We argue that a household may register a greater demand for a toilet if it is headed by a woman. We show using pan-India survey datasets, that women members in female headed households exercise greater decision-making power within the household, enjoy greater autonomy in public participation, and indeed, participate more in the local political activities within the village. Therefore, we hypothesize that if demand is important, then the gender quota effect among either Hindu or Muslim Sarpanches is likely to be greater for households headed by women. Consistent with this, we find that gender quota among Hindu Sarpanches leads to a differentially large increase in the probability of toilet allocation to *Muslim* female headed households.[6] The Hindu female headed households also experience a differentially greater increase in the allocation probability. The effect, however, is smaller in magnitude and statistically imprecise. The observation that Hindu female Sarpanches allocate more toilets to Muslim female headed households, relative to Hindu ones, establishes the demand mechanism. It also rules out the possibility that female Sarpanches being more concerned about (or focused on) female headed households in general drives the result.[7] Since female headed households comprise only 4% of eligible households, the average gender quota effect among Hindu Sarpanches becomes small and statistically insignificant.

Among Muslim Sarpanches, we find that gender quota differentially increases toilet provision to *both* Hindu and Muslim female headed households, but the effect is larger in magnitude and is statistically significant for the latter. These results show that household demand is an important factor in explaining heterogeneity in gender quota effect across GPs. Moreover, we show that the result is stronger in GPs with high population share of Muslims, i.e., in GPs that are likely to have a higher collective demand for toilets. Therefore, we find that the effect of household demand on the likelihood of it receiving the good partly depends on the overall level of demand in the population. Additionally, we examine the effect of elect-

[5] We do not examine heterogeneous treatment effect across Hindu and Muslim households, as Muslim Sarpanches allocating greater number of toilets to Muslim households can be due to own-group favoritism–a supply side consideration. Moreover, both Hindu and Muslim Sarpanches could have electoral motives to allocate more toilets to Hindus, the majority group in both samples, which may lead to underestimation of the demand effect.

[6] We identify the religion of households by using a machine learning algorithm developed by Chaturvedi and Chaturvedi (2020) that predicts with high degree of accuracy the religion of a household based on the name of the household head. We discuss this in further detail in Section 2.4.4.

[7] The observation could be consistent with Muslims being over-represented among the *eligible* female headed households, i.e., those that did not own toilet before 2016—17. Any emphasis by Hindu female Sarpanches to provide toilets to female headed households then may imply greater benefits going to Muslim female headed households. However, we do not find support for this argument in the data. The share of Muslim households among eligible female headed households is 11.76%, while that in the entire population of female headed households is 13.71%.

ing female Sarpanches in open elections (i.e., elections without female quotas) using close elections between male and female candidates and find consistent patterns.

To identify and estimate the supply mechanism, we focus on GPs that had a close election between a Hindu and a Muslim candidate (both in GPs with and without female quotas). Within this sample, the religion of the Sarpanch is, in effect, randomly assigned. Hence, if we compare the GPs with Hindu and Muslim Sarpanches within this sample, they would on average, have the same demographic composition and voter preferences. Consequently, the demand effect induced by the female quota would, on average, be same across these two samples. Therefore, any differential female quota effect in the GPs with Muslim Sarpanches (relative to Hindu ones) would be evidence of the supply mechanism. We, however, find a null effect, implying that the supply mechanism is not present in our context. We, therefore, conclude that the overall effect of gender quota and more importantly its heterogeneity across villages can be explained exclusively by demand side factors.

The pioneering study by Chattopadhyay and Duflo (2004) that led to the emergence of this literature, considers both demand and supply mechanisms.[8] They indeed find that women voters express their preferences more under female leadership, suggesting presence of the demand channel. They, however, do not examine—possibly due to identification challenges—whether greater voicing of preferences by women voters causally influenced the treatment effect. Their results show that differential preference of female Sarpanches is likely to be the primary driver of the differential allocation of public goods. Subsequent papers in this literature do not explore these mechanisms.[9] We contribute to this literature by providing an identification strategy for estimating the demand effect and empirically demonstrating its importance.

While the literature emerged from the Indian context, several papers estimate the effect of female leaders on policy in other countries (Van der Windt et al., 2018 (Congo); Devlin and Elgie, 2008 (Rwanda); Franceschet and Piscopo, 2008 (Argentina); Braga and Scervini, 2017 (Italy); Campa, 2011 (Spain); Ferreira and Gyourko, 2014 (USA)), as well as using cross-country comparisons (Dollar et al., 2001; Swamy et al., 2001; Barnes and Burchard, 2013; Hicks et al., 2016; Bhalotra et al., 2020). The literature, however, finds mixed evidence on the presence of gender quota effect, both in India as well as internationally. Bardhan et al. (2010), for example, do not find that female quotas lead to differential public good provision in GPs in West Bengal.[10] Our results help reconcile the mixed findings in the literature

[8] They make distinctions between two related but separate supply mechanisms—female leaders implementing their differential preferences and female leaders being more responsive to female voters' preferences. They however do not find support for the second mechanism.

[9] Beaman et al. (2007) point out the possibility of demand effects, but do not test for it.

[10] Internationally, Van der Windt et al., 2018; Devlin and Elgie, 2008; Campa, 2011; Ferreira and Gyourko,

by showing explicitly that variation in demand across regions within the same state can critically shape the effect of gender quota on a politically salient good for which preference is starkly gendered.

2.2 Background

2.2.1 Swachh Bharat Mission

Sanitation policies have been around in India since 1986 when the Government of India launched the Central Rural Sanitation Program. The program was rechristened as the Total Sanitation Campaign (TSC) in 2000 and again as Nirmal Bharat Abhiyan (NBA) in 2012. These efforts were largely unsuccessful in achieving the desired reduction in open defecation rates in India. The proportion of households having toilets stood at 37% in 2001 which increased to 47% in 2011 (Census 2001, 2011). In rural India these numbers stood at 22% and 30.7% in 2001 and 2011 respectively—a paltry increase of 8 percentage points over a decade. The state of Uttar Pradesh, the context of our study, fared worse than the national average, with 19% and 22% rural households having access to toilets in 2001 and 2011 respectively.

In the most recent efforts to increase access to toilet access for households, Swachh Bharat Mission (SBM) or the 'Clean India Mission' was launched in October 2014 by the newly formed central government, led by the Bharatiya Janata Party (BJP), that came to power in May 2014. Providing access to toilets was one of the important policy platforms of the BJP during its election campaign. Therefore, after winning the elections, the government allocated substantial resources under the SBM towards construction of household and community toilets across the country. We specifically examine toilet construction in rural Uttar Pradesh (UP) because the lack of access to toilets and the incidence of open defecation is primarily concentrated in rural India, and especially so in UP.[11] The main thrust under the rural component of SBM was to provide subsidy towards the construction of household toilets. A subsidy of ₹12,000 was provided for construction of each Individual Household Latrine (IHHL).[12] A baseline survey was conducted in 2012 by the Government of India

2014 find null effects while Franceschet and Piscopo, 2008; Braga and Scervini, 2017; Bhalotra et al., 2020 find that women leaders indeed pursue different policies.

[11] According to Census 2011, 92 percent of households without access to a toilet or latrine were rural. Similarly, the Indian Human Development Survey (IHDS) 2011–12 reports that 90 percent of households practicing open defecation live in rural areas.

[12] The SBM promotes construction of twin-pit structures for toilets so that the feces decompose by themselves and no frequent manual cleaning of fecal sludge is required. This subsidy, on an average, is sufficient to cover all costs related to a twin-pit toilets in India. However, if a household wanted it could build a better quality toilet. The subsidy amount would remain unchanged.

to identify households without toilets and determine the household eligibility for subsidy towards construction of IHHL.

The subsidy was provided in the form of a reimbursement, which a household could apply for after initiating toilet construction. It was paid in two installments of ₹6,000 each. The first installment was paid when the household reported that a pit was dug and filled out an agreement form. The second installment was paid after completion of toilet with the structure and submission of a completion form. These forms had to be submitted to the District Panchayat Office. All the households below the poverty line (BPL) were eligible for this subsidy. Among the above poverty line (APL) households, those belonging to Scheduled Castes, Scheduled Tribes, small and marginal farmers, landless laborers with homestead, physically handicapped and women headed households were eligible for subsidy.

Data from the Ministry of Drinking Water and Sanitation show that there has been a steady expansion of toilets since the start of the program in 2014. The proportion of rural households having toilets has increased from 38% in 2014 to 84% in 2018. In UP, the same has increased from 30% in 2014 to 66% in 2018. We comment on this data and its quality in greater detail in Section 2.4.2.

2.2.2 Gram Panchayat

Structure and Responsibilities

The village councils or Gram Panchayats (GPs) are the lowest tier of the local governance structure in rural India. The council in a GP consists of elected members and a president or head of the council, who is known as the Sarpanch of the GP. The members are elected from individual wards within a GP, and the Sarpanch is either directly or indirectly elected depending on the state. In UP the Sarpanch is elected directly following a presidential system. The council is responsible for provision of local public goods such as hand pumps, toilets, local roads etc. and employment under public works. Moreover, the Sarpanch within the council enjoys substantial executive power in deciding the public good priorities of the council and its overall expenditure pattern (Besley et al., 2004; Das et al., 2017; Gulzar et al., 2020).

The GPs played a pivotal role in the implementation of the SBM program. While the program implementation was monitored at the district level by the district magistrate, the primary implementing agency at the village level was the GP. The council was in charge of the identification of potential beneficiaries, fund flow and maintenance of records, mobilization of demand for construction of toilets, actual construction of toilets, and social audits.

Elections

The GP elections in Uttar Pradesh were held during November–December 2015. We focus on the election of Sarpanches in the GPs, as they are the key decision-makers within the council. The Sarpanch elections are subject to affirmative action policies or quotas for various caste (or ethnic) groups as well as for women. This is known as the "reservation policy" which sets aside a certain number of Sarpanch positions within each block for three disadvantaged caste groups —Scheduled Tribes (STs), Scheduled Castes (SCs) and Other Backward Classes (OBCs), where only members of these groups can run as candidates in the elections. If a GP is not reserved for any caste group then we refer to it as an unreserved GP. Within each of the four categories of GPs (reserved for ST, SC, OBC and unreserved), at least one-third are again reserved for women. If the Sarpanch position in a GP is reserved for SC and woman, for example, then only female SC candidates can run in the Sarpanch election in that GP. For unreserved GPs that are reserved for women, any woman can contest the election.

Due to the delimitation of constituencies of GPs based on the 2011 census, the assignment of SC, ST and OBC reservation, as well as female reservation for the 2015 elections was done afresh, disregarding the allotment status in the previous elections.[13] Unlike some other states, the allocation of reservation positions across GPs in UP is not randomized but is based on a deterministic algorithm that we describe below.

Caste Reservation

The reservation status of GPs in the 2015 Sarpanch elections in UP was decided separately for each block. The proportion of Sarpanch positions to be reserved for a caste group in a block equals the rural population share of that caste group in that block. The number of GPs to be reserved for a caste group in that block is then determined by rounding off the the calculated share. The GPs within a block were first arranged in the descending order of the ST population share of the GPs. Then the top ranked GPs were selected to be reserved for ST, where the number of reserved GPs is given by the rounded off number calculated in the previous step. Then the remaining GPs within the block were arranged in descending order of their SC population shares. Again, the top ranked GPs were selected to be reserved for SC. For the remaining GPs within the block, OBC reservation status was decided using the same procedure.

[13] This was in accordance with the 10th amendment of Uttar Pradesh Panchayat Raj (Reservation and Allotment of Seats and Offices) Rules 1994.

Gender Reservation

At least a third of Sarpanch positions allocated to every caste group in each block are reserved for females of that caste. For this, the GPs reserved for a caste group are ranked in descending order of the population share of that caste. Top one-third of these GPs, rounded off to the higher integer, are reserved for women. The unreserved GPs in a block are listed in descending order of the "general" category population in the GPs.[14] We implement the caste and gender reservation algorithms using the population figures from the 2011 census, the same figures that the state government officials used. We are able to correctly predict the allotment of female reservation in almost 98 percent of GPs using the above algorithm.

2.3 Preference for Household Toilets

We use toilet allocation as our primary policy outcome. In this section, therefore, we discuss the existence of a gender gap in preferences for toilets and systematic differences in the gender gap in preference across well-identified population groups —across religions, i.e., Hindus and Muslims, and across male and female headed households.

2.3.1 Gender Gap in Preferences

We first argue that women have a greater preference for household toilets than men (Coffey et al., 2014; Khanna and Das, 2016; Stopnitzky, 2017). This is motivated by the observation that notions of shame and dignity are associated more strongly with women, especially in rural India, leading to their lower preference for defecating in the open. Men, on the other hand, consider defecating in the open a sign of strength and masculinity. Moreover, women face a greater risk of being harassed or attacked while defecating in the open (Mahajan and Sekhri, 2020). Additionally, menstruating women may prefer to use toilets due to hygiene concerns. Coffey et al. (2014) show direct evidence of gender gap in toilet preference by capturing within household variation in open defecation rates in the SQUAT survey (2014).[15] They find that among the households owning a latrine, men are twice as likely to defecate in the open than women. This suggests a lower revealed preference for open defecation among women. We validate this using the SQUAT survey data and find that women are 9

[14] The "general" category is the group of upper castes, i.e., those who are not STs, SCs, or OBCs.

[15] The SQUAT survey was carried out in rural areas of northern states of India, namely in Bihar, Uttar Pradesh, Haryana, Madhya Pradesh and Rajasthan, and captures information about household ownership and individual usage of toilets. The states in the survey are culturally similar in terms of their gender attitude and include the state that we study. The dataset is publicly available here: https://riceinstitute.org/data/2014-and-2018-rural-sanitation-surveys/.

percentage points more likely to use a toilet, conditional on the household owning one, even within the same household (Appendix Table 2.5: Panel A, columns (1) and (2)). We also find that for the sample of households that do not own toilets, women report having higher preference for toilets than men (Appendix Table 2.5: Panel A, columns (3), (4) and (5)). We discuss the detailed results in Appendix Section 2.A.1.

2.3.2 Gender Gap in Preferences across Religions

Having established the existence of gender gap in the preference for household toilets, we now argue that this gender gap is larger among Muslims than Hindus. This observation is motivated by existing evidence in the literature. For instance, Coffey et al. (2017) discuss the reasons for high rates of open defecation among Indian rural population than Sub-Saharan countries, despite higher per capita incomes, education and water availability.[16] Their findings show that cultural factors affecting notions of 'purity and pollution' among the Hindus, associated with defecating within a house, are an important factor for higher open defecation rates in India. These concerns, existing among both Hindu men and women, increase their alignment on the issue of toilet ownership and partially narrow the gender gap in preferences among Hindus. Such purity concerns are absent among Muslims.[17] Consequently, they have a higher likelihood of owning and using toilets than Hindus, in spite of being poorer than Hindus on average. Coffey et al. (2017) report that only 4% of rural Hindu households used inexpensive pit latrines, compared to 15% rural Muslim households. Data from the National Family Health Survey (NFHS) 2015–16 also show that Muslim households are 21% less likely than Hindu households to defecate in the open. In fact, Geruso and Spears (2018) argue that differential sanitation practices of Muslims and Hindus can explain the longstanding puzzle that in India, "[...]Muslim children are substantially more likely than Hindu children to survive to their first birthday, even though Indian Muslims have lower wealth, consumption,

[16] In 2015, the proportion of rural population in India defecating in the open stood at 43% while this figure was 32% for rural Africa and the world average in rural areas was 20%. The figures in rural parts of comparable economies like Bangladesh stood at 1.7% and China at less than 1% (World Development Indicators).

[17] Alternatively, Muslim women face greater restrictions on mobility than Hindus, resulting in greater demand for household toilets from them. We find evidence of this claim in the National Family Health Survey (NFHS) 2015 that asked women if they can visit various public spaces in their village on their own. Muslim women are 7 percentage points less likely to say yes, relative to Hindu women (Appendix Table 2.6: Panel B, column (2)). Often these mobility restrictions are imposed by Muslim men living in the same household. This could decrease the gender gap in preference for toilets among Muslims if Muslim men believe that construction of toilets within households can improve compliance with the mobility restrictions. Using SQUAT survey, we find that purity concerns dominate and that gender gap is lower between Hindu men and women. Though Muslim men are more likely to use household toilets than Hindu men conditional on ownership, they are not more likely than Hindu men to report household toilets among their top 3 priorities (Appendix Table 2.5: Panel B, columns (1), (3), (4), (5)).

educational attainment, and access to state services."

We again validate this claim using the SQUAT survey data. We find that among households owning toilets, the gender gap in usage of toilets (i.e., in revealed preference for toilets) is 50% higher among Muslims compared to Hindus within the same village (Appendix Table 2.5: Panel B, columns (1) and (2)). Revealed preference indicates the absolute importance attributed to household toilets by each subgroup. The usage patterns confirm the role played by cultural factors in determining the importance given to household toilets by Muslim and Hindu women vis à vis men. Among households that do not own toilets, gender gap in preference for toilets is also significantly higher among Muslims (Appendix Table 2.5: Panel B, columns (3), (4) and (5)). Specifically, Muslim women are 22–24 percentage points more likely than Muslim men to report household toilets among their top 2 and top 3 priorities (p-values 0.055 and 0.053 respectively) while we see no difference in stated preference among Hindu women and men. Appendix Section 2.A.2 discusses these results in detail.

2.3.3 Demand across Male and Female Headed Households

Toilet is a household level public good, even though preference for it differs starkly across male and female members of the household. Therefore, conditional on the preferences of the members of a household, whether the household publicly expresses a demand for toilets to the local government may depend on who is at the helm of the household. For instance, Coffey et al. (2014) discuss that young women are likely to have the largest demand for latrine use but have the least decision making power within the household. Therefore, the capacity or willingness of a household to register its demand may depend on the relative bargaining power and autonomy of women within the household as well as their political participation.[18] We argue that, conditional on the religious identity of a household, which captures the gender gap in preference within the household, it will publicly express greater demand for a toilet if it is headed by a woman. This is because, women belonging to a female headed household enjoy more decision-making power and autonomy as well as participate more in the local political activities.

We test this claim using two nationally representative survey datasets, namely the National Family Health Survey 2015 (NFHS-4) and Rural Economic and Demographic Survey 2006 (REDS) data. We show that women from female headed households are significantly more likely to make household decisions on their own, venture out alone in public spaces, attend village meetings, and participate in activities of the political parties. We report the

[18] Bargaining power within the household may be important in our context, since household often has to begin construction, i.e., pay an upfront cost, before seeking the subsidy for toilet under the SBM scheme, as described in Section 2.2.1.

results in Appendix Table 2.6: Panel A and discuss them in Appendix Section 2.A.3. Moreover, this difference between female and male headed households exists across both Hindus and Muslims (Appendix Table 2.6: Panel B).[19]

2.4 Data

2.4.1 GP Election

Detailed results for the 2015 GP head (or *Sarpanch*) election come from the State Election Commission of Uttar Pradesh. We focus on the 2015 election since this is the first election after the current central government came into power in 2014 and gave the sanitation campaign a major push by launching the SBM program in October 2014. There are over $59,000$ GPs in UP and about $470,000$ candidates contested the Sarpanch election in 2015. The election data contains information on candidate characteristics, such as their gender, caste etc., as well as the votes received by each of them. The election data also provides the reservation status of the Sarpanch position for each GP.

2.4.2 Toilet Construction

We compile household level toilet construction data using detailed information available on the official website of SBM maintained by the Ministry of Drinking Water and Sanitation, Government of India.[20] We scraped the data during March to April, 2018. The website provided the full list of households living in each GP and for each household, it indicated whether the household had a toilet at the time of the baseline survey, carried out in 2012. It subsequently tracked each household from 2013–14 onwards, and indicated whether it had a toilet at the end of each financial year.[21] This information is given along with the name of the household head, the name of the parent or spouse of the head, and certain characteristics of the household such as whether it was Below Poverty Line (BPL), landless, whether the household head is a woman etc. For our purpose, we focus on toilet construction in the year 2016–17, the year following the Sarpanch elections in UP. This was the first financial year

[19] There is mixed evidence on whether female headed households (FHH) in rural India are poorer than male headed households. While Dreze and Srinivasan (1997) find no evidence that this is true, Gangopadhyay and Wadhwa (2004) show using more recent data that FHH are indeed poorer. We discuss the relevance of this issue for our results in Section 2.6.2. Dreze and Srinivasan (1997) also report that FHH have smaller household size than average. If it is easier to make a collective decision when household size is small, then this can also contribute to FHH expressing greater demand for toilet under female Sarpanches.

[20] The official website of SBM is `https://swachhbharatmission.gov.in/`.

[21] From 2019 onward, the website only provided the list of households that received toilets in a year, i.e., it stopped showing the full list of households.

post the GP elections in UP. Moreover, in 2017, the ruling party at the center (the BJP) won the UP state election. The state government subsequently heavily pushed the SBM scheme in the state, especially in areas where construction was lagging, since the target set by the central government was to reach 100% toilet access by the end of 2019. Hence, toilet construction during 2016–17 is more likely to reflect political will of the Sarpanch and household demand in the GP, rather than being driven by the policy priorities of the state government. In rural Uttar Pradesh, the proportion of households having a toilet increased from 32% in 2012–2013 to 37.5% in 2015–2016 and further to over 44% in 2016–2017 (i.e., by 7 percentage points in just one year post the launch of SBM).

Data Quality: We provide evidence to show that the administrative data is of good quality for the period we consider. First, Mahajan and Sekhri (2020) show that the correlation between district level toilet coverage in the administrative data for 2015–16 with toilet ownership reported by households in NFHS-4, also conducted in 2015–16, is fairly high (0.70). We plot the proportion of households in a district with toilet access from the two data sources for the state of Uttar Pradesh in Appendix Figure 2.3 and demonstrate a similarly high association between them in our context.[22] In another study, Gupta et al. (2019) resurveyed a subset of households from the SQUAT survey of 2014 in 2018, to examine toilet construction under SBM and changes in open defecation during 2014–18.[23] They report that 74% of households in rural UP had toilets in October, 2018.[24] According to the administrative data, 64% of households in rural UP had toilets at the end of 2017–18, i.e., by March 2018. These figures are highly comparable. Moreover, they also observe a large drop (26 percentage points) in open defecation during 2014–18, and find that "nearly the entire change in open defecation between 2014 and 2018 comes from increases in latrine ownership, rather than from changes in behaviour." This further validates the administrative data and confirms that toilet construction under SBM in the initial years led to perceptible changes in the practice of open defecation in these states.[25]

[22]One source of noise is the fact that NFHS proportions are estimates. Moreover, there is a small difference in coverage periods for the two data sources. The administrative data gives coverage at the end of fiscal year 2015–16 while NFHS reports an average over 2015–16. Thus, the official data is likely to report slightly higher coverage on an average.

[23]The 2018 resurvey was carried out in the states of Rajasthan, Madhya Pradesh, Uttar Pradesh and Bihar.

[24]This is an estimate for the entire rural UP arrived at using census population weights. The average estimate for the four states in the survey is 71%.

[25]It is possible that the administrative data is potentially manipulated in 2019, as there was an emphasis from the government to attain 100% toilet coverage by the end of 2019. Our period of study, 2016–17, however, is safely removed from such manipulation concerns.

2.4.3 Census

We complement the data on election results and toilet construction with the Census 2011 village population and amenities data. The census villages are smaller geographic units than GPs. For instance, in UP there are about $106,000$ census villages and $59,000$ GPs. We use mapping between census villages and the GPs prepared by the Ministry of Panchayati Raj, Government of India to construct GP level figures from the census data.[26] Moreover, the state government of UP shared with us the caste group wise population (for General caste, OBCs, SCs, and STs) at the GP level for the entire state.

2.4.4 Religion Identification in Rural U.P.

We wish to identify preference for toilets using religious identity of beneficiaries and politicians. We focus on Hindus and Muslims, who are the primary religious groups in UP comprising more than 99% of the rural population (Census, 2011). The data on election results and the toilet construction do not provide information on religion. For this reason, we identify the religion using the names of household heads and candidates in village elections. We take advantage of the fact that the names of Hindus and Muslims in India are distinct. For example, Bhalotra et al. (2014) and Heath et al. (2015) infer religion of electoral candidates in India from their names. Therefore, we manually classify religion of candidates as Muslim or non-Muslim (which we refer to as Hindu) based on their names and the names of their parent or spouse.

For identifying the religious affiliation of the households, we use the names of the household heads and their fathers' or spouses' names. There are about 25 million households in rural Uttar Pradesh. Therefore, manual classification is not feasible. To overcome this, we use a new and highly accurate algorithm proposed by Chaturvedi and Chaturvedi (2020) who infer religion from names using character sequence based machine learning models used in Natural Language Processing (NLP).[27] Character sequence based models have the advantage that they can exploit differences in linguistic origins of the two religions, and hence can classify unseen names with a high degree of accuracy. In contrast, dictionary based methods which use string matching to identify religion can only classify names that exist in a predefined name list.

Within a random sample of manually annotated 20,000 households in rural UP, the model correctly identifies over 97% of true Hindus as well as true Muslims. Another test of the algorithm's accuracy is the correlation between the predicted Muslim household share at

[26] The mapping is accessed from the website www.lgdirectory.gov.in maintained by the ministry.

[27] For our purpose, we use a deep learning architecture which combines convolutional neural network with long short-term memory recurrent neural network.

the sub-district or *tehsil* level in our data with that of Census 2011 population share of Muslims.[28] Figure 2.4 shows the relation between Muslim population share and the Muslim household share estimated by the algorithm at the tehsil level. The correlation between the two is 0.978. This shows that the model predicts religious affiliation of households very well in our data.

2.5 Overall Effect of Gender Quota

2.5.1 Identification

We use a fuzzy regression discontinuity design (RDD) strategy to find the causal effect of female reservation on construction of toilets. As discussed above, no less than one-third of the Sarpanch positions reserved for each caste group in every block are reserved for women of that caste. The procedure for allotment of gender quota described in Section 2.2.2 creates discontinuities in the mapping between a GP's rank in the ordered list and its female reservation status. We exploit these discontinuities to estimate the overall gender quota effect since the gender reservation status of the GPs is essentially randomized around the discontinuity for each caste group in each block. We define the running variable $X^c_{g,b}$ in GP g in block b reserved for caste group c in the following manner:

$$X^c_{g,b} = \frac{Pop^c_{g,b} - Pop^c_{threshold,b}}{\sigma^c}$$

where $Pop^c_{g,b}$ is the population share of caste c in GP g in block b, where $c \in \{SC, ST, OBC\}$. $Pop^c_{g,b}$ is the total general category population in GP g in block b for $c \in \{Unreserved\}$. The threshold value of caste c in block b is given by $Pop^c_{threshold,b}$. It is the mean of the lowest $Pop^c_{g,b}$ at which the GP Sarpanch position should be reserved for a woman of caste c in block b and the $Pop^c_{g,b}$ of the next GP in the ranking for caste c within the block. Obtaining the threshold by taking the mean in this way is standard in political economy literature using RDD (see, for example, Hyytinen et al., 2018). The denominator (σ^c) is the standard deviation of $Pop^c_{g,b}$ across the entire state. We follow this procedure to generate the running variable for GPs reserved for each caste category, i.e. ST, SC, OBC, as well as for the unreserved GPs. Since each GP can only be reserved for one caste or remain unreserved, this procedure gives a unique running variable for each GP depending on its caste reservation status.

[28]Sub-district or tehsil is the lowest geographic unit for which religion wise population figures are available in the 2011 Indian Census.

McCrary test: We test for manipulation at the threshold using the test proposed by McCrary (2008). This tests for the null hypothesis that the density of the underlying running variable that defines the assignment at the discontinuity is continuous at the cutoff, against the alternative of a jump in the density function at that point, which can reflect manipulation in treatment assignment. Figure 2.1 shows that there is indeed no jump in the density of the running variable at the threshold.

Balance test: We perform balance tests on a large number of covariates around the discontinuity to test the validity of the RDD approach. The results for the balance tests performed for bandwidth 0.1 are shown in column (1) of Table 2.9: Panel A.[29] We compile a large number of development indicators using the census 2011 village amenities dataset and also test for balance in proportion of Muslim population and female headed households. As column (1) shows, there is no discontinuity in these covariates at the threshold.

Moreover, we also test for discontinuity in toilet provision at the threshold in the three years before the 2015 Sarpanch elections. Column (1) of Table 2.9: Panel B reports the results. We notice that the proportion of uncovered households that were provided a toilet in 2013–14, 2014–15 and 2015–16 does not change discontinuously at the threshold. These results demonstrate that the female reservation assignment was indeed exogenous at the threshold value of the running variable.

2.5.2 Estimation Strategy

We use the running variable defined above to predict the treatment status of a GP in the first stage. The treatment variable ($Q_{g,b}$) equals one if the GP has a quota for women. The assignment variable ($A_{g,b}$) takes the value one if the female reservation algorithm predicts that the Sarpanch position should be reserved for a woman ($X^c_{g,b} \geq 0$) and 0 otherwise ($X^c_{g,b} < 0$). We restrict the sample to GPs where the running variable is near zero, i.e., $X^c_{g,b} \in [-t, t]$ for some small $t > 0$. We use the following specification:

$$Y_{h,g,b} = \alpha_0 + \tau Q_{g,b} + \alpha_1 X_{g,b} + \alpha_2 X_{g,b} * A_{g,b} + u_{g,b} \tag{2.5.1}$$

where $Y_{h,g}$ is a dummy variable that takes value one if a household h in GP g received a toilet during 2016–2017 under the SBM scheme, and zero otherwise. Our sample is the set of households which did not have toilets at the end of 2015–2016 and were eligible for the program. The treatment variable ($Q_{g,b}$) is instrumented with the assignment variable ($A_{g,b}$)

[29] These balance tests hold at smaller bandwidths as well. We discuss bandwidth selection in Section 2.5.2.

in the following first stage equation:

$$Q_{g,b} = \beta_0 + \gamma A_{g,b} + \beta_1 X_{g,b} + \beta_2 X_{g,b} * A_{g,b} + \epsilon_{g,b} \tag{2.5.2}$$

We use the number of eligible households in a GP as weights in our regressions to give equal consideration to all GPs in the household level data. The standard method used to select bandwidths in RDD is the one proposed by Calonico et al. (2014). However, the "CCT bandwidth" is 0.475 for the second stage, which is particularly wide as it includes about 64% of GPs in the state. We, therefore, estimate the results using three manually chosen narrower bandwidths—0.100, 0.075 and 0.050 and cluster the standard errors at the GP level. The bandwidths correspond to 17.3%, 13.6% and 9.9% of GPs, respectively. Also, since we test for heterogeneity in the gender quota effects across different sub-samples, manually chosen fixed bandwidths maintain consistency and facilitate comparison of estimates across specifications.

2.5.3 Results

We graphically show the first stage results in Figure 2.1. It shows that the likelihood of female reservation is zero on the left of the threshold and jumps to about 0.6 at the threshold. Appendix Table 2.10: Panel A reports the regression result for specification (2.5.2). It shows that we have a strong first stage across the three different bandwidths and the estimated discontinuity in the probability of female reservation is in the range of 0.53–0.59 across specifications. This lines up well with Figure 2.1.

We now discuss the second stage results using the household level data. Table 2.1 reports the average effect of female reservation on the probability that any household without a toilet received one in 2016–2017. The coefficients at the three bandwidths—0.1 (column (1)), 0.075 (column (2)) and 0.05 (column (3))—are 0.0152, 0.0206 and 0.0282 respectively, i.e., they are all positive. The effect sizes vary from 15% (in column (1)) to 28% (in column (3)) of the mean allocation probability, i.e., they are moderate in size. But all the coefficients are statistically insignificant due to large standard errors. The average treatment effect is, therefore, positive but noisy. In the following section we show that the average noisy result subsumes significant heterogeneity across GPs.

2.5.4 Heterogeneity in Gender Quota Effect

Estimation: The effect of gender quota could be heterogeneous across GPs either due to supply or demand related factors. To test for heterogeneity in treatment effect, we estimate specification (2.5.1) separately for GPs with Muslim and non-Muslim (or, Hindu) Sarpanches.

The gender gap in preferences is higher for Muslims than for Hindus as discussed in Section 2.3. Hence, the effect of gender quotas among Muslim Sarpanches on toilet provision could be higher than among Hindu Sarpanches due to differences in leader's preference (i.e., supply). However, Muslim Sarpanches are more likely to win in GPs with higher Muslim population shares. The average Muslim population share in the sample of GPs with Hindu Sarpanches is 10% while it is 49% in GPs with Muslim Sarpanches (see Appendix Table 2.8: Panel A). Appendix Figure 2.5 that plots the distribution of Muslim population shares for the two samples, also shows this clearly. If women from Muslim households express greater demand for toilets in the presence of female Sarpanches, then the effect of female reservation among Muslim Sarpanches could be higher due to changes in demand as well.

The sample of GPs with either Hindu or Muslim Sarpanches is obviously endogenous. However, within each sample, we can still estimate the causal effect of a female Sarpanch. This is because within each of the samples, it is still the case that female reservation status changes discontinuously at zero threshold value of the running variable. Appendix Figure 2.7 shows that this is indeed true. Additionally, Appendix Figure 2.6 shows the McCrary tests for the two samples separately. For both samples, we observe that the density of the running variable is continuous at the threshold, signifying non-manipulation in both the samples. Moreover, columns (2) and (3) in Appendix Table 2.9: Panel A report that the baseline characteristics of GPs do not show any jump at the threshold for the two samples. In Panel B, columns (2) and (3) show that toilet provision in previous years also does not change discontinuously at the threshold for either of the two samples. Therefore, we conclude that the same identification strategy is valid for each of the two samples. We hypothesize that if the gender gap in toilet preferences matters through either demand or supply, then we should expect:

Hypothesis 1. τ *(for Muslim Sarpanch)* $>$ τ *(for Hindu Sarpanch).*

Results: We have strong first stages for Hindu and Muslim Sarpanches separately, as depicted in the Appendix Figure 2.7 and estimated in Appendix Table 2.10 Panels B and C. The likelihood of female reservation is zero on the left of the threshold and jumps to approximately 0.6 and 0.5 at the threshold for Hindu and Muslim Sarpanches respectively. Table 2.2 columns (1)–(3) report the estimation results of specification (2.5.2) for the sample of GPs with Hindu Sarpanches while columns (4)–(6) report the results for the Muslim Sarpanches. As before, we show the results at the three bandwidths—0.1 (columns (1) and (4)), 0.075 (columns (2) and (5)) and 0.05 (columns (3) and (6)). The results show that female reservation within Hindu Sarpanches has no effect on toilet construction. The coefficients are small and statistically insignificant in all the three columns. On the other hand, among

Muslim Sarpanches, female reservation has a large and statistically significant positive effect on construction of toilets for all bandwidths. The estimated coefficients are 0.15, 0.21 and 0.26 for the three bandwidths respectively. The magnitudes are large considering the mean probability of toilet provision is between 0.13–0.14 in the estimating sample.[30] This finding clearly validates Hypothesis 1.[31]

Figure 2.2 shows the second stage results graphically. It confirms the positive effect of gender quota among Muslim Sarpanches while we do not observe any jump in the probability of allocation in GPs having Hindu Sarpanches. In an alternate specification we pool all GPs in one sample and run a difference-in-discontinuities specification (Grembi et al., 2016) to test whether the effect of female reservation is heterogeneous across GPs with Hindu and Muslim Sarpanches. Appendix Table 2.11 reports the result. Consistent with Table 2.2, we find that the effect is small and statistically insignificant for Hindu Sarpanches and is positive and statistically significant among Muslim Sarpanches.

2.6 Isolating Demand and Supply Mechanisms

2.6.1 Identification

In order to understand whether the heterogeneity in the gender quota effect is driven by demand or supply mechanism, we perform the following analysis. Let $T_g(L_g, D_g)$ be the expected proportion of eligible households that receive toilets from the Sarpanch in GP g. The allocation depends on two features of the GP—(i) preference of the elected leader, captured by his/her gender identity (L_g), and (ii) the aggregate expressed demand from the households (D_g) in the GP. Let $T_{h,g}$ be the indicator of a household $h \in \mathcal{H}_g$ receiving toilet in GP g. Then,

$$T_g(L_g, D_g) \;=\; \mathbb{E}\left[\frac{1}{H_g}\sum_{h\in\mathcal{H}_g} T_{h,g}(L_g, D_g)\right]$$

where $\mathcal{H}_g$ is the set of eligible households in g and H_g is the number of such households.[32] For reasons argued above, T_g can also depend on the religion of the Sarpanch. For simplicity,

[30] The results in Table 2.2 and all other results are robust to having a quadratic specification for the running variable on both sides of the threshold.

[31] The result could also have been consistent with the possibility of Muslim female Sarpanches exhibiting greater in-group favoritism (i.e., allocating more toilets to their own group than Hindus) than Hindu female Sarpanches. The empirical results in the next section, however, rule this out. We also rule out in Section 2.7 that the result is driven by higher gender gap in ability among Muslim Sarpanches (relative to Hindu ones).

[32] We allow the function $T_{h,g}(.,.)$ to be different across households within the same GP. This is because, $T_{h,g}$ can be a function of household specific characteristics, such as its religion, gender of the household head, wealth level etc. Hence two different households in the same GP can have different likelihoods of receiving toilets, even though they have the same values of L_g and D_g.

we suppress this information for the most part of this section. We implicitly assume that when we change the leader's gender, we keep his/her religion the same.

The source of complexity is that demand itself can change in response to leader's gender (and hence, preference). For instance, when the leader changes from male ($L_g = 0$) to female ($L_g = 1$) due to the quota, female voters can express greater demand for toilets, increasing D_g. Therefore, D_g is also a function of L_g and $D_g(1) > D_g(0)$.[33] Moreover, only $T_{h,g}$ and L_g are observable, while D_g is not. Using our identification strategy explained above, we can generate random variation in the leader's gender identity. We, therefore, can compute the average treatment effect (ATE) of gender quota, given by

$$\Delta T_g \;\equiv\; [T_g(1, D_g(1)) - T_g(0, D_g(0))]. \tag{2.6.1}$$

As is evident from the expression, the average effect is driven by changes in both supply and demand. We can, therefore, decompose the total effect into demand and supply in the following manner:

$$\begin{aligned} \Delta T_g &= \underbrace{[T_g(1, D_g(1)) - T_g(1, D_g(0))]}_{\textit{Demand}} + \underbrace{[T_g(1, D_g(0)) - T_g(0, D_g(0))]}_{\textit{Supply}} \\ &\equiv \Delta^d T_g \;+\; \Delta^s T_g \end{aligned} \tag{2.6.2}$$

where $\Delta^d T_g$ denotes the expected change in the allocation due to greater demand, conditional on the gender quota being in place. We refer to this as the demand effect. $\Delta^s T_g$ denotes the expected change in allocation induced by change in the leader's identity, conditional on the demand remaining what it would be under male leadership. We refer to this as the supply effect.

Demand: To identify whether demand is important we rely on the argument, provided in Section 2.3 above, that male and female headed households would register differential increase in demand in the presence of female leadership. Let's denote the expected proportion of eligible male and female headed households receiving toilets by T_g^m and T_g^f, respectively. Hence,

$$T_g^m(L_g, D_g) \;=\; \mathbb{E}\left[\frac{1}{H_g^m}\sum_{h\in\mathcal{H}_g^m} T_{h,g}(L_g, D_g)\right], \quad \text{and} \quad T_g^f(L_g, D_g) \;=\; \mathbb{E}\left[\frac{1}{H_g^f}\sum_{h\in\mathcal{H}_g^f} T_{h,g}(L_g, D_g)\right]$$

[33] Here, $D_g(1)$ represents the demand under a female leader and $D_g(0)$ represents demand under a male leader.

where $\mathcal{H}_g^m$ ($\mathcal{H}_g^f$) is the set of male (female) headed households and H_g^m (H_g^f) is the number of such households. Hence, we can write that

$$T_g \;=\; (1-\lambda_g)T_g^m + \lambda_g T_g^f \;=\; T_g^m + \lambda_g(T_g^f - T_g^m) \tag{2.6.3}$$

where λ_g is the proportion of female headed households in g. Following a similar definition as in equation 2.6.1 we can then compute:

$$\Delta T_g^m \;\equiv\; [T_g^m(1, D_g(1)) - T_g^m(0, D_g(0))], \quad \text{and} \quad \Delta T_g^f \;\equiv\; [T_g^f(1, D_g(1)) - T_g^f(0, D_g(0))].$$

where ΔT_g^m (ΔT_g^f) is the ATE among male (female) headed households. Similar to the decomposition carried in equation 2.6.2, we can write

$$\Delta T_g^m \;=\; \Delta^d T_g^m \;+\; \Delta^s T_g^m, \quad \text{and} \quad \Delta T_g^f \;=\; \Delta^d T_g^f \;+\; \Delta^s T_g^f$$

where $\Delta^d T_g^m$ ($\Delta^d T_g^f$) and $\Delta^s T_g^m$ ($\Delta^s T_g^f$) are the demand and supply side effects for male (female) headed households. If the supply side effect is completely determined by the leader's characteristics then it would be the same across both male and female headed households.[34] Therefore, $\Delta^s T_g^m = \Delta^s T_g^f$. Hence we get that,

$$\Delta T_g^f - \Delta T_g^m \;=\; \Delta^d T_g^f - \Delta^d T_g^m.$$

The L.H.S. can be estimated which, therefore, would allow us to estimate the difference in demand across female and male headed households. Using equation 2.6.3 we get:

$$\Delta^d T_g \;=\; \Delta^d T_g^m + \lambda_g(\Delta^d T_g^f - \Delta^d T_g^m) \;=\; \Delta^d T_g^m + \lambda_g(\Delta T_g^f - \Delta T_g^m). \tag{2.6.4}$$

Therefore, we partially identify $\Delta^d T_g$ by estimating $(\Delta T_g^f - \Delta T_g^m)$. Even though the proportion of female headed households in a GP is small, this identification is crucial to establish whether increased demand under female leadership can cause differential allocation.

Supply: To identify supply, we rely on the fact that the gender gap in preference is higher among Muslims than Hindus. Therefore, the supply effect of a Muslim female Sarpanch (relative to a Muslim male Sarpanch) would also be higher than that of a Hindu female Sarpanch (relative to a Hindu male Sarpanch). Analysis in Section 2.5.4 shows that the

[34] It is certainly possible for female leaders to be differentially focused on female headed households, leading to differential allocation to them driven by supply side considerations. We, however, rule out this possibility during our estimation.

overall effect of Hindu female Sarpanch is effectively zero, implying that her supply effect would also be negligible. Therefore, if supply is at all an important mechanism, we should expect it to be positive for Muslim female Sarpanches. Therefore, we can write

$$\Delta^s T_g^I \ > \ \Delta^s T_g^H$$

where $\Delta^s T_g^I$ ($\Delta^s T_g^H$) is the supply effect in the sample of GPs with Muslim (Hindu) Sarpanches. Let ΔT_g^I and ΔT_g^H denote the ATE of a Muslim and a Hindu female Sarpanch respectively. Then, we can write

$$\Delta T_g^H \ = \ \Delta^d T_g^H \ + \ \Delta^s T_g^H \quad \text{and} \quad \Delta T_g^I \ = \ \Delta^d T_g^I \ + \ \Delta^s T_g^I$$

where $\Delta^d T_g^H$ and $\Delta^d T_g^I$ are the demand effects of female reservation in the two kinds of GPs. Now, generally $\Delta^d T_g^H$ and $\Delta^d T_g^I$ would be different, as the GPs where Muslims become Sarpanches would be very different from those with Hindu Sarpanches. However, if the religion of Sarpanches is randomly assigned (say, in a sample of GPs with close elections between a Hindu and a Muslim), then the GPs with Hindu and Muslim Sarpanches would be the same on average, implying that they would have the same demand effect due to gender quota. This is because the two samples would have the same demographics including same population share of Muslims and same preference for a Muslim Sarpanch (as the vote shares for a Muslim candidate would be almost identical). However, the overall effect of female reservation could still be different due to difference in the supply effect. Therefore, under the assumption that the religion of Sarpanch is randomly assigned, we get that

$$\Delta^d T_g^H \ = \ \Delta^d T_g^I$$

Therefore,

$$\Delta T_g^I - \Delta T_g^H \ = \ \Delta^s T_g^I - \Delta^s T_g^H$$

Suppose ω is the share of GPs with a Muslim Sarpanch. Then,

$$\Delta^s T_g \ = \ \Delta^s T_g^H + \omega(\Delta^s T_g^I - \Delta^s T_g^H) \ = \ \Delta^s T_g^H + \omega(\Delta T_g^I - \Delta T_g^H) \tag{2.6.5}$$

Similar to the identification of demand, we partially identify supply effect $\Delta^s T_g$ by ($\Delta T_g^I - \Delta T_g^H$). Moreover, $\Delta^s T_g^H$ is zero for reasons argued above. Hence, in practice, we are able to identify the full supply effect. Substituting equations (2.6.4) and (2.6.5) in equation (2.6.2) we get,

$$\Delta T_g \;=\; \overbrace{\Delta^d T_g^m + \underbrace{\lambda_g(\Delta T_g^f - \Delta T_g^m)}_{\text{Identifiable demand}}}^{\text{Demand effect}} + \overbrace{\underbrace{\Delta^s T_g^H}_{=0} + \underbrace{\omega(\Delta T_g^I - \Delta T_g^H)}_{\text{Identifiable supply}}}^{\text{Supply effect}} \tag{2.6.6}$$

The equation above clearly spells out our identification strategies for demand and supply. We estimate heterogeneous treatment effect across female and male headed households to identify demand and across GPs with (randomly assigned) Muslim and Hindu Sarpanches to identify supply.[35]

2.6.2 Demand Estimation Strategy and Results

Estimation: We use the *difference-in-discontinuities* approach proposed by Grembi et al. (2016) to estimate $(\Delta T_g^f - \Delta T_g^m)$. We do it separately for the samples of GPs with Hindu and Muslim Sarpanches and within each case, for Hindu and Muslim households. We do this to guard against an alternate interpretation that our results are driven by favorable concern or focus that female Sarpanches may have for female headed households. We explain this point in greater detail after describing our empirical strategy.

Let $R \in \{H, I\}$ denote the religion of the Sarpanch and $r \in \{h, i\}$ denote the religion of any household, with H and h denoting Hindu and I and i, Muslim (Islam). Let $S_{R,r}$ denote the set of households belonging to religion r in GPs where the religion of the Sarpanch is R. For each combination (R, r), we then estimate following specification:

$$\begin{aligned} Y_{h,g} &= \alpha_0 + \tau^{R,r} Q_g + \alpha_1 X_g + \alpha_2 X_g * A_g \\ &+ F_{h,g} * [\theta_0 + \rho^{R,r} Q_g + \theta_1 X_g + \theta_2 X_g * A_g] + u_g, \quad \text{for } h \in S_{R,r} \end{aligned} \tag{2.6.7}$$

where $F_{h,g}$ is a dummy variable that indicates whether a household is headed by a woman or not. All parameters in equation (2.6.7) vary by (R, r). However, for notational simplicity we give the superscript (R, r) only to the main coefficients of interest ρ and τ. As before, Q_g is instrumented with A_g in the following first stage equations, which are again estimated for each combination (R, r):

$$Q_g = \beta_0 + \gamma A_g + \beta_1 X_g + \beta_2 X_g * A_g + F_{h,g} * [\delta_0 + \lambda D_g + \delta_1 X_g + \delta_2 X_g * A_g] + \epsilon_g \tag{2.6.8}$$

[35] We can not identify demand by testing for heterogeneous treatment effects across Muslim and Hindu households because Muslim Sarpanches allocating more toilets to Muslim households could be due to own-group favoritism as well. Additionally, Hindus are the majority group on average even in GPs with Muslim Sarpanches. Therefore, Muslim Sarpanches may have electoral incentives to allocate toilets to Hindus, even if they do not demand it as much as Muslims. This may lead to underestimation of the demand effect. Female headed households constitute a small fraction of the GP population, and therefore, electoral concerns are absent in their case. We rule out the favoritism channel in our empirical results.

$$Q_g * F_{h,g} = \beta_0^{'} + \gamma^{'} A_g + \beta_1^{'} X_g + \beta_2^{'} X_g * A_g + F_{h,g} * [\delta_0^{'} + \lambda^{'} A_g + \delta_1^{'} X_g + \delta_2^{'} X_g * A_g] + \epsilon_g^{'} \quad (2.6.9)$$

Here, $\rho^{R,r}$ estimates the differential allocation to female headed households relative to male headed ones by female Sarpnches. It therefore estimates $(\Delta T_g^f - \Delta T_g^m)$ for each combination of (R, r). Since female headed households are more likely to register their demand under a female leader than male headed households of the same religion, we hypothesize that if demand is important for provision of toilets then:

Hypothesis 2. $\rho^{R,r} > 0$ *for all combinations of* (R, r).

Moreover, since the gender gap in preference is higher for Muslims, we hypothesize that:

Hypothesis 3. *(i)* $\rho^{H,i} > \rho^{H,h}$ *and (ii)* $\rho^{I,i} > \rho^{I,h}$

On the other hand, if differential allocation to female headed households is completely driven by supply side considerations, then we should expect $\rho^{H,i} = \rho^{H,h}$ and $\rho^{I,i} = \rho^{I,h}$. Additionally, if female Sarpanches are more focused on female headed households from their own religious group, then we should expect $\rho^{H,i} < \rho^{H,h}$. Therefore, validation of Hypothesis 3 would allow us to establish that our estimates indeed capture demand. From equation (2.6.6), we get that the estimate for τ captures $\Delta^d T_g^m + \Delta^s T_g$. Therefore, we need a separate estimate for the supply mechanism $\Delta^s T_g$ to infer what part of τ is supply.

Results: To test Hypotheses 2 and 3 we need to know the gender identity of the household head. This information is available only for the Above Poverty Line (APL) households. This is because all the Below Poverty Line (BPL) households are eligible for the scheme while a subset of APL households are eligible. One eligibility criteria is whether an APL household is headed by a woman.[36] Hence, we restrict our attention to eligible APL households for this exercise.[37] We estimate specification (2.6.7) in four sub-samples of eligible APL households: $S_{H,h}$, $S_{H,i}$, $S_{I,h}$ and $S_{I,i}$, i.e., Hindu (h) and Muslim (i) households in GPs with Hindu (H) and Muslim (I) Sarpanches. Table 2.3 columns (1)–(4) report the results for the four sub-samples respectively. We present the results at 0.1 bandwidth, but the results are similar for smaller bandwidths as well (reported in the Appendix Table 2.12).

The result in column (1) shows that female reservation within Hindu Sarpanches increases the probability of toilet provision to Hindu female headed households by 0.05. The coefficient is noisily estimated, even though the effect size is 45% of the mean. On the other hand,

[36] We discuss the eligibility criteria for the SBM scheme among APL households in Section 2.2.1.

[37] The result in Table 2.2 remains the same if we restrict the sample to eligible APL households instead of all eligible households.

the probability increases by 0.35 (column (2)) for *Muslim* female headed households. The estimate is statistically significant at 5% level of significance and is a considerably large effect given the mean of 0.12. Moreover, the column (2) coefficient is larger than the column (1) coefficient (p-value = 0.02), which validates part (i) of Hypothesis 3.

Within GPs with Muslim Sarpanches, female reservation leads to a jump in the probability of provision by 0.36 among Hindu female headed households (column (3), row (2)). The coefficient, however, has a high standard error and is not statistically significant. The magnitude of the effect, nonetheless, is large—suggesting greater allocation to Hindu female headed households. The corresponding coefficient in column (4) is 0.49 and is statistically significant at 1% level. Therefore, we find that Muslim female headed households experience a meaningful and statistically significant increase in the likelihood of toilet provision due to female reservation among Muslim Sarpanch GPs. The difference between the coefficients in columns (4) and (3), is also positive (0.13=0.49-0.36), though not statistically significant. However, the positive difference is consistent with the part (ii) of the Hypothesis 3.[38] Moreover, all the four coefficients for the heterogeneous effect on female headed households (row 2), across columns (1)–(4) are positive and economically large, though only two are statistically significant, implying that Hypothesis 2 is validated as well. An alternate explanation for the result could be that female headed households are poorer and hence may need the toilet subsidy more, resulting in greater allocation by the Sarpanch. We, however, find that the coefficients for female headed households in row (3) across all the columns are negative and large in magnitude, implying that those households are *less likely* to receive toilets relative to male headed households when the Sarpanch is a man. This is not consistent with the above argument.[39] Moreover, the result is in line with our demand mechanism: under a male Sarpanch, these households are less likely to register their demand for toilets compared to male headed households—leading to lower allocation. Finally, the coefficient for Female reservation (τ) in row 1 is small in columns (1) and (2), and large but noisy columns (3) and (4). This is also consistent with demand: Muslim Sarpanch GPs have significantly higher share of Muslim households and hence higher level of collective demand for toilets, which can result in higher allocation to male headed households among both Hindus and Muslims.

[38]Our results are not driven by initial lower toilet provision to either Muslim households or households in GPs with Muslim Sarpanches. In fact, we find that 44% of Muslim households had access to toilets at the end of 2015-16 while it was 36% for Hindu households. Consistent with this, the average ownership of toilets in GPs with Muslim Sarpanches was also higher than those with Hindu GPs in the pre-treatment period (see Appendix Table 2.8: Panel B).

[39]Appendix Table 2.7: Panel A reports the differences in characteristics between the two types of households using NFHS 2015–16 data. We find that FHHs are indeed poorer on average. However, Panel B reports that Muslim female headed households are less likely to be poor than their Hindu counterparts. Therefore, the heterogeneity in allocation across religious groups towards female headed households can't be attributed to Muslim female headed households being poorer.

[40] We, therefore, conclude that household demand is vitally important in understanding the effect of gender quota in elections on provision of goods by the leader.

2.6.3 Supply Estimation Strategy and Results

Estimation: We now estimate $(\Delta T_g^I - \Delta T_g^H)$ in equation (2.6.6). We first select the sample of GPs where the top two candidates were a Hindu and a Muslim and the election was close. Within this sample, the religion of the Sarpanch would effectively be randomly assigned (see, for example, Meyersson, 2014). The close election regression discontinuity design is a popular method to generate random variation in the identity of elected leader (Eggers et al., 2015). We, therefore, use the differences-in-discontinuities strategy to estimate $(\Delta T_g^I - \Delta T_g^H)$—where the (sharp) discontinuity around close elections generates random variation in the religion of Sarpanch and the quota for women is the difference variable. We prefer this strategy over an alternate one, where we use regression discontinuity in the running variable associated with female reservation (i.e., X_g in equation (2.6.7)) to generate exogenous variation in female reservation status. This method would require us to use RDD across both religion and female reservation in the same specification. While it would generate exogenous variation in both variables, it would force us to focus on a sample of GPs that are just around the threshold value of X_g as well as experienced a close election between a Hindu and a Muslim candidate. The sample of such GPs is small and very special. Moreover, the practice of using RDD strategies on both differencing variables is not common. Consequently, we believe our strategy is easier to interpret and gives a more reliable estimate due to the larger sample size.

We run the following specification on the sample of close election GPs to estimate the differential effect of having a Muslim Sarpanch (relative to a Hindu Sarpanch) in female quota GPs vis-a-vis non-female quota GPs:

$$\begin{aligned} Y_{h,g} &= \alpha_0 + \beta I_g + \alpha_1 V_g + \alpha_2 I_g * V_g \\ &+ Q_g * [\theta_0 + \phi I_g + \theta_1 V_g + \theta_2 I_g * V_g] + u_g \end{aligned} \tag{2.6.10}$$

where I_g is a dummy that takes value one if the Sarpanch in GP g is Muslim and zero otherwise. V_g is the margin of victory for a Muslim candidate, i.e., it is defined as (vote share of Muslim - vote share of Hindu). Q_g, as before, is an indicator of female quota in g. Our coefficient of interest is ϕ, which estimates $(\Delta T_g^I - \Delta T_g^H)$. Of course, the GPs with and without female reservation are not the same. Female reservation status is, however, almost

[40] The coefficient estimates in row (1), columns (3) and (4) are similar in magnitude, implying that Muslim female Sarpanches do not exhibit significant own-group favoritism in allocation of toilets.

completely determined by X_g. Hence, we show robustness of our result using an alternate specification where we control for X_g in the following way:

$$\begin{aligned} Y_{h,g} &= \alpha_0 + \beta I_g + \alpha_1 V_g + \alpha_2 I_g * V_g + \beta_1 X_g + \beta_2 Q_g * X_g \\ &+ Q_g * [\theta_0 + \phi I_g + \theta_1 V_g + \theta_2 I_g * V_g] + u_g \end{aligned} \tag{2.6.11}$$

Results: Table 2.4 reports the results from estimation of specification 2.6.10 in the sample of GPs with close election between a Hindu and a Muslim candidate.[41] The three columns refer to three definitions of close election. Column (1) reports the result when the absolute value of margin of victory is at most 0.1, while columns (2) and (3) report it for bandwidths 0.075 and 0.05 respectively.[42] We find that the female reservation dummy has a small and statistically insignificant coefficient across all the columns, implying null effect of the female quota among Hindu Sarpanches. This is consistent with the result in Table 2.2 columns (1)–(3). For all the bandwidth specifications, we find that the difference-in-discontinuity estimate is also very small and statistically insignificant. Therefore, election of a female Muslim Sarpanch vis-a-vis a female Hindu Sarpanch does not change the probability of toilet construction. This may seem surprising given the result in Table 2.2 that showed that the effect of female reservation is large and positive in GPs with Muslim Sarpanches (columns (4)–(6)). However, as we argued in Section 2.5.4, the results in Table 2.2 can be due to either demand or supply, as the GPs with Muslim Sarpanches have high population share of Muslims relative to GPs with Hindu Sarpanches. We remove these differences across the two samples in Table 2.4 by conditioning on there being a close election between candidates belonging to different religions. The consequent null effect implies that supply mechanism can not explain the differential effects across the two samples. For robustness, we estimate equation (2.6.11). The results, reported in Appendix Table 2.15, remain the same.

2.7 Robustness

2.7.1 Demand Estimation in High Muslim Share GPs

The demand estimation results show that female headed households receive on average more toilets under a female Sarpanch than a male Sarpanch. However, the marginal effect of a female headed household demanding toilet may be greater when a greater share of the GP population is also demanding it, due to complementarities associated with collective

[41] Appendix Figure 2.8 provides a McCrary Test for manipulation of the running variable in this sample of GPs. We find that there is no discontinuity in the density of running variable at the threshold value of zero.
[42] The CCT bandwidth in this case is 0.082.

action. Then the effect is likely to be even greater in GPs with high population share of Muslims—who have a greater preference for toilets. To test this, we estimate specification 2.6.7 on the sample of eligible APL households belonging to GPs with high Muslim population shares. We define a GP to have a high Muslim share if its share is greater than the 85^{th} percentile of the Muslim share distribution in the full sample (= 27% of Muslim share). Table 2.16 reports the results. We find that the estimates of ρ are significantly larger across all the four columns in Table 2.16 relative to Table 2.3. These results confirm that the demand effect is stronger in GPs where Muslim population share is higher. We also observe that the main effect of female reservation in Muslim Sarpanch GPs is similar across Hindu and Muslim households. The coefficient estimate is 0.213 for Hindus (columns (3)) and 0.189 for Muslims (column (4)) in Table 2.16. This further rules out the possibility that the heterogeneous effect of gender quota in Muslim Sarpanch GPs is driven by greater own-group favoritism exercised by Muslim female Sarpanches.

2.7.2 Female Sarpanches in Open Elections

In our main analysis we have examined gender quota and the mechanisms behind its effect. In some of the GPs that are not reserved for women, i.e., where the Sarpanch election is open to both genders, female candidates also win and become Sarpanches. In this section, we test whether electing a female Sarpanch in an open election results in effects that are consistent with our main findings. While female Sarpanches that win open elections can be very different from those than come through quotas, the demand mechanism may still be at work in this case. We use regression discontinuity design method in the sample of open election GPs that had a close election between a man and a woman to generate exogenous variation in the gender of Sarpanch. To estimate the demand mechanism, we then run the following specification:

$$\begin{aligned} Y_{h,g} &= \beta_0 + \pi\mathbb{I}[MoV_g > 0] + \beta_1 MoV_g + \beta_2 MoV_g * \mathbb{I}[MoV_g > 0] \\ &+ F_{h,g} * [\theta_0 + \delta\mathbb{I}[MoV_g > 0] + \theta_1 MoV_g + \theta_2 MoV_g * \mathbb{I}[MoV_g > 0]] + \epsilon_g \end{aligned} \tag{2.7.1}$$

where MoV_g is the margin of victory for a woman candidate in GP g, defined as the difference between the vote shares of the female and male candidates in GPs where the top two candidates are a man and a woman. The dummy $\mathbb{I}[MoV_g > 0]$ therefore is an indicator of female Sarpanch. $F_{h,g}$ is an indicator of female headed household, as before. Our coefficient of interest is δ that captures the demand mechanism. We separately estimate δ for Hindu and Muslim female sarpanches against any male Sarpanch. We use any male Sarpanch in the control group as opposed to Hindu and Muslim male Sarpanch for the two samples separately because of sample size and power considerations. As before, within each case, we run separate regressions for Hindu and Muslim households to verify whether the heterogeneous treatment effect is indeed demand.

Appendix Table 2.17 reports the results. We maintain the same bandwidth choice of 0.1 in this case as well.[43] We find that even in open elections, Hindu and Muslim female Sarpanches allocate additional toilets to female headed households, except in column (1), where the coefficient is small and negative. The result, therefore, mostly validates Hypothesis 3. Additionally, estimate of δ in row (2) of Table 2.17 is larger in magnitude for Muslim households (columns (2) and (4)) compared to Hindu households (columns (1) and (3)), which is consistent with Hypothesis 3. Therefore, we find that the importance of demand mechanism driving the overall effect of female Sarpanches broadly generalizes to open elections as well.

2.7.3 Placebo: Effect Across SC/ST households

We do not find any evidence of differential effect of female reservation on SC/STs, among APL households (Appendix Table 2.18).

2.7.4 Alternate Mechanism : Ability of Female Sarpanches

Here we consider an alternate hypothesis that the differential female reservation effect across GPs with Hindu and Muslim Sarpanches could be driven by differential ability of Hindu and Muslim women who come to the leadership position. To examine this, we look at employment provision under National Rural Employment Guarantee Scheme (NREGS), which is the largest expenditure head in a GP budget and constitutes an overwhelming majority of a GP's annual expenses.[44] Greater provision under this scheme could potentially signal higher ability of Sarpanches to implement public projects. We estimate the effect of female reservation on two outcomes in Appendix Table 2.19—expenditure on NREGS scheme per capita in a GP (columns (1) and (3)) and person days of employment generated under the scheme per capita (columns (2) and (4)). We estimate it separately for GPs with Hindu and Muslim Sarpanches. We find that female reservation does not lead to greater provision under the scheme among either Hindu or Muslim Sarpanches. Thus, our results for toilet provision are unlikely to be driven by differential gender gap in the ability of female Sarpanches across religions.

[43] The CCT bandwidths are 0.103 and 0.111 for Hindu and Muslim female leader GPs, respectively.

[44] This scheme was launched in 2006 and aims to provide 100 days of employment per year to every adult in rural India.

2.8 Conclusion

This chapter uses a novel identification strategy and data to infer whether the effect of gender quotas in elections on public goods provision can be driven by greater demand expressed by female voters in the presence of female leadership. We identify both demand and supply side channels by using provision of a good that is targeted at the household level and for which women exhibit a greater preference, namely household toilets. We document stark differences in gender gap in preference for toilets across two well-identified and salient population groups—Hindus and Muslims. We also show that conditional on the religion of the household, female headed households are likely to express greater demand for toilets due to greater autonomy in decision making and higher political participation of women in these households. Consistent with the demand mechanism, we find that female reservation leads to a significantly higher allocation for female headed households, and the result is even stronger among Muslims households. The religion of the female Sarpanch, on the other hand, has no effect on toilet provision, implying absence of the supply mechanism. Our results establish that there is a large heterogeneity in the effect of female quota on toilet provision across GPs within a state and the demand side factors drive all of the heterogeneity. The result, therefore, highlights the importance of variation in the preference for public goods across regions in potentially explaining the mixed evidence found in the literature on the effect of gender quotas in elections. More importantly, it suggests that policies that empower women voters and encourage them to participate in political processes can make gender quotas more effective and can significantly improve the substantive representation of women.

Figures and Tables

Figure 2.1: RD plots: McCrary Plot and First Stage

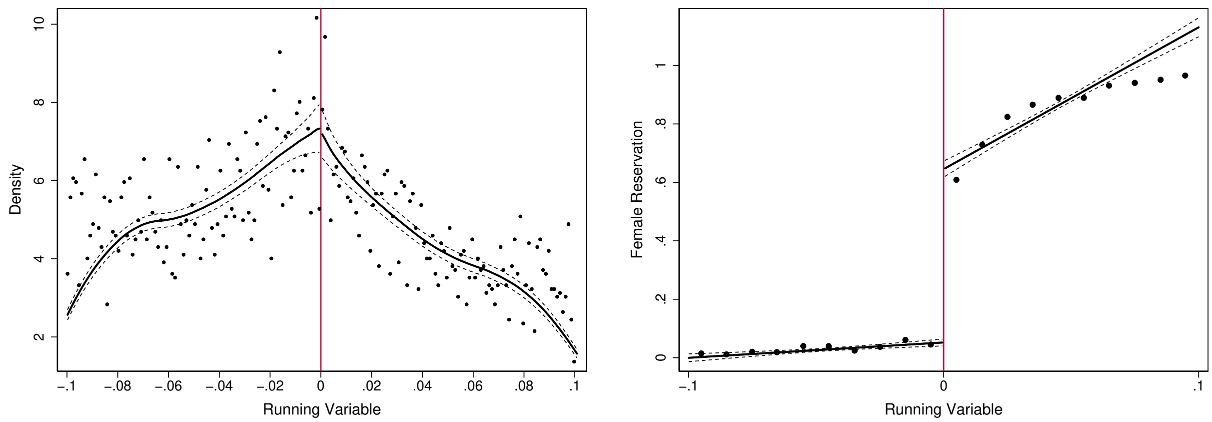

Overall McCrary Test

Overall First Stage

Figure 2.2: RD plots: Second Stage for Female Reservation on Toilet Construction

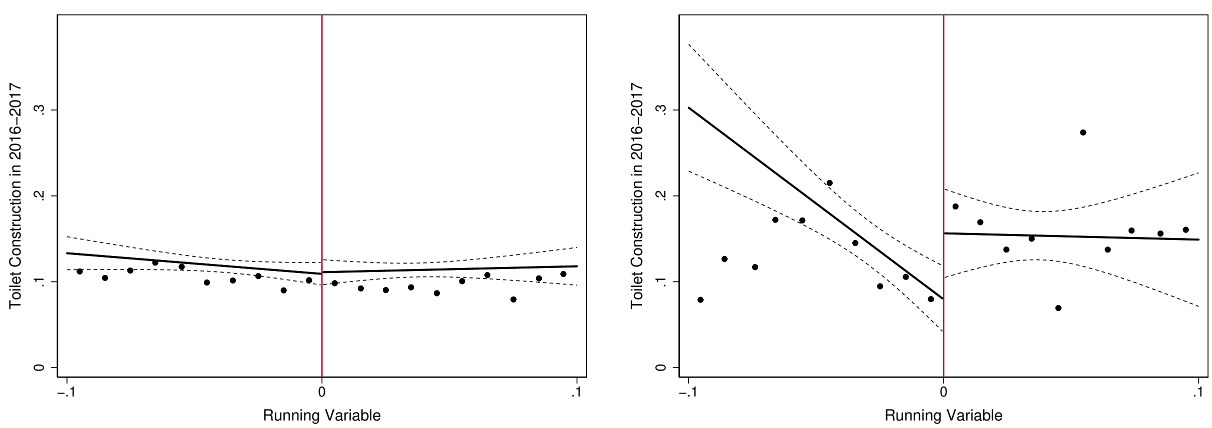

Hindu Sarpanch

Muslim Sarpanch

Table 2.1: Gender Quota Effect

	HH Received Toilet in 2016–17		
	(1)	(2)	(3)
Female reservation	0.0152	0.0206	0.0282
	(0.0162)	(0.0192)	(0.0240)
Mean dep. var.	0.104	0.104	0.099
Observations	2,470,191	1,962,725	1,457,226
Number of GPs	9,179	7,234	5,277
Polynomial order	1	1	1
Bandwidth	0.100	.075	0.050

Notes: The dependent variable is a dummy that takes value one if the household received a toilet in 2016-17, and zero otherwise. The sample includes eligible households,i.e., those that did not have toilet at the end of 2015-16 and were eligible to receive toilet under the SBM scheme. The polynomial order is 1. The bandwidths are manually chosen. Standard errors clustered at gram panchayat level and reported in parentheses.

Table 2.2: Gender Quota Effect and Its Heterogeneity

	Dep. Var.: HH Received Toilet in 2016–17					
	Hindu Sarpanch GPs			**Muslim Sarpanch GPs**		
	(1)	(2)	(3)	(4)	(5)	(6)
Female reservation	0.00325	0.00383	0.00598	0.153**	0.212***	0.263***
	(0.0166)	(0.0197)	(0.0246)	(0.0658)	(0.0792)	(0.0968)
Mean dep. var.	0.101	0.100	0.096	0.138	0.140	0.129
Observations	2,256,016	1,796,577	1,336,812	214,175	166,148	120,414
Number of GPs	8,278	6,541	4,774	901	693	503
Polynomial order	1	1	1	1	1	1
Bandwidth	0.100	.075	0.050	0.100	0.075	0.050

Notes: The dependent variable is a dummy that takes value one if the household received a toilet in 2016-17, and zero otherwise. The sample includes eligible households,i.e., those that did not have toilet at the end of 2015-16 and were eligible to receive toilet under the SBM scheme. The first and last three columns are households in GPs with Hindu and Muslim Sarpanches, respectively. The polynomial order is 1. The bandwidths are manually chosen. Standard errors clustered at gram panchayat level and reported in parentheses.

Table 2.3: Gender Quota Effect: Identifying Household Demand

	HH Received Toilet in 2016–2017			
	Hindu Sarpanch		Muslim Sarpanch	
	Hindu HH (1)	Muslim HH (2)	Hindu HH (3)	Muslim HH (4)
Female reservation	0.0173	-0.0334	0.121	0.163*
	(0.0209)	(0.0475)	(0.0909)	(0.0861)
Female reservation * Female Headed HH	0.0500	0.346**	0.360	0.489***
	(0.0487)	(0.146)	(0.262)	(0.177)
Female Headed HH	-0.0318	-0.105***	-0.132*	-0.107*
	(0.0202)	(0.0301)	(0.0706)	(0.0612)
Mean dep. var.	0.111	0.120	0.135	0.202
Observations	1,330,303	131,334	83,770	57,155
Number of GPs	8010	7041	848	837
Bandwidth	0.1	0.1	0.1	0.1
Polynomial order	1	1	1	1

Notes: The sample is Above Poverty Line (APL) households who are eligible for the SBM program and did not have toilets at the end of 2015–2016. The sample for column (1) is Hindu households under Hindu leaders while that for column (2) is Muslim households under Hindu leaders. The samples for columns (3) and (4) are defined similarly under Muslim leaders. Female Headed HH is a dummy that takes value one if the household head is a woman and zero otherwise. The polynomial order is 1. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.4: Muslim vs Hindu Sarpanches in Close Election GPs

	HH Received Toilet in 2016–2017		
	(1)	(2)	(3)
Female Reservation	-0.0106	-0.0154	-0.0120
	(0.0258)	(0.0293)	(0.0358)
Muslim Sarpanch	-0.0145	-0.0181	-0.00571
	(0.0228)	(0.0257)	(0.0309)
Female Reservation*Muslim Sarpanch	-0.00429	-0.00136	-0.00235
	(0.0368)	(0.0417)	(0.0505)
Mean dep. var.	0.099	0.099	0.104
Observations	943,640	777,713	569,103
Number of GPs	3,263	2,666	1,941
Bandwidth	0.100	0.075	0.050
Polynomial order	1	1	1

Notes: The sample is restricted to households which did not have toilets at the end of 2015–2016. The polynomial order is 1. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Appendix

2.A Preferences for Toilets: Evidence from Survey Data

2.A.1 Gender Gap in Preference for Toilets

We use the SQUAT survey to analyze preference for toilets. First, we examine the revealed preference for toilets by analyzing how likely a person is to use a toilet conditional on the household owning one. In Table 2.5, columns (1) and (2), we look at the probability of using a toilet for defecation, by gender, conditional on household owning a toilet using village and household fixed effects respectively. Results in Panel A, columns (1) and (2), clearly show that women are more likely to use in-home toilets. For the set of households that do not own a toilet, the survey asks a randomly chosen member of the household about his or her top three priorities for the household from a list of assets that the household does not possess, if money is not a constraint. We create three indicator variables that capture whether the respondent states toilet ownership to be among the top most, one of top two or top three priorities. We test whether respondent gender matters for the responses. The results are reported in columns (3), (4), (5) of the same table. We find that women are more 5% points more likely to state that their most top-most priority is a toilet for the house (column (3)). The estimates for the other two specifications are also positive, but are statistically insignificant. Overall our analysis demonstrates that, consistent with existing evidence, women have a stronger preference for a household toilet than men.[45]

2.A.2 Gender Gap in Preference between Hindus and Muslims

In panel B of Table 2.5, we examine the heterogeneity in the gender gap in preference for toilets across Hindus and Muslims, using the SQUAT survey. Columns (1) and (2) show that conditional on having a toilet, Muslim households are more likely to use a toilet than Hindu households. Moreover, the usage gap between Muslim women and men is almost double of

[45]Similar gender gap in toilet preference has been documented in other countries as well (Jenkins and Curtis (2005), Santos et al. (2011)).

the gap between Hindu women and men. Focusing on the households that do not yet own a toilet, columns (3), (4) and (5) show that relative to Muslim men, Muslim women are 22–24% more likely to report that it is among the top two or three priorities for them. The gender gap in reported preference among Hindus is positive but statistically insignificant. These results show that on the whole, the gender gap in preference for toilets is much higher for Muslims relative to Hindus.

2.A.3 Demand across Male and Female Headed Households

To examine female autonomy across male and female headed households, we create two indicators using the NFHS-4 survey. The first indicator uses the set of questions on whether the woman can make any of the following decisions on her own –accessing health care, major household purchases and visits to family. The second indicator captures whether the woman is allowed alone to any of these places –market, health facility and outside village. We regress these indicators on whether a woman belongs to a household headed by a woman and report the results in Appendix Table 2.6 Panel A, columns (1) and (2) respectively. We control for various individual and household characteristics and village fixed effects. The results show that women residing in households headed by women are more likely to take decision and go alone to places, indicating greater decision making power and autonomy for women in these households. We also find that female memebers who are not the head of the family themselves also enjoy this greater agency in female headed households (table not reported).

Additionally, we also test if women are more politically active if they live in a woman headed household by using the REDS data. Appendix Table 2.6 Panel A, columns (3) and (4) show the results for whether a woman has attended any of the last four village meetings and whether a woman is actively involved with any political party, respectively. We find a positive association between female headship of the household and women's active role in community and political engagement, controlling for religion and other household and individual level characteristics and village fixed effects. These results show that, conditional on women having greater demand for toilets, households headed by women are more likely to express the demand for it.

In Panel B of the same table we test if the positive relationship we find is driven by Hindus. We interact the indicator of female headed households with the indicator of the household being Muslim. We find that the relationship (both in the NHFS-4 and REDS) is similar across Hindus and Muslims. If anything, women in female headed Muslim households enjoy greater bargaining power within the household (column (1)) and attend village meetings more (column (3)) compared to their counterparts in Hindu female headed households.

2.B Additional Figures and Tables

Figure 2.3: Comparing Toilet Coverage: Administrative Data versus Survey Data (2015–16)

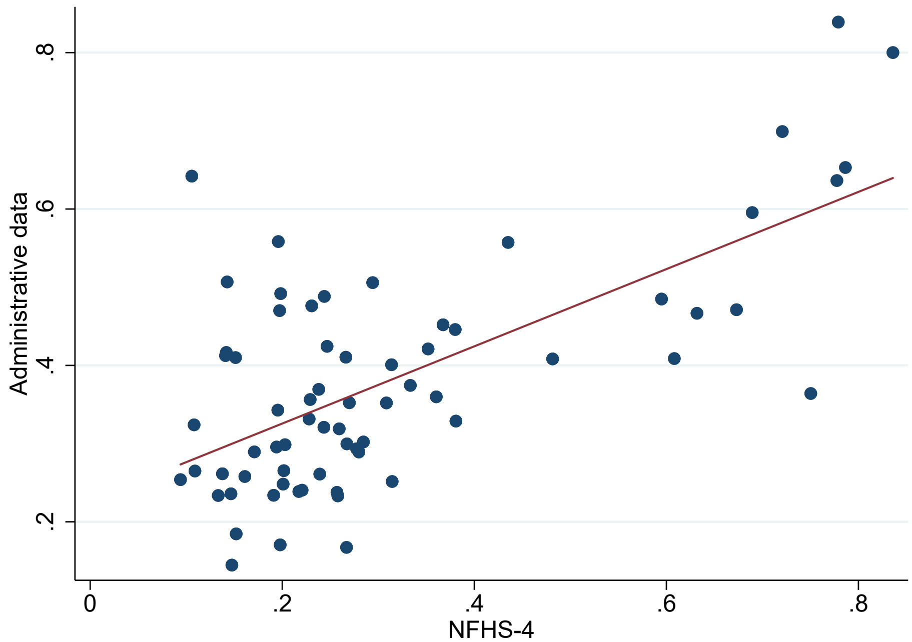

Source: Ministry of Drinking Water and Sanitation (MDWS), India for administrative toilet data and National Family Health Survey (NFHS)-4 (2015–16) for sanitation survey data.
Notes: The figure plots the proportion of households having a toilet in the administrative data and the proportion of households having a toilet in the NFHS data for the same district. A linear relationship between the two is fitted. Correlation = 0.7.

Figure 2.4: Comparing true and predicted Muslim population share

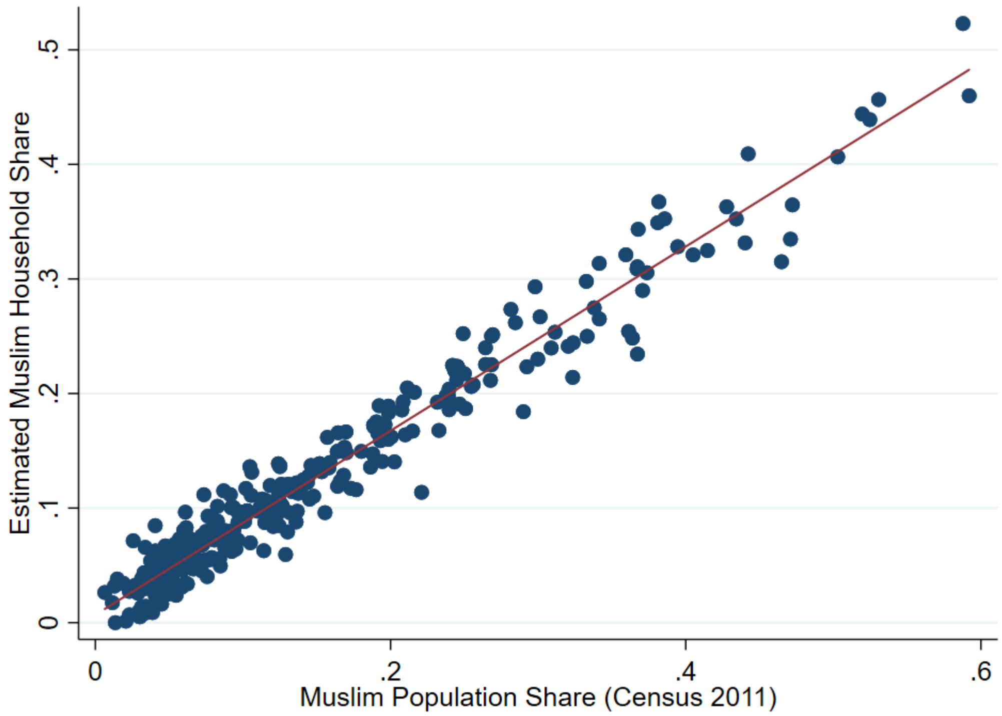

Note: The true Muslim population share based on 2011 census is on the x-axis. Population share estimated by the algorithm for 312 tehsils in U.P. is on y-axis. Correlation = 0.9776.

Figure 2.5: Density of Muslim Population Share

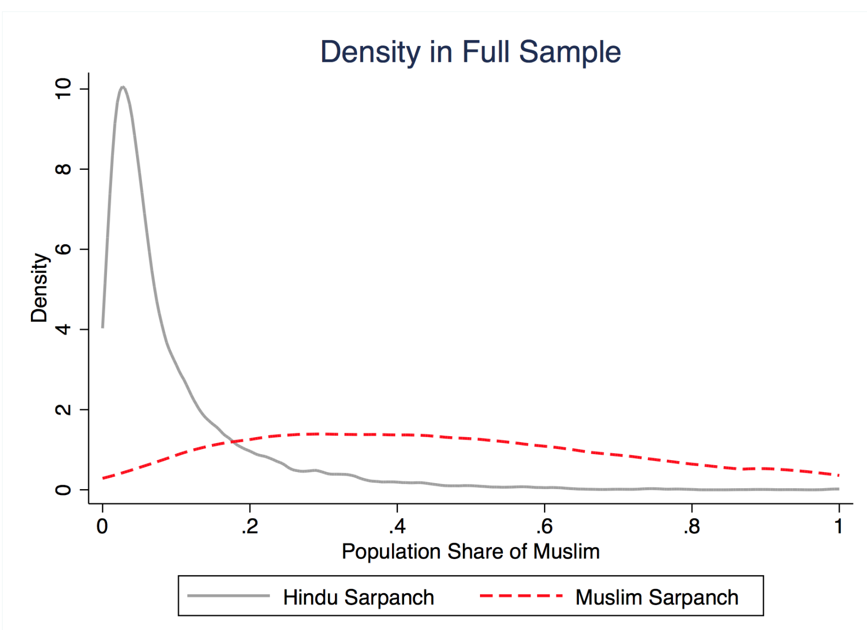

Figure 2.6: McCrary Plots for Discontinuity at the Cut-off

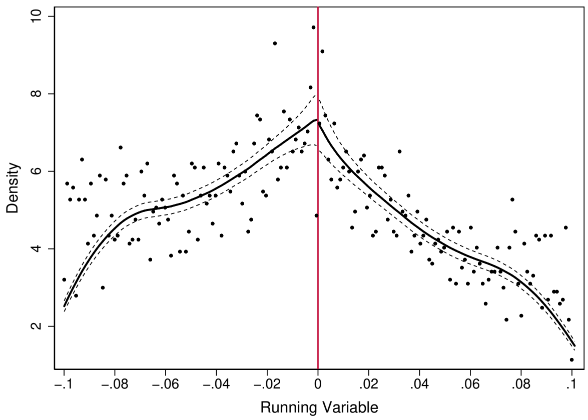

Hindu Sarpanch

Muslim Sarpanch

Figure 2.7: RD plots: First Stage for Gender Reservation

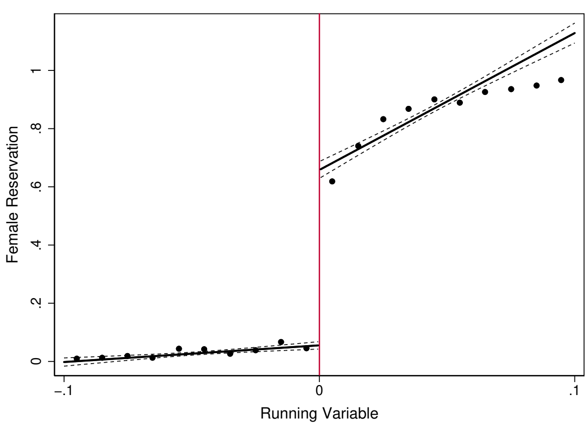

Hindu Sarpanch

Muslim Sarpanch

Figure 2.8: McCrary Plots for Discontinuity at the Cut-off in Close Elections

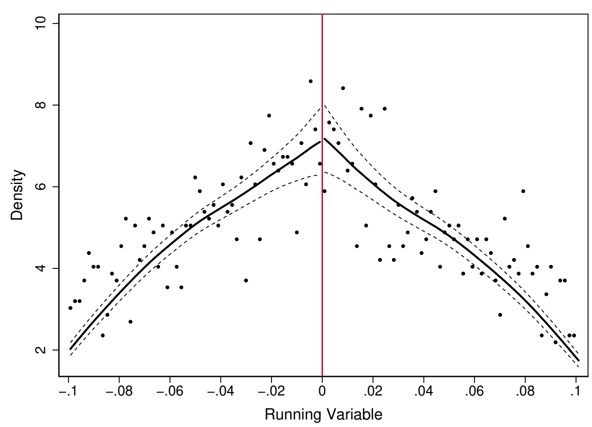

Hindu vs Muslim Close Election

Table 2.5: SQUAT Survey: Gender Gap in Latrine Preference

	Latrine Usage		Latrine Preference		
			Top	Top 2	Top 3
	(1)	(2)	(3)	(4)	(5)
			Panel A: Overall Gender Gap		
Female	0.0933***	0.0940***	0.0507*	0.0114	0.0215
	(0.00792)	(0.00620)	(0.0300)	(0.0282)	(0.0252)
			Panel B: Heterogeneity in Gender Gap		
Female	0.0889***	0.0903***	0.0456	0.00322	0.0150
	(0.00841)	(0.00651)	(0.0309)	(0.0290)	(0.0257)
Muslim * Female	0.0549**	0.0475**	0.0811	0.238*	0.221*
	(0.0234)	(0.0209)	(0.141)	(0.124)	(0.114)
Muslim	0.0903***		0.0328	0.0128	0.00193
	(0.0252)		(0.119)	(0.107)	(0.106)
Mean Dep. Var.	0.80	0.80	0.46	0.64	0.78
Observations	7,731	7,717	1,472	1,472	1,472
Fixed Effect	Village	HH	Village	Village	Village

Notes: The dependent variable for columns (1) and (2) is a dummy that takes value one if the individual uses the latrines for defecation. The samples in the first two columns only include households which have latrines. The dependent variable in each of the columns (3), (4), and (5) is a dummy that takes value one if the respondent reports a toilet as being the topmost, top two or top three priorities, respectively, if money was not a constraint, from a list of assets that the household does not have. The samples in the last three columns only include households which do not have latrines. The SQUAT states are Haryana, Uttar Pradesh, Madhya Pradesh, Bihar, and Rajasthan. Columns (1) and (3)-(5) have village fixed effects and control for household level assets and household's main source of income. Columns (2) and (4) have household level fixed effects. The number of observations in columns (1) and (2) differ with village and household fixed effects as single member households are dropped from the analysis. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.6: Female Autonomy in Female Headed Households

	NFHS Dataset		REDS Dataset	
	Decision	Allowed alone	Village Meeting Attendance	Political Party Activity
	(1)	(2)	(3)	(4)
		Panel A: Overall		
Female Headed HH	0.111***	0.084***	0.036***	0.022***
	(0.006)	(0.005)	(0.014)	(0.007)
		Panel B: Heterogeneity by Religion		
Female Head HH	0.106***	0.083***	0.028**	0.023***
	(0.007)	(0.006)	(0.014)	(0.008)
Female Head HH * Muslim	0.044**	0.011	0.099**	-0.005
	(0.019)	(0.018)	(0.047)	(0.019)
Muslim	0.001	-0.072***	0.005	0.005
	(0.009)	(0.011)	(0.016)	(0.005)
Mean Dep. Var.	0.16	0.59	0.12	0.01
Observations	51,541	70,369	10,342	10,342
Fixed Effect	Village	Village	Village	Village

Notes: The samples in columns (1) and (2) come from the NFHS 2015-16 dataset, while that in columns (3) and (4) come from the REDS 2006 dataset. Column (1) includes all married rural Hindu or Muslim women aged 15-49 while column (2) includes all rural Hindu or Muslim women aged 15-49. Columns (3) and (4) include all (rural) Hindu or Muslim women aged 18 years and above. *Decision* is an indicator variable that takes value one if a respondent can take any of these decision on her own: decisions about health care, decisions about making major household purchases, decisions about visits to family. *Allowed alone* is an indicator variable that takes value one if a respondent is allowed alone to any of these places: market, health facility, outside village. The dependent variable in column (3) is an indicator of whether a woman has attended any of the past four village meetings. The dependent variable in column (4) is an indicator of whether a woman is actively involved with a political party. *Female Headed HH* is an indicator that takes value one if the women respondent belongs to a household headed by a woman. Other controls in all the four columns include age, age squared, years of schooling and its square, religion, caste and village fixed effects. Robust standard errors clustered at village level are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.7: Female vs Male Headed Households

	(1) Poor	(2) Education of members	(3) Household members	(4) Age of members	(5) Child proportion	(6) Old proportion
			Panel A: Overall			
Female Headed HH	0.072***	-0.796***	-1.174***	1.724***	-0.011***	0.054***
	(0.002)	(0.014)	(0.010)	(0.079)	(0.001)	(0.001)
			Panel B: Heterogeneity by Religion			
Female Headed HH	0.074***	-0.816***	-1.160***	1.869***	-0.014***	0.057***
	(0.002)	(0.015)	(0.011)	(0.083)	(0.001)	(0.002)
Muslim*Female Headed HH	-0.019***	0.242***	-0.161***	-1.316***	0.023***	-0.024***
	(0.006)	(0.041)	(0.036)	(0.232)	(0.004)	(0.004)
Muslim	-0.008*	-0.904***	0.493***	-2.657***	0.045***	-0.020***
	(0.004)	(0.029)	(0.022)	(0.113)	(0.002)	(0.002)
Observations	425,563	425,532	425,563	425,549	425,563	425,563
Fixed Effect	Village	Village	Village	Village	Village	Village

Notes: The sample includes rural data from the NFHS 2015-16 dataset for all states. *Female Headed HH* is an indicator variable that takes value one if a household is headed by a woman, and is zero otherwise. *Poor* is an indicator variable that takes value one if a household belongs to the poorest two quintiles of the wealth index distribution. *Education of members* is average years of education of all household members. *Age of members* is average years of age of all household members. *Child proportion* and *Old proportion* refer to proportion of children (age 0-14) and old (age 60 and above) among all household members, respectively. *Muslim* is an indicator variable that takes value one if a household's head belongs to Muslim religion, else zero. Village fixed effects are included as a control in all specifications. Robust standard errors clustered at village level are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.8: Summary Statistics (GP Level)

	Overall (1)	Hindu Sarpanch (2)	Muslim Sarpanch (3)
	Panel A: Demographics		
Total Households	462.12 (342.52)	455.88 (334.90)	513.37 (396.13)
% APL hh	66.81 (27.34)	66.63 (27.15)	68.30 (28.84)
% Muslim hh	13.97 (18.48)	9.73 (11.51)	48.89 (26.35)
% Women headed hh	2.17 (4.80)	2.16 (4.78)	2.25 (4.97)
	Panel B: Toilet Ownership by Year		
% HH owning toilets end of 2012–13	29.97 (23.31)	29.39 (23.17)	34.69 (23.98)
% HH owning toilets end of 2013–14	30.59 (23.62)	29.99 (23.45)	35.53 (24.42)
% HH owning toilets end of 2014–15	32.55 (24.43)	31.86 (24.22)	38.21 (25.37)
% HH owning toilets end of 2015–16	35.08 (25.38)	34.29 (25.13)	41.52 (26.50)
% HH owning toilets end of 2016–17	42.25 (28.89)	41.41 (28.70)	49.12 (29.49)
Observations	54,012	48,153	5,851

Notes: The table reports the summary statistics at the GP level. Column (1) reports the results for the full sample of GPs, while columns (2) and (3) report it for samples with Hindu and Muslim Sarpanches respectively. Standard deviation reported in parentheses.

Table 2.9: Covariates and Pre-treatment Outcomes Balanced (GP)

	Overall (1)	Hindu Sarpanch (2)	Muslim Sarpanch (3)
	Panel A: Covariates		
Total population	84.17 (96.37)	106.7 (101.2)	-146.0 (315.0)
Primary school within 5 km	0.00970 (0.0175)	0.0143 (0.0184)	-0.0362 (0.0582)
Middle school within 5 km	0.0314 (0.0239)	0.0427* (0.0246)	-0.0955 (0.0981)
Secondary school within 5 km	0.0169 (0.0321)	0.00456 (0.0332)	0.154 (0.123)
Tap water	-0.0219 (0.0242)	-0.00779 (0.0252)	-0.181** (0.0894)
Closed drainage	0.00310 (0.0137)	0.00486 (0.0140)	-0.0138 (0.0559)
Waste disposal	-0.00379 (0.0173)	-0.00506 (0.0181)	0.0119 (0.0594)
All weather roads	0.0116 (0.0299)	-0.00388 (0.0310)	0.189* (0.113)
Domestic power	0.0115 (0.0177)	0.0168 (0.0180)	-0.0496 (0.0784)
Irrigation	0.00283 (0.0138)	-0.00402 (0.0143)	0.0840 (0.0544)
% APL hh	-0.00341 (0.0192)	0.00284 (0.0197)	-0.0707 (0.0795)
% Muslim hh	-0.00183 ((0.0113)	0.00335 (0.00764)	0.0114 (0.0649)
% Women headed hh	0.00200 (0.00326)	0.00141 (0.00326)	0.00835 (0.0156)
	Panel B: Pre-treatment Outcomes		
Covered to Uncovered 2013–14	-0.00344 (0.00412)	-0.00361 (0.00447)	-0.00153 (0.00571)
Covered to Uncovered 2014–15	-0.00558 (0.00640)	-0.00348 (0.00669)	-0.0256 (0.0223)
Covered to Uncovered 2015–16	0.00269 (0.00825)	0.00671 (0.00843)	-0.0388 (0.0355)
Polynomial order	1	1	1
Bandwidth	0.10	0.10	0.10

Notes: The variables from “Total Population” to “Irrigation” in Panel A are obtained using the Census 2011 data on village amenities. The remaining variables are obtained using the SBM administrative data. The rest of The polynomial order is 1. The bandwidths are manually chosen. Standard errors clustered at gram panchayat level and reported in parentheses.

Table 2.10: First Stage

	Female Sarpanch		
	Panel A: Overall		
	(1)	(2)	(3)
Female instrument	0.593***	0.563***	0.528***
	(0.0153)	(0.0177)	(0.0213)
Mean dep. var.	0.395	0.393	0.383
Observations	2,470,191	1,962,725	1,457,226
Number of GPs	9,179	7,234	5,277
Polynomial order	1	1	1
Bandwidth	0.100	.075	0.050
	Panel B: Hindu sarpanch		
	(1)	(2)	(3)
Female instrument	0.603***	0.572***	0.535***
	(0.0161)	(0.0186)	(0.0225)
Mean dep. var.	0.397	0.396	0.389
Observations	2,256,016	1,796,577	1,336,812
Number of GPs	8,278	6,541	4,744
Polynomial order	1	1	1
Bandwidth	0.100	.075	0.050
Estimated mean at the threshold	0.0548	0.0568	0.0590
	Panel C: Muslim sarpanch		
	(1)	(2)	(3)
Female instrument	0.504***	0.483***	0.465***
	(0.0487)	(0.0557)	(0.0663)
Mean dep. var.	0.376	0.361	0.315
Observations	214,175	166,148	120,414
Number of GPs	901	693	503
Polynomial order	1	1	1
Bandwidth	0.100	.075	0.050

Notes: The polynomial order is 1. The bandwidths are manually chosen. Standard errors clustered at gram panchayat level and reported in parentheses.

Table 2.11: Heterogeneous Effect of Female Reservation across GPs

	HH Received Toilet in 2016–2017		
	(1)	(2)	(3)
Female Reservation	0.00325	0.00383	0.00598
	(0.0166)	(0.0197)	(0.0246)
Muslim Sarpanch	-0.0332	-0.0539**	-0.0584*
	(0.0222)	(0.0254)	(0.0301)
Female Reservation*Muslim Sarpanch	0.149**	0.209**	0.257**
	(0.0678)	(0.0816)	(0.0999)
Mean dep. var.	0.104	0.104	0.098
Observations	2,470,191	1,962,725	1,457,226
Number of GPs	9,179	7,234	5,277
Bandwidth	0.100	0.075	0.050
Polynomial order	1	1	1

Notes: The sample is restricted to households which did not have toilets at the end of 2015–2016. The polynomial order is 1. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.12: Gender Quota Effect: Differentiating between Demand and Supply

	HH Received Toilet in 2016–2017			
	Hindu Sarpanch		Muslim Sarpanch	
	Hindu HH	Muslim HH	Hindu HH	Muslim HH
	(1)	(2)	(3)	(4)
	Panel A: Bandwidth 0.075			
Female reservation	0.0232	-0.0277	0.179	0.214**
	(0.0247)	(0.0587)	(0.111)	(0.0963)
	(0.0189)	(0.0403)	(0.0877)	(0.0820)
Female reservation * Women Headed HH	0.0532	0.298*	0.263	0.455***
	(0.0585)	(0.156)	(0.242)	(0.166)
Women Headed HH	-0.0298	-0.0995***	-0.105**	-0.0429
	(0.0230)	(0.0349)	(0.0489)	(0.0469)
Mean dep. var.	0.110	0.120	0.135	0.209
Observations	1,049,926	104,433	66,117	44,114
Number of GPs	6,325	5,542	656	647
Bandwidth	0.075	0.075	0.075	0.075
Polynomial order	1	1	1	1
	Panel B: Bandwidth 0.05			
Female reservation	0.0323	-0.00946	0.261*	0.320***
	(0.0305)	(0.0758)	(0.139)	(0.112)
Female reservation * Women Headed HH	0.0251	0.131	0.163	0.333*
	(0.0715)	(0.226)	(0.232)	(0.172)
Women Headed HH	-0.0122	-0.0470	-0.0956**	-0.0134
	(0.0257)	(0.0414)	(0.0457)	(0.0440)
Mean dep. var.	0.105	0.114	0.118	0.201
Observations	784,725	77,530	47,307	29,647
Number of GPs	4,622	4,055	474	466
Bandwidth	0.05	0.05	0.05	0.05
Polynomial order	1	1	1	1

Notes: The data is at the household level. The sample for column (1) is Hindu households under Hindu leaders while that for column (2) is Muslim households under Hindu leaders. The samples for columns (3) and (4) are defined similarly under Muslim leaders. Female Headed HH is a dummy that takes value one if the household head is a woman and zero otherwise. The polynomial order is 1. The bandwidth is manually chosen to be 0.075 in Panel A and 0.05 in Panel B. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.13: Gender Quota Effect: Identifying Household Demand (Common GPs)

	HH Received Toilet in 2016–2017			
	Hindu Sarpanch		Muslim Sarpanch	
	Hindu HH	Muslim HH	Hindu HH	Muslim HH
	(1)	(2)	(3)	(4)
Female reservation	0.00351	-0.0330	0.115	0.167*
	(0.0211)	(0.0475)	(0.0908)	(0.0868)
Female reservation * Female Headed HH	0.0597	0.254*	0.366	0.485***
	(0.0508)	(0.134)	(0.263)	(0.177)
Female Headed HH	-0.0346	-0.100***	-0.136*	-0.106*
	(0.0218)	(0.0302)	(0.0708)	(0.0612)
Mean dep. var.	0.111	0.120	0.135	0.201
Observations	1,280,507	131,297	83,457	57,085
Number of GPs	7,026	7,026	823	823
Bandwidth	0.1	0.1	0.1	0.1
Polynomial order	1	1	1	1

Notes: The sample is Above Poverty Line (APL) households who are eligible for the SBM program and did not have toilets at the end of 2015–2016. The sample for column (1) is Hindu households under Hindu leaders while that for column (2) is Muslim households under Hindu leaders. The samples for columns (3) and (4) are defined similarly under Muslim leaders. The sample of GPs for columns (1) and (2), and for columns (3) and (4) is common. Female Headed HH is a dummy that takes value one if the household head is a woman and zero otherwise. The polynomial order is 1. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.14: Gender Quota Effect: Identifying Household Demand (GPs with Hindu and Muslim Female Headed Households)

	HH Received Toilet in 2016–2017			
	Hindu Sarpanch		Muslim Sarpanch	
	Hindu HH (1)	Muslim HH (2)	Hindu HH (3)	Muslim HH (4)
Female reservation	-0.0505	-0.153	0.301**	0.274**
	(0.0471)	(0.152)	(0.146)	(0.115)
Female reservation * Female Headed HH	0.142**	0.371**	0.292	0.434**
	(0.0667)	(0.170)	(0.288)	(0.178)
Female Headed HH	-0.0354	-0.161***	-0.137	-0.0867
	(0.0288)	(0.0595)	(0.0856)	(0.0623)
Mean dep. var.	0.100	0.130	0.133	0.181
Observations	337,411	45,035	37,069	27,101
Number of GPs	1,443	1,443	276	276
Bandwidth	0.1	0.1	0.1	0.1
Polynomial order	1	1	1	1

Notes: The sample is Above Poverty Line (APL) households who are eligible for the SBM program and did not have toilets at the end of 2015–2016. The sample for column (1) is Hindu households under Hindu leaders while that for column (2) is Muslim households under Hindu leaders. The samples for columns (3) and (4) are defined similarly under Muslim leaders. The sample restricted to GPs having both Hindu and Muslim female headed households. Female Headed HH is a dummy that takes value one if the household head is a woman and zero otherwise. The polynomial order is 1. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.15: Muslim vs Hindu Sarpanches in Close Election GPs: Robustness

	HH Received Toilet in 2016–2017		
	(1)	(2)	(3)
Female Reservation	-0.00389	-0.00955	-0.00474
	(0.0275)	(0.0313)	(0.0385)
Muslim Sarpanch	-0.0136	-0.0172	-0.00501
	(0.0228)	(0.0258)	(0.0310)
Female Reservation*Muslim Sarpanch	-0.00557	-0.00264	-0.00338
	(0.0369)	(0.0418)	(0.0507)
Mean dep. var.	0.099	0.099	0.104
Observations	943,224	777,713	569,103
Number of GPs	3,262	2,666	1,941
Bandwidth	0.100	0.075	0.050
Polynomial order	1	1	1

Notes: The sample is restricted to households which did not have toilets at the end of 2015–2016. The polynomial order is 1. All regressions additionally control for the reservation assignment variable and its interaction with gender reservation. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 2.16: Gender Quota Effect in High Muslim Population Share GPs

	Household Received Toilet in 2016–2017			
	Hindu Sarpanch		Muslim Sarpanch	
	Hindu HH (1)	Muslim HH (2)	Hindu HH (3)	Muslim HH (4)
Female Sarpanch	0.0195	-0.0568	0.213*	0.189**
	(0.145)	(0.253)	(0.110)	(0.0941)
Female Sarpanch * Female Headed HH	0.509	0.906***	0.510*	0.508***
	(0.316)	(0.350)	(0.289)	(0.185)
Female Headed HH	-0.0825	-0.202*	-0.127	-0.0847
	(0.0567)	(0.105)	(0.137)	(0.0674)
Mean dep. var.	0.139	0.142	0.158	0.211
Observations	64,693	31,485	49,453	49,798
Number of GPs	554	532	613	616
Bandwidth	0.1	0.1	0.1	0.1
Polynomial order	1	1	1	1

Notes: The sample is Above Poverty Line (APL) households who are eligible for the SBM program and did not have toilets at the end of 2015–2016. Samples for all the columns only include GPs with population share of Muslims higher than the 85^{th} percentile of the Muslim population share distribution in the full sample. The sample for column (1) is Hindu households under Hindu leaders while that for column (2) is Muslim households under Hindu leaders. The samples for columns (3) and (4) are defined similarly under Muslim leaders. Female Headed HH is a dummy that takes value one if the household head is a woman and zero otherwise. The polynomial order is 1. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses.

Table 2.17: Close Election Gender Effect: Identifying Household Demand

	Household Received Toilet in 2016–2017			
	Hindu Female, any male		Muslim Female, any male	
	Hindu HH (1)	Muslim HH (2)	Hindu HH (3)	Muslim HH (4)
Female Sarpanch	0.0391**	0.0920**	0.0107	-0.000459
	(0.0193)	(0.0357)	(0.0762)	(0.0731)
Female Sarpanch * Female Headed HH	-0.0417	0.134	0.296*	0.636***
	(0.0507)	(0.124)	(0.164)	(0.205)
Female Headed HH	0.0682*	0.0376	0.0101	-0.142**
	(0.0364)	(0.0654)	(0.0940)	(0.0670)
Mean dep. var.	0.113	0.121	0.147	0.171
Observations	659,787	67,756	41,404	27,744
Number of GPs	4,116	3,646	423	405
Bandwidth	0.1	0.1	0.1	0.1
Polynomial order	1	1	1	1

Notes: The sample is Above Poverty Line (APL) households who are eligible for the SBM program and did not have toilets at the end of 2015–2016. The sample for column (1) is Hindu households under Hindu females, any males leaders while that for column (2) is Muslim households under Hindu females, any male leaders. The samples for columns (3) and (4) are defined similarly under Muslim female, any male leaders. Female Headed HH is a dummy that takes value one if the household head is a woman and zero otherwise. The polynomial order is 1. The bandwidth is manually chosen to be 0.1. Standard errors clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.18: No Gender Quota Effect for SC/ST Households

	Dep. Var.: HH Received Toilet in 2016–17					
	Hindu Sarpanch GPs			**Muslim Sarpanch GPs**		
	(1)	(2)	(3)	(4)	(5)	(6)
Female reservation	0.0109	0.0150	0.0201	0.165**	0.219**	0.312***
	(0.0240)	(0.0288)	(0.0360)	(0.0839)	(0.0975)	(0.118)
Female reservation * SC/ST HH	0.0257	0.0322	0.0428	-0.0774	-0.0847	-0.0885
	(0.0291)	(0.0341)	(0.0418)	(0.108)	(0.137)	(0.166)
SC/ST HH	-0.0492***	-0.0522***	-0.0591***	-0.0328	-0.0244	0.0150
	(0.0124)	(0.0138)	(0.0161)	(0.0341)	(0.0398)	(0.0407)
Mean dep. var. in control GPs	0.118	0.117	0.110	0.146	0.148	0.137
Observations	1,461,637	1,154,359	862,255	140,925	110,231	76,954
Number of GPs	8,025	6,339	4,632	862	666	481
Polynomial order	1	1	1	1	1	1
Bandwidth	0.100	.075	0.050	0.100	0.075	0.050

Notes: The dependent variable is a dummy that takes value one if the household received a toilet in 2016–17, and zero otherwise. The sample includes Above Poverty Line (APL) eligible households,i.e., those that did not have toilet at the end of 2015–16 and were eligible to receive toilet under the SBM scheme. The first and last three columns are households in GPs with Hindu and Muslim Sarpanches, respectively. The polynomial order is 1. Standard errors are clustered at Gram Panchayat level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.19: Gender Quota Effect in NREGS Implementation: No Heterogeneity

	Expenditure/Capita 2016–17			Person-days/Capita 2016–17		
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Hindu sarpanch					
Female reservation	0.00485	0.00562	0.0111	0.0277	0.0520	0.0771
	(0.0249)	(0.0300)	(0.0386)	(0.117)	(0.141)	(0.183)
Observations	7,834	6,195	4,520	7,891	6,238	4,548
Polynomial order	1	1	1	1	1	1
Bandwidth	0.100	.075	0.050	0.100	.075	0.050
Estimated mean at the threshold	0.290	0.294	0.297	1.292	1.300	1.309
	Panel B: Muslim sarpanch					
Female reservation	0.0490	0.0531	0.0524	0.327	0.261	0.340
	(0.0800)	(0.0972)	(0.119)	(0.353)	(0.438)	(0.541)
Observations	879	673	490	884	677	493
Polynomial order	1	1	1	1	1	1
Bandwidth	0.100	.075	0.050	0.100	.075	0.050
Estimated mean at the threshold	0.255	0.253	0.264	1.093	1.112	1.156

Notes: The expenditure reported is in thousands of INR. The polynomial order is 1. The bandwidths are manually chosen. Standard errors clustered at gram panchayat level and reported in parentheses.

Chapter 3

Words Matter: Gender, Jobs and Applicant Behavior

3.1 Introduction

Gender disparities in labor force participation and wages exist in both developed and developing countries, albeit to a varying degree. In developing country labor markets it is particularly difficult to investigate the causes and mechanisms behind such disparities due to a lack of large scale and high quality data on vacancies and job seekers. Recent studies (Kuhn and Shen, 2013; Helleseter et al., 2020; Ningrum et al., 2020) use online job portal data from China, Mexico and Indonesia to surmount these difficulties and document the prevalence of explicit gender requests in job ads to understand their role in observed gender disparities. In a recent paper, Kuhn et al. (2020) use data from China to explore how job seekers respond to explicit gender requests by employers. Nevertheless, much remains unknown about the role of explicit gender preferences in observed gender disparities within labor market outcomes—particularly in developing countries where gender inequalities may be far higher than in China. Moreover, employers can choose words within job ads to effectively recruit particular kinds of workers (e.g. men vs women). These words also reveal gender stereotypes associated with job roles held by employers. However, there is little research looking into how words in job ads are associated with particular stereotypes, how they relate to different characteristics of jobs and the posted wage, or how they direct the job search behavior of men and women who are looking for work.[1] We investigate these questions in this paper.

We use proprietary data from an online job portal in India to investigate explicit gender

[1] A detailed survey of search-theoretic models of the labor market is provided by Rogerson et al. (2005), while Wright et al. (2021) provide a recent survey of directed search.

preferences exhibited by employers in job ads, as well as their impact on job search behaviors of young and inexperienced job seekers. A distinctive feature of the Indian labor market today is the presence of large gender gaps in labor force participation (Fletcher et al., 2018). Despite high economic growth, increases in educational attainment and a decline in fertility over time, female labor force participation rates have remained stagnant among urban households (Klasen and Pieters, 2015; Afridi et al., 2018). In 2017–2018 only 20.6% of working age Indian women (age 15–65) in urban households were part of the labor force, compared to 78.9% of working age Indian men in urban households.[2] In fact, India's female labor force participation rates rank among the lowest in the world today and the country also has one of the largest gender wage gaps in the world.[3]

Though there are provisions within the Indian legal framework that could prohibit employers from posting job ads that explicitly request a male or female, the implementation of labor laws is generally inadequate.[4] Therefore, it is not unusual for employers to express explicit gender preferences in the job ads they post online. We examine such preferences by using data on 0.16 million job ads posted on an Indian job portal between July 2018 and February 2020 together with all applications made in response to these ads. We also use explicit gender preferences to derive an implicit gender preference (*femaleness* or *maleness*) signalled by the job text and examine their effect on wages and applications. Our estimations control for detailed occupation fixed effects derived from job titles, and thus exploit variation in explicit gender requests and implicit gender associations *within* an occupation and location (Indian states).

Around 8.3% of the job ads exhibit an explicit gender preference in the data, with slightly more ads exhibiting an explicit female preference than an explicit male preference. We find that jobs with an explicit gender preference tend to be low-skill jobs (in terms of lower education requirements and advertised wages). We also find that an employer's explicit female preferences dramatically reduce the number of applications to a job ad. At the same time, they increase the *share* of female applications by 16 percentage points (or 50%)

[2]Based on own calculations from the Periodic Labor Force Survey (PLFS) 2017–2018. Labor force status is defined using activity status over the previous year. A person is in the labor force if they are self-employed, an unpaid family worker, a regular salaried employee, a casual worker or unemployed.

[3]According to the World Bank, India's current female labor force participation rates are only better than those in Yemen, Iraq, Jordan, Syria, Algeria, Iran, and West Bank and Gaza and are comparable to Morocco, Afghanistan, Somalia, Pakistan, Egypt, Saudi Arabia, Lebanon and Tunisia. The ILO Global Wage Report 2018-19 documents the large gender wage gap in India.

[4]Article 16 of the Constitution of India prohibits discrimination on the basis of sex in public employment, while Article 39 guides the state to direct its policy towards ensuring "equal pay for equal work for both men and women". The Equal Remuneration Act, 1976 implements the provisions of Article 39 and prohibits sex based discrimination in payment of salary for same work (or work of similar nature) as well as in recruitment, promotion, training and transfer.

while an explicit male preference reduces this share by 9 percentage points (or 28%). Hence, explicit gender preferences have a substantial impact on the gender mix of the applicant pool. Jobs with an explicit female preference also have the lowest advertised wage, on average. Even among jobs without any explicit gender preference, we find that job ads containing text predictive of an employer's female preference (*femaleness*) have a significantly lower advertised wage. We also find that a higher fraction of women apply to these low-wage jobs.

We find that gender segregation can occur not only due to the employer's explicit gender preferences, but also because the applicant pool for a job may be proportionately larger for a given gender if the text contained in a job ad displays an implicit gender association. In line with findings in Kuhn et al. (2020) for China, we also find that women's share in the applicant pool increases when an explicit request for women is made than when it is not made for the same implicit gender association, showing that women may be more "ambiguity-averse". However, unlike China, we find that the gap in the share of female applicants who apply to female targeted ads and gender non-targeted ads initially increases and then remains constant as the implicit *femaleness* associated with a job ad increases. These results show that implicit gender associations contained in the job text matter, and that they matter for Indian women even when the ad includes an explicit preference for women. Lastly, we uncover gender word associations under the broad categories of hard skills, soft skills, personality and job flexibility to investigate how these affect applicant behavior. In general, we find that compliance with gender requests increases as words associated with skills, flexibility and personality associated with the requested gender increase and vice versa. For jobs that do not express any explicit gender request, skills and flexibility related words matter the most for increasing the female applicant share to a job ad.

Our work is inspired by a series of papers which investigate explicit gender preferences in the Chinese labor market (Kuhn and Shen, 2013; Helleseter et al., 2020). Kuhn and Shen (2013) were the first to document explicit gender preferences in job ads. They found evidence of a persistent negative skill-targeting relationship, i.e. as a job's skill requirements increased (as measured by required education, required experience or advertised wages) the share of ads expressing explicit gender preferences declined. Using additional data (including from a job board in Mexico) Helleseter et al. (2020) found that firms' explicit gender requests shifted from female to male workers as they sought older (vs younger) workers, a feature also referred to as the 'age twist'. We find a similar pattern in the urban Indian labor market regarding the negative skill-targeting relationship but we find only limited evidence of the age twist. Ningrum et al. (2020) also examine employer's explicit gender preferences using data from a job portal in Indonesia. In contrast to these papers, we derive implicit gender associations from explicit gender requests and are able to observe applicants' behavior in our data; this allows us to address additional important research questions such as how explicit gender requests by employers influence search behavior.

We contribute to the literature on several fronts. We make an important contribution to the large literature examining gender disparities and discrimination in the labor market. In addition, our work provides a first comprehensive analysis of gender targeting in job ads for India. In related work, Chowdhury et al. (2018) examine employer's gender preferences using 0.8 million job ads posted between 2011 and 2017 on the Indian job portal *Babajob* but do not use detailed occupation controls or data on applications. We use detailed occupation controls (with 501 occupation categories) derived using job titles in our estimations and make use of data on applications to examine how job seeker behavior is affected by explicit gender preferences. We also construct measures of implicit gender associations i.e. *femaleness* (*maleness*) from textual analysis of job ads to gain new insights into how the wording of a job ad influences labor market outcomes.

We build on work by Kuhn et al. (2020) who use applications data from an online job portal in China. We additionally employ several techniques from the literature on machine learning. First, we use a short text topic model to classify job titles to detailed occupation categories. Second, we construct implicit associations (*femaleness* and *maleness*) attached to a job from the text that appears in the job title and description using a different machine learning algorithm, which we argue improves upon their method. Third, we examine the effect of implicit gender associations on gender wage gaps.

This chapter also extends the recent literature on job attributes, wage penalty and the gender wage gap (Goldin and Katz, 2011; Goldin, 2014; Mas and Pallais, 2017; Bustelo et al., 2020; He et al., 2019). We derive job attributes associated with each gender using recent research in explainable artificial intelligence, hitherto not applied in the field of economics, to arrive at a word list that may be used to identify gender stereotypes. This word list is likely to be useful to researchers who are interested in uncovering gender associations in a similar labor market setting but who do not have information on explicit gender preferences. We use the word list derived from job text to examine how these affect applicant behavior and how the advertised wages vary by these attributes.

Lastly, and more generally, we add to a growing literature on various aspects of labor markets using high frequency data from online job portals. Hershbein and Kahn (2018) use data from job vacancies to investigate how skills demand changed over the Great Recession in the US. A number of recent papers examine high frequency changes in employment and skill demand with the onset of the Covid-19 pandemic by making use of job portal data.[5]

[5] These include Forsythe et al. (2020) for the US, Chiplunkar et al. (2020) for India, Hayashi and Matsuda (2020) for Bangladesh and Sri Lanka, and Campos-Vazquez et al. (2020) for Mexico.

While we use a similar data source, the research questions we study are quite distinctive from these papers.

Our work is also relevant for a growing empirical literature motivated by directed search models which investigates where job seekers send their applications. As opposed to random search models where wages are bargained ex-post and have no impact on the number of applicants, directed search models predict that job ads posting higher wages attract more applicants (Moen, 1997). Marinescu and Wolthoff (2020) use data from the US job portal *Careerbuilder* to find that job titles and posted wages affect the applicant pool that a firm attracts. Banfi and Villena-Roldan (2019) use data from a Chilean job portal to find that job ads with higher wages attract more applicants while Banfi et al. (2019) use the same dataset to document novel facts related to job search behavior of employed and unemployed job seekers. We also find that that higher posted wages attract more applicants in the urban Indian labor market. Moreover, we find that higher posted wages attract better or more able applicants which is consistent with the theory and evidence in Dal Bó et al. (2013) and Marinescu and Wolthoff (2020). A further implication of directed search models with heterogeneity is endogenous segmentation with applicants targeting their search towards sub-markets where they meet the selection criteria set by employers (Shi, 2002; Menzio et al., 2016). Our work confirms this by providing evidence that applicants direct their search based on the explicit and implicit gender requirements of employers as given in job ad text. Several studies have also used field experiments to examine similar research questions.[6]

This chapter provides a first comprehensive examination of the nature and consequences of explicit gender requests in job ads within the distinctive Indian context, where female labor force participation rates are low and gender disparities in the labor market are larger in comparison to China, Indonesia or Mexico. We construct implicit gender associations from textual analysis to further investigate how gender wage disparities and applicant behavior are shaped by the interaction of explicit gender requests *and* implicit gender associations. Our work is the first to identify words associated with explicit female and male gender preferences in job ads, using machine learning methods hitherto not applied in the field of economics. These words indicate underlying gender associations or stereotypes and are informative for researchers studying similar labor market contexts. Importantly and more generally, our results highlight how the text contained in a job ad matters for search behaviors.

[6] Belot et al. (2017) set up a field experiment to find that experimentally manipulated high wage jobs receive significantly more applications. Ibanez and Reiner (2018) examine the effect of affirmative action statements on application decisions using three field experiments in Colombia. Flory et al. (2015) use a field experiment to examine how workers' application decisions respond to competitive work environments while Mas and Pallais (2017) use a field experiment to examine how these decisions respond to non-wage job attributes.

In the next section we provide a detailed description of our data set, as well as different variables that we construct from the underlying job ad text. In Section 3.3 we discuss our empirical methodology while section 3.3.2 presents our estimation results. In Section 3.4 we describe the gender word associations we identify by carrying out further textual analysis and their effect on applicant behavior. Section 3.5 concludes.

3.2 Data

We analyse data from a leading job portal in India which primarily caters to young job seekers. Job seekers can create a profile for free and start applying to posted ads while employers need to pay a fee to post ads and view applicants (≈ USD 20). We use data on the population of jobs advertised on the portal with a last date of application between 24th July 2018 and 25th February 2020 together with data on all applications made to these ads. In our analysis we use data on 'active' job ads and job seekers; so we use job ads to which at least one male or female job seeker applied, and job seekers who applied to least one ad during this time.

Job seekers can view all jobs advertised on the portal and sort these by date of posting or popularity. They can also filter jobs based on job role, sector, location, education and type of job (govt/private). Job seekers who additionally register for a premium service are provided with specific job recommendations and alerts on new jobs by e-mail. The proportion of job seekers who registered for this service in our data was ≈ 0.5%; hence, the chances that applications are driven by matching algorithms used by the portal are negligible.

3.2.1 Job ads

There were a total of 196,821 job ads posted on the portal during this time. We exclude ads which had a location outside India and which had an application window of less than a day or more than four months (120 days), leaving us with 1,88,857 job ads. Next we drop duplicate job ads, where a duplicate ad was posted within a month of the original ad; this leaves us with 1,75,126 unique job ads.[7] When examining applicant behaviors we aggregate applications across duplicated ads to ensure we use data on *all* job seekers who apply to a job ad. We further drop job ads which had no male or female applicants which leaves us with 1,71,960 ads. We also restrict the sample to job ads that explicitly mention an education and experience requirement (which reduces the sample to 1,71,940 ads) and job

[7]Approximately 70% of duplicate job ads were posted within a month of the original ad. We keep duplicate job ads posted more than a month after the original job ad as separate job ads since these are likely to reference new vacancies.

ads that specify cities within a single Indian state as the location of the job (which reduces the sample to 1,58,249 ads).[8] We further restrict the sample to those jobs for which we obtain an occupational classification based on the method described in sub-section 3.2.2, leaving us with a final sample of 1,57,890 job ads.

We construct variables indicating an employer's gender and other requirements by carrying out a text search using the job title and description for each job ad in our sample.[9] In constructing a variable for gender requirement we make use of the text contained in the job title and description since words such as 'female only' or 'female preferred' (conversely 'male only' or 'male preferred') tend to appear here. We search the text for the following words which indicate an explicit female preference: 'female', 'females', 'woman', 'women', 'girl', 'girls', 'lady' or 'ladies'. Similarly, we undertake a search for the following words which indicate an explicit male preference: 'male', 'males', 'man', 'men', 'guy', 'guys', 'boy', 'boys', 'gent' and 'gents'. Some job ads include words related to both genders in the job title or job description. We categorise such job ads as having no explicit gender preference, together with ads which did not include words related to either gender. About 4.5% of the job ads in our sample have an explicit female preference (F jobs), 3.8% have an explicit male preference (M jobs) and the rest have no explicit gender preference (N jobs).[10]

To construct a variable indicating whether an employer has an age requirement, we split a job description if the following words appear: 'years of age', 'years old', 'years to', 'age', 'age limit'. We examine the 25 characters before and after the split. We search for numerals starting from 18 to 45 (since 45 is the maximum numeral found across all vacancies) among these characters and create variables for each number. If an ad has two numbers, the minimum of these is taken to specify the minimum age requirement and the maximum is taken to specify the maximum age requirement. In jobs where only one numeral appears, we combine it with words such as 'above', 'below', 'more than' and 'not above', 'not below', 'not less' to determine whether the age specified is a minimum or maximum required age.

Finally, we create dummy variables indicating the presence of a beauty requirement in an ad and indicating whether a job requires working a night shift since these features may be correlated with gender preferences in a job. To identify whether a job has a preference

[8] We find that restricting the sample of jobs to those that specify cities in a single state as the location of a job does not change the distribution of observable characteristics of the sample of job ads; this comparison is available on request.

[9] The portal does not have a separate field that allows employers to directly state the preferred gender for an advertised job to job seekers.

[10] The fraction of F and M jobs we find are smaller than those reported by Chowdhury et al. (2018) using data from *Babajob*. This could be because, unlike the job portal we use, *Babajob* had a separate field where employers could directly state the preferred gender to job seekers. Chowdhury et al. (2018) found that a third of all employers used this field. Of the total job ads on *Babajob*, 21% preferred men and 14% preferred women.

for beauty, we undertake a word search for 'height', 'weight', 'beautiful', 'charming', 'delightful', 'pretty', 'attractive' (ignoring a combination of words that specify an attractive salary or package), 'good looking', 'nice looking', 'complexion', 'pleasing', 'appearance' and 'handsome' within the job description of an ad.[11] Similarly, we create a dummy variable for the presence of a night shift requirement in a job ad by carrying out a word search for: 'night-shift', 'night shift' or 'night'; we exclude ads if the job description specifically mentions 'no night shift' or that transportation will be made available like 'night shift fully secured', 'cab drop', 'drop facility', 'cab facility', 'dropping available', 'pick' and 'drop'.

A very small fraction of jobs advertised on the portal either did not specify an education requirement or specified it as none (or illiterate). We keep these ads in our estimation sample, and group these together with ads requiring a secondary education or less as the base category in our empirical analysis. In general, N jobs are *less* likely to require only a secondary and senior secondary education (as opposed to higher education categories of graduate and postgraduate) than F or M jobs (Appendix Table 3.8). A lower fraction of M jobs require a graduate or postgraduate degree. Consistent with the portal catering primarily to young job seekers, we find that most job ads (at $\approx$ 67%) require less than one year of experience. We also find that N jobs are more likely to list two or more years of experience and M jobs are somewhat more likely to require a higher experience than F jobs. Ads specifying a gender preference are also more likely to specify other preferences, such as those related to age or beauty. We find that M jobs are more likely to also specify an age preference and these jobs specify a higher minimum and maximum required age than F or N jobs (this may also be seen in Appendix Figure 3.3). Finally, F jobs are most likely to also specify a beauty requirement while M jobs are most likely to require working night shifts.

Employers include a wage range in 88% of the job ads advertised on the portal that are in our sample. We find that wages are more likely to be missing for jobs requiring higher education and experience. Thus, the sample of job ads with wage information is a somewhat selected sample of lower skill jobs; nevertheless, we are able to observe wages for a far higher fraction of job ads in our sample than existing studies.[12] The mean of the mid-point of the wage range is 221 thousand rupees per year. This is higher than the national average of

[11] To find words related to beauty, we started out with an initial list of beauty related words such as 'beautiful' and 'handsome'. We then appended to this list by considering the cosine similarity of vector representation of these words with other words using the unsupervised GloVe algorithm (Pennington et al., 2014). The 300 dimensional pre-trained word vectors have been obtained by training the algorithm on web data from common crawl, and comprise 2.2 million unique words. Cosine similarity between any two vectors is a score $\in [0, 1]$, which in this case indicates the relatedness of any two words in terms of the context in which they appear on the internet, and can essentially help identify synonyms.

[12] In comparison wages are advertised in just 16.4% of job ads in Kuhn and Shen (2013) using a Chinese job portal and 20% of job ads in Marinescu and Wolthoff (2020) using *Careerbuilder*.

salaries earned by urban Indian workers with an age distribution similar to candidates on the portal, indicating that our estimation sample consists of relatively high skill urban jobs. N jobs have the highest mean wage while M jobs have a higher mean wage than F jobs, a pattern that can also be seen in Appendix Figure 3.4.[13]

The share of female applicants to N jobs is 32%. This is because, as shown in Appendix Table 3.9, there are fewer female applicants on the portal compared to male applicants. For F jobs this share rises to 51% while for M jobs it falls to 14%. This indicates that there is some compliance with explicit gender requirements in job ads but that this compliance is far from perfect. The overall compliance with gender in F and M jobs i.e. percent applications that are of the requested gender is 68%. In order to account for compliance that can occur by chance (expected compliance) due to the distribution of job and candidate characteristics on the portal, we use Cohen's kappa.[14] Cohen's kappa κ for compliance with gender requirements is 32%. Compliance with other requirements such as education and experience i.e. percent applications that have at least as much education or experience as requested, across jobs ads, is 98% ($\kappa = 97\%$) and 32% ($\kappa = 25.47\%$) respectively.[15] Thus, compliance with gender requirements is lower than with education requirements but higher than with experience requirements.

There are about 42 applications per ad, on average. The average number of applications to F jobs is less than half of this, at about 18, while the average number of applications to M jobs is about 32. This indicates that explicit gender preferences lead to a substantial reduction in the number of applications, particularly by job seekers of the opposite gender to the preferred one.

3.2.2 Job titles and occupations

Job ads also include information on which role a particular job belongs to, out of 33 job roles pre-specified by the portal.[16] However, these job roles are too coarse to characterize

[13] Though there are non-wage elements that may be crucial for applications such as allowances, paid and unpaid leaves, work-from-home etc. we cannot directly account for them because of a lack of systematic data on these variables. We, however, are able to discuss some of these aspects in Section 3.4 where we discuss gendered words in job descriptions.

[14] Cohen's kappa is defined as $\kappa \equiv \frac{Compliance_{observed} - Compliance_{expected}}{1 - Compliance_{expected}}$. The component of compliance on gender that is expected to occur merely by chance is 53%.

[15] These compliance figures are calculated after dropping jobs which have no education requirement and a minimum experience requirement of 0 years respectively.

[16] Appendix Figure 3.5 gives the relative percentage of ads among F, N and M jobs, across these job roles, helping us understand broad patterns of gender segregation (we group the category 'Physio Therapist' with 'Others' as there were only 7 job ads with this job role). It shows that a relatively high fraction of F jobs belong to job roles 'Beautician/Spa', 'Receptionist/Front Office', and 'Teacher/Trainer' while a high fraction of M jobs belong to job roles 'Delivery Executive', 'IT Hardware Engineer' and 'Mechanic/Fitter/Production'.

occupation for a job ad. Marinescu and Wolthoff (2020) have shown that job titles can provide a much finer classification of occupations since titles not only capture the job role, but also the hierarchy and specialization within a role. In line with this, Marinescu and Wolthoff (2020) find that words contained in job titles are predictive of wages as well as applications using data from *Careerbuilder* in the US. Figure 3.1 shows word clouds of job titles in F, M and N jobs in our sample. As may be seen, job titles such as 'telecaller' and 'office executive' occur with high frequency among F jobs while titles such as 'delivery boy' and 'sales executive' occur with high frequency among M jobs. This confirms that explicit gender preferences operate to maintain existing occupational gender stereotypes.

However, job titles might also reflect noise in the word choice of an ad without any meaningful differentiation. Therefore, we additionally use an unsupervised machine learning technique to classify semantically similar job titles into disaggregate occupation categories. Specifically, we use the collapsed Gibbs Sampling algorithm for the Dirichlet Multinomial Mixture model (GSDMM) proposed by Yin and Wang (2014), and apply it to the text contained in job titles.[17] Qiang et al. (2020) show that GSDMM is very effective for short text topic modeling and outperforms Latent Dirichlet Allocation (LDA) and several other methods at this task. GSDMM makes the assumption that each document (or in our case, job title) comprises a single topic. This assumption is suitable for short texts which are unlikely to contain multiple topics. Each document is conceptualized as a set of words. The GSDMM algorithm probabilistically combines documents into groups such that documents in the same group contain a similar set of words, whereas documents in different groups contain a different set of words. We relegate the details of the data pre-processing steps, the algorithm, and our hyperparameter choice to the Technical Appendix section 3.A.1. The final number of topics (or dis-aggregate occupation categories) discovered by GSDMM for our sample of job ads is 501.

Our empirical results are largely robust to an alternative manual clustering of job ads based on the existence of word unigrams, bigrams and trigrams in job titles which has been used in the literature (Marinescu and Wolthoff, 2020; Banfi and Villena-Roldan, 2019).[18] To implement the alternative manual clustering we calculate n-gram counts after removing duplicate job ads.[19] We then classify all jobs on the basis of the most frequently occurring trigrams in job titles, subject to the trigram existing in at least 50 job titles. We then classify

This confirms that explicit gender preferences operate to maintain existing (and broadly defined) occupational gender stereotypes.

[17] We use the python implementation of GSDMM available at https://github.com/rwalk/gsdmm.

[18] We discuss estimation results using the alternative categorization in sub-section 3.3.4.

[19] For the purpose of this classification job ads made by the same employer, with the same job title and job description are considered duplicates.

the remaining jobs based on the most frequently occurring bigrams and then unigrams in the job title, with the restriction that the bigrams and unigrams occur in at least 100 job titles. The precedence given to higher order n-gram followed by their frequency of occurrence ensures that each job posting is classified into no more than one cluster or occupation group. This way we obtain a total of 747 occupation groups.[20]

We prefer the short text topic model to the manual classification of job titles, since it provides dimension reduction based on co-occurrence of words in the corpus of job titles. The algorithm accomplishes this by probabilistically clustering together documents which do not share any common word between them, but are linked together through sharing common word(s) with some other documents that act as a bridge between the two. For instance, the jobs titled 'english transcriber' and 'japanese translator' are assigned the same cluster as they are linked through 'transcriber translator'. These jobs cannot be assigned the same cluster using the manual classification as they do not share any common word. This also ensures that most of the job ads in the topic model get assigned to meaningful clusters. In contrast, over 5,800 jobs could not be assigned to any cluster using the manual classification because of the low frequency of word n-grams contained in them across the corpus.[21]

3.2.3 Job seekers

We also use data on 1.06 million job seekers who applied to at least one ad using the portal; descriptive statistics for job seekers by gender are given in Appendix Table 3.9. There are 0.37 million female and 0.68 million male job seekers. The smaller number of female job seekers is consistent with lower female labor force participation rates in urban India compared to males (Appendix Table 3.11). Notably, while the labor force participation rate of men is about three times that of women, there are only slightly less than twice as many male job seekers on the portal as female job seekers. Also, once female job seekers start searching

[20] We could not classify around 3% of job postings to any occupation using this method.

[21] There is no direct way to assess objectively whether short text topic model or manual clustering performs better. Existing measures such as homogeneity and completeness used in the literature are not appropriate in our context since the true occupation categories are not known. The variable depicting job roles (see Figure 3.5 for the job roles in the data) has very few categories to reflect true occupation categorization. In many cases two jobs involving similar tasks can often be assigned two or three different job roles. For example, the job ads titled "customer care executive" and "customer care professional" are both assigned job roles "BPO/Telecaller" as well as "Customer Service/Tech Support". While our topic model assigns them to the same cluster, the manual classification assigns them to different topics—"customer care executive" and "customer care" respectively. Similarly, "software engineer" and "software test engineer" are both assigned job roles "IT Software Engineer" as well as "Engineer (Core, Non IT)". These are assigned to same cluster by our topic model, but again assigned different occupations by the manual classification. Therefore, job role is an imperfect gold standard for measuring homogeneity. Nonetheless, we compute the homogeneity score and find that it has a value of 74.44% for the short text model. This indicates that job ads within a cluster largely belong to the same job role.

for jobs using the portal, they make a similar number of job applications, on average, as male job seekers. Most job seekers on the portal (or more than 85%) have a graduate or post-graduate degree with female job seekers being more likely to have higher education. Job seekers are also relatively young, with an average age of 24 years; female job seekers are slightly younger than male job seekers and have less experience. About 76% of job seekers have less than a year of experience, again indicating that the portal caters primarily to young job seekers. Lastly, female job seekers (unconditionally) apply to job ads with slightly higher posted wages than male job seekers. However, female job seekers tend to have more education than male job seekers. In fact, conditional on candidate characteristics, women apply to job ads with 3% lower posted annual wage than men.[22]

We compare job seekers on the portal with the urban working age population in India using the Periodic Labor Force survey (PLFS) from 2017–18, which is a nationally representative survey of employment in India. Appendix Table 3.11 Panel A uses the PLFS to give the average annual earnings for those in casual or salaried employment among working age adults (age 16-60) in urban Indian districts (with $\geq 70\%$ urban population).[23] Advertised wages on the portal are higher than this nationally representative sample by $\approx$ Rs. 20,000 per annum. However, wages in PLFS could also be high because it has older, more experienced workers. To make the PLFS sample comparable to the age group catered to by the online job portal we only keep adults who are 18-32 years old in Appendix Table 3.11 Panel B.[24] The gap in annual earnings increases to Rs 45,000 per annum, or the average advertised wage is now 25% higher on the job portal. Thus, the job portal caters to younger, inexperienced but more educated and skilled workers in urban India.[25]

[22]We also estimate a different specification at the application rather than candidate level to estimate the gender wage gap in applications in which we also control for occupation. We regress the log of the posted wage for the applied job on job ad and candidate characteristics, giving each candidate equal weight, and find that women, on average, apply to jobs with 1.8% lower wages than men.

[23]Annual earnings are obtained by multiplying monthly earnings by 12 for salaried workers and weekly earnings by 52 for daily wage workers.

[24]Approximately 95% of the job seekers on the portal are 18-32 years old.

[25]In the sample of workers in PLFS survey men earn higher a annual wage than women. We regress the log of wage on a gender indicator and find that women earn 18% lower annual wages than men. To keep a comparable sample to the portal, we then keep only the candidates having more than school education and those aged between 18-32. Controlling for worker education, age and occupation this wage difference then decreases to 8%. On the portal, women applicants apply to jobs that on an average offer 1.8% lower annual wages after including occupation times location controls. These findings show that the gender wage gap among educated individuals in the Indian labor market is partly driven by applications of female job seekers to lower wage jobs.

3.2.4 Implicit *femaleness* and *maleness*

The text contained in a job ad may also convey an implicit signal to a candidate about whether the employer posting the ad prefers a female or a male candidate for the job even in the absence of an explicit gender preference. Similar to Kuhn et al. (2020), we define the implicit "*femaleness*" (F_p) and "*maleness*" (M_p) of a job as:

$$F_p \equiv \text{Prob(explicit female request} \mid \text{job text)}$$

$$M_p \equiv \text{Prob(explicit male request} \mid \text{job text)}$$

We use a supervised machine learning approach to infer F_p and M_p associated with each job ad based on the job text. We implement a Logistic Regression (LR) classifier with balanced class weights.[26] For this purpose, we concatenate the job title and the job description to reflect the complete job text.[27] We follow the standard pre-processing steps as in the natural language processing literature which we outline in the Technical Appendix 3.A.2. We then convert our corpus of processed documents to their bag-of-n-grams representation using term frequency-inverse document frequency ($TF - IDF$) vectors—which we use as inputs to the model. $TF - IDF$ captures how important a token (or a set of words) is to a document with respect to its importance in the corpus based on its frequency. Therefore, it improves text classification by scaling down the weights of common tokens which are likely to be uninformative in capturing employers' preferences. We consider word unigrams, bigrams and trigrams, i.e., $n \in \{1, 2, 3\}$. The output class in the model takes the value one if a job ad shows an explicit preference for females (males) and 0 otherwise, when calculating F_p (M_p). We perform stratified 10-folds cross-validation wherein we split the data in 10 parts and preserve the percentage of sample that belongs to each class. Each document will belong to the test set only once. F_p and M_p are then the estimated probabilities of a document belonging to the positive class when it belongs to the test set.

Conditional on an employer making an explicit gender request, the model correctly predicts requests for females and males in 78.07% and 75.13% of job ads when they are part of the test set. The corresponding figures when employers do not explicitly request a gender are 90.74% and 94.73% in the model for F_p and M_p respectively. Furthermore, F_p and M_p capture employer requests very well, with correlations of 0.46 and 0.52 with binary variables capturing explicit female and male requests respectively.[28] F_p takes high values for jobs with

[26] We use balanced class weights since the classes are highly imbalanced with only a small fraction of total jobs explicitly requesting a female or a male.

[27] We use a total 196,857 jobs which include an additional set of jobs provided to us by the portal to increase data points for the classification model.

[28] The correlations using balanced class weights, i.e. weighted by inverse frequency of observations belong-

titles such as ‘receptionist’ and low values for jobs with titles such as ‘mechanical engineer’ while M_p takes high values for jobs with titles such as ‘electrician’ and low values for jobs with titles such as ‘telecaller’.

We differ from Kuhn et al. (2020) in calculating F_p and M_p in several ways. First, as opposed to using only the job title, we make use of both the job title as well as job description in predicting implicit gender association of a job posting. This is reasonable as we expect that candidates will make use of both to infer their hiring prospects, as well as suitability, when applying to a particular job. This is validated when we include occupation fixed effects based on job titles in the specification where we infer the impact of F_p and M_p on applicant behavior (sub-section 3.3.3); including occupation fixed effects based on job titles reduces the responsiveness of job seekers to implicit gender associations but does not eliminate it. This shows that candidates use information contained in the job description in conjunction with job titles to infer their suitability for a job.

Second, and perhaps more importantly, we are able to improve our measures of F_p and M_p using an LR classifier instead of the Bernoulli Naive Bayes (NB) classifier. To compare the two we implement the NB classifier as well on our data using the methodology in Kuhn et al. (2020). We find that the NB classifier does not perform as well in our context. It gives a much worse measure of F_p and M_p in our data with correlations of 0.23 and 0.22 with explicit employer requests for women and men.[29] LR also does a better job of explaining applicant behaviors, as discussed in section 3.3.2. These results demonstrate the fact that even though NB is a reasonable classifier, it does a poor job of estimating probabilities associated with the classes.

3.3 Explicit gender preferences

3.3.1 Empirical methodology

We first examine characteristics of jobs where employers exhibit explicit gender preferences. The regressions we estimate are variations of the following specification:

$$Y_{ijst}^{k} = \alpha^{k} + \beta^{k} X_{ijst} + \gamma_{j\times s} + \phi_t + \epsilon_{ijst}^{k} \tag{3.3.1}$$

ing to the two classes are 0.75 and 0.76 respectively.

[29] The correlations using balanced class weights are 0.51 and 0.51. The requests for females and males conditional on an employers’ explicit gender request are correctly predicted in 71.92% and 73.33% of job ads. For jobs that make no explicit gender request, correct predictions are made in 73.90% and 69.43% job ads in the model for F_p and M_p respectively. These figures are much lower than those for the LR model discussed previously.

where the superscript $k \in \{FM, M\}$ indicates two different dependent variables capturing the *presence* and *direction* of gender preferences. The first dependent variable Y_{ijst}^{FM} is a binary outcome which takes the value one if there is either an explicit male or female preference exhibited in job ad i which advertises for a job of occupation j in state s at time (or month and year) t. The second dependent variable Y_{ijst}^{M} takes on three values: minus one if there is an explicit female preference, zero if there is no gender preference and one if there is an explicit male preference exhibited in a job ad. X_{ijst} is a set of job ad specific variables including a set of dummy variables for education requirements, a set of dummy variables for experience requirements, dummy variables for the presence of age and beauty requirements, a dummy variable for the presence of a night shift requirement and a quadratic in log advertised wage. In our preferred specification we include occupation and state fixed effects ($\gamma_{j\times s}$) as well as (month, year) fixed effects (ϕ_t). We employ a detailed categorisation of jobs to occupations, as described in sub-section 3.2.2, with 501 distinct occupation categories derived from job titles. The use of fixed effects ensures we use *within* occupation and state variation only to identify the effect of different variables on whether a vacancy exhibits a gender (or male) preference. We cluster standard errors by occupation and state.

The wage difference between job ads that explicitly request men or women can be obtained from estimates of equation 3.3.1 when Y_{ijst}^{M} is the dependent variable. In a separate set of regressions we also examine whether wage differences exist when the text of a job ad is predictive of explicit gender preferences by the employer (sub-section 3.2.4), separately for F, N and M jobs. The regressions we estimate are variations of the following specification:

$$\ln W_{ijst} = \alpha^W + \lambda^W F_{p,ijst} + \nu^W M_{p,ijst} + \beta^W X_{ijst} + \gamma_{j\times s} + \phi_t + \varepsilon_{ijst} \tag{3.3.2}$$

where $\ln W_{ijst}$ is the log wage in a job ad, $F_{p,ijst}$ is a measure of implicit *femaleness* and $M_{p,ijst}$ is a measure of implicit *maleness*. The coefficients on these variables (λ^W and ν^W) tell us how the advertised log wage changes as predicted *femaleness* (*maleness*) of a job ad increases by one i.e. the probability of a job being a stereotypical female (male) increases by one, everything else equal. X_{ijst} is a set of job ad specific variables (dummy variables for education and experience requirements).[30] In our preferred specification we include occupation and state fixed effects ($\gamma_{j\times s}$) as well as (month, year) fixed effects (ϕ_t). This allows us to control for wage differentials across occupations in the same location. As before, we cluster standard errors by occupation and state.

We also examine how explicit gender preferences affect job seeker's responses to a job ad; we first estimate variations of the following specification:

[30]We do not include dummies for the presence of age, beauty or a night shift requirement since the words used to construct these variables are also highly predictive of an explicit gender preference.

$$Y_{ijst}^{TA} = \alpha^{TA} + \pi^{TA} F_{ijst} + \theta^{TA} M_{ijst} + \beta^{TA} X_{ijst} + \gamma_{j\times s} + \phi_t + \mu_{ijst} \tag{3.3.3}$$

where Y_{ijst}^{TA} is the total number of applications to a job ad. F_{ijst} is a binary variable taking the value one if ad i has an explicit female preference and zero otherwise. Similarly, M_{ijst} is a binary variable taking the value one if ad i has an explicit male preference, and zero otherwise. The coefficients on the binary variables (π^{TA} and θ^{TA}) give the difference in total applications sent to ads that exhibit an explicit female or male preference in comparison to ads that exhibit no such preference (the base category), everything else equal. X_{ijst} is a set of job ad specific variables which are the same as those specified in equation (3.3.1). In our preferred specification we include occupation and state fixed effects ($\gamma_{j\times s}$) as well as (month, year) fixed effects (ϕ_t); as before, we also cluster standard errors by occupation and state.

In order to examine job seekers compliance with the gender preference requirement set by the employer, we estimate variations of the following specification:

$$Y_{ijst}^{C} = \alpha^{C} + \pi^{C} F_{ijst} + \theta^{C} M_{ijst} + \beta^{C} X_{ijst} + \gamma_{j\times s} + \phi_t + \mu_{ijst} \tag{3.3.4}$$

where Y_{ijst}^{C} is the share of female applicants to a job ad. Apart from the difference in outcomes, the specification of equation (3.3.4) is similar to (3.3.3); coefficients on the binary variables (π^{C} and θ^{C}) give the difference in the share of female applicants across ads that exhibit an explicit female or male preference and those that exhibit no such preference (the base category), everything else equal. Since we exploit *within* occupation variation, our estimates capture causal effects of gender requests within an occupation on the share of female applicants. These regressions are also weighted by the total number of male and female applications made to a job ad. In additional estimations we report the effect of explicit gender preferences on applicant and match quality by using specifications similar to (3.3.4) but with applicant and match quality as the dependent variables of interest.

3.3.2 Results

Table 3.1 gives estimation results from estimation of equation (3.3.1) when the dependent variable is Y_{ijst}^{FM}. Column (I) includes education and experience requirements only as the set of explanatory variables as well as (month, year) fixed effects. Column (II) adds occupation and state fixed effects to the estimations; columns (III) and (IV) add more explanatory variables. Since wages are not advertised for all jobs, some observations are lost when including a quadratic in log advertised wage as explanatory variables in column (IV). The results support a negative skill-targeting relationship even within the same occupation and state i.e. jobs with a higher skill requirement (a higher education requirement or log advertised wage) are *less* likely to have an explicit gender preference. We find mixed results for experience. When occupation fixed effects are not included (column (I)), jobs that specify a higher experience category are associated with a lower probability of exhibiting a gender preference. In column (IV), once occupation and wage controls are included, higher experience is associated with an increased probability of exhibiting an explicit gender preference. This reversal primarily occurs due to inclusion of controls for advertised wage. Experience is positively correlated with wage and wages have a negative, though insignificant, effect on the probability of a job ad exhibiting a gender preference. We also find that the presence of an age requirement, a beauty requirement or working night shift requirement leads to an increased probability that a vacancy has an explicit gender preference (columns (III) and (IV), Table 3.1).

Table 3.2 gives estimation results when the outcome of interest is male preference in a job ad. Here we find that jobs having an explicit male preference are less likely to belong to higher education categories (compared to the base category), but are more likely to require higher experience ($>$ 2 years, column (III), Table 3.2). However, the experience requirement becomes insignificant once we further control for a quadratic in log wage. Jobs with an explicit male preference also offer higher wages than those with an explicit female preference; this is evident from our finding that a higher advertised wage is associated with an increased preference for men. We also find that the presence of age requirements and working night shifts leads to increased preferences for men while the presence of a beauty requirement leads to a reduced preference for men (columns (III) and (IV), Table 3.2).[31]

The results in Table 3.2 show that jobs with an explicit preference for men offer higher wages than jobs with an explicit preference for women. We next examine estimation results of equation (3.3.2) (estimated separately for F, N and M jobs) to evaluate the effect of the *femaleness* (*maleness*) measures on the advertised wage; the results are reported in Table 3.3. As expected, higher education and experience requirements lead to an increase in the advertised wage (for all kinds of jobs). For N jobs we find that an increase in *femaleness* from 0 to 1 leads to a reduction in the offered wage by 39%, with no occupation controls (column (III)). Once detailed occupation and location controls are added, the effect of *femaleness* on offered wages drops to 29% but remains highly statistically significant (column (IV), Table 3.3). This coefficient estimate translates to a decrease in advertised wage of 5.8% for a one standard deviation increase in the *femaleness* measure ($SD = 0.2$). On the other hand, an

[31] To further investigate whether a male preference in a vacancy is associated with a higher maximum age requirement (or to check for evidence of the 'age twist' in explicit gender preferences) we also estimated regressions on the sub-set of vacancies which specify a maximum required age and used maximum required age instead of any age requirement as the explanatory variable of interest. While we found a positive effect of maximum required age on preference for men, we did not find that this effect to be statistically significant. These results are available on request.

increase in *maleness* is associated with a small change in the log wage which is statistically indistinguishable from zero. This provides evidence that jobs with higher female association (or jobs where applicants are likely to infer that the employer would prefer a female from reading the job text) offer systematically lower wages even when the job does not exhibit any explicit gender preference. We find a similar pattern for F and M jobs, but the negative effect of the *femaleness* measure on log wage is smaller in these jobs than for N jobs, although it is still statistically significant. In later results we also find that a higher share of women apply to these low-wage jobs (column (VI), Table 3.4).

To examine the effect of explicit gender preferences on applicant behaviors we estimate and report regressions specified by equation (3.3.3) where the outcome variable is the total number of applications to a job ad; the results are reported in columns (I)-(III) of Table 3.4. We find that the number of applications are reduced by a statistically significant number ($\approx 6-8$) if the job ad exhibits an explicit female preference. On the other hand, the change in number of applications to job ads that exhibit an explicit male preference compared to no gender preference is generally not statistically significantly different from zero, although there is a large reduction (≈ 4.6) once we include a full set of controls such as a quadratic in log advertised wage (column (III), Table 3.4).

We also find that the number of applications to jobs with higher education requirements tend to increase and then fall. Notably, job postings which specify a graduate degree in a STEM subject see the largest increase in the number of applications compared to the base category —at 58 (column (III), Table 3.4). However, the number of applications to jobs that require a post graduate degree (in either a STEM or non-STEM subject) is reduced once we include a full set of controls, although the coefficient is not always statistically significant. These results reflect the fact that the number of graduate job seekers with a STEM degree in the Indian urban labor market outnumber the number of jobs which require a specialization in STEM. 54% of the applicants on the portal have a graduate degree in STEM (Table 3.9) while the number of jobs that require only graduates with a STEM degree are 9%. Hence, these fewer jobs attract a much larger applicant pool.

The number of applications to job ads with higher experience requirements is much smaller, and these coefficients are always statistically significant. The number of applications to jobs that specify an age requirement or a beauty requirement is also reduced, while there isn't a statistically significant effect of advertised wages on the number of applicants. This could possibly reflect a slack youth labor market in India where over the last decade the unemployment rate, especially among the educated youth, has been on a rise.[32]

Next, we estimate and report the regressions specified by equation (3.3.4) where the

[32]The unemployment rate for urban young men reached 18% in 2017-18. See: Mint Report

outcome is the fraction of female applicants to a job ad; the results are reported in columns (IV)-(VI) of Table 3.4. We find that the fraction of female applicants to a vacancy increases by $15.7 - 16.1$ percentage points when the vacancy exhibits an explicit female preference and reduces by $8.8 - 9.0$ percentage points when the vacancy exhibits an explicit male preference. These translate to an increase of 50% and decrease of 28% in the share of female applicants to a vacancy, which are substantially large effects. In addition we find that a higher fraction of women apply to vacancies which have higher education and lower experience requirements. This is likely to be driven by more educated and younger women on the portal (Appendix Table 3.9). A smaller fraction of women apply to vacancies which specify an age requirement or working night shifts as part of the job description. We also find that a statistically significant higher fraction of women apply to vacancies with a lower advertised wage. This could very well be driven by *femaleness* association with a job, strikingly even within occupations, which are likely to attract more female applicants and which advertise significantly lower wages.

We also examine the effect of explicit gender requests on applicant and match quality; the results are reported in Appendix Table 3.10. We use two measures of applicant quality—completed years of schooling and the percentage marks (out of 100) obtained by a candidate in secondary school (matriculation) examination.[33] To the extent that gender requests are likely to be made in low skill jobs, applicant quality can be lower. On the other hand, gender requests may decrease or increase applicant quality by deterring otherwise highly or less qualified candidates of the non-preferred gender. We find that explicit male requests are associated with a decline in applicant quality; however, we find that explicit female requests are not associated with a reduction in applicant quality as measured by completed years of education and are actually associated with a statistically significant increase in applicant quality as measured by matriculation marks, although the effect size is economically small (Columns (I) and (IV)). These associations are driven by a higher fraction of female applicants applying to jobs with an explicit female preference (on average women on the portal are more educated and have higher matriculation scores than men). Once we control for the share of female applicants, an explicit female preference has a larger negative impact on applicant quality than an explicit male preference (Columns (II) and (V)). However, these effects continue to be economically very small. At the same time, we find that applicant quality improves with higher posted wages, consistent with the theory and evidence in Dal Bó et al. (2013) and Marinescu and Wolthoff (2020) (Columns (III) and (VI)). We also find the effect of explicit gender requests on match quality (in terms

[33] We do not use experience as a measure of quality since the portal mostly caters to inexperienced graduates.

of the share of applicants complying with the job ad's minimum education and experience requirements) to be economically small (Columns (VII) and (VIII)).

3.3.3 Do implicit gender associations matter for applications?

We find that explicit gender requests in a job ad lead to differential application rates by men and women, though compliance is not perfect. Next, we examine how F_p (M_p) derived from the text of a job ad affect applicant behaviors. Remember, though F and M indicate the presence of explicit female and male requests in job ads, F_p and M_p capture the likelihood that text indicates the presence of such requests based on similarity with job ads with explicit female and male requests as returned by the classification model based on job description. Therefore, it is possible for a job ad with an explicit male request (M job) to have high F_p and vice versa. We regress the share of female (male) applicants to a job on explicit gender requests as well as quartics in F_p and M_p, following the strategy adopted by Kuhn et al. (2020). We include the set of controls in equation (3.3.4), focusing on specifications that either include the full occupation $\times$ state fixed effects or state fixed effects only.[34] Further, we interact the quartics in F_p and M_p with explicit gender requests and use these as additional explanatory variables. Using these regression estimates, we predict the share of female (male) applicants as a function of F_p (M_p) for each type of job (F, M and N jobs). Sub-figure 3.2(a) gives the predicted share of female (male) applicants as F_p (M_p) changes while keeping M_p (F_p) constant at its mean level using a specification that includes occupation $\times$ state fixed effects. Sub-figure 3.2(b) gives the same variation but uses a specification with state fixed effects only.

Sub-figure 3.2(a) shows that as F_p associated with a job increases (or as we switch to jobs with an increasingly female job description *within* the same occupation), the predicted share of female applicants increases for F and N jobs but decreases for M jobs. For low values of F_p, the predicted share of female applicants is similar in F, N and M jobs. As F_p increases there is an initial rapid increase (decrease) in the predicted share of female applicants in F (M) jobs and a more gradual increase in N jobs. Our results show that an explicit gender preference in jobs containing text predictive of female preference (with higher F_p) results in a larger share of female applicants. Similarly, as M_p associated with a job increases (or as we move along jobs with an increasingly male job description *within* the same occupation) the predicted share of male applicants increases for M and N jobs, while it initially decreases but then increases for F jobs. Notably, the difference in predicted share of male applicants

[34] We do not include wage controls in these regressions to ensure we use the full sample of job ads. We also do not include dummies for age, beauty or night shift requirements since the words used to construct these measures are also highly predictive of explicit gender preferences.

across M and N jobs, for a given level of implicit male stereotype, is generally smaller than the difference in the predicted share of female applicants across F and N jobs. The difference in predicted share of male applicants further declines as M_p increases, while the difference in predicted share of female applicants increases and then remains constant as F_p increases.[35] We discuss the implications of these findings later.

Since occupational classification itself is expected to contain gender associations, one would expect that not controlling for occupations would lead to an even higher response of applications to F_p and M_p in a job ad; sub-figure 3.2(b) confirms this. We see that the predicted share of female applicants increases more rapidly as F_p rises (or as we switch from 'mechanical engineer' to 'receptionist' jobs) in F and N jobs. The predicted share of male applicants also increases, more rapidly than in sub-figure 3.2(a), as M_p increases (as we switch from 'telecaller' to 'electrician/IT hardware engineer' jobs) for M and N jobs. As before, the difference in predicted share of male applicants across M and N jobs is generally smaller and further declines as M_p increases than the difference in predicted share of female applicants across F and N jobs as F_p increases.

We also re-construct our measures of F_p and M_p using a Bernoulli NB classifier, as used in Kuhn et al. (2020) (Appendix Figure 3.7). We estimate similar regressions as before to find the predicted share of female (male) applicants using state fixed effects rather than occupation $\times$ state fixed effects since F_p and M_p are now constructed using text in job titles only and these job titles are also used to assign jobs to different occupations. We find that the confidence intervals are larger when F_p and M_p lie in (0,1) for F and M jobs compared to when we used the LR classifier, possibly showing more noise when gender associations are assigned to jobs ads which do not show an extreme association with either gender. We also find that the predicted female (male) applicant shares increase, as F_p (M_p) increases, but only after the threshold of 0.5 is crossed, and particularly for F (M) and N jobs. In comparison, we found a more gradual response of applications as F_p or M_p increased for N jobs when we used the LR classifier for gender association assignment.

Our results bear similarities and differences from those reported by Kuhn et al. (2020). We also find that predicted male shares in job ads with text predictive of a male preference, without an explicit gender preference, are quite similar to predicted male shares when these jobs also have an explicit male preference label. On the other hand, predicted female shares in job ads with text predictive of a female preference, but without an explicit gender preference tend to be lower than the predicted female shares if these jobs were to have an explicit female preference label. This indicates that explicit female requests matter more for female

[35] Predictions at very high and very low values of F_p and M_p are not estimated very precisely since there are few observations at the endpoints.

applicant shares than explicit male requests matter for male applicant shares. However, our findings, unlike those in Kuhn et al. (2020), show that gender associations play a role in increasing female applicant shares even when an explicit female request is made in a job ad. For instance, in F jobs, we find that the share of female applications increases as F_p rises. In fact, the gap in female applicant share, between F and N jobs first increases and then remains constant as F_p increases. This gap falls in China because the female applicant share increases for N jobs but remains relatively constant for F jobs. What could potentially explain the differences between the findings for China and our results for India?

One, the methodology for deriving F_p differs from Kuhn et al. (2020) but even when we use their methodology (Figure 3.7), we continue to find the same differences. Second, while Kuhn et al. (2020) use explicit employer preferences for a female in a separate field, we cull out this explicit preference from the job description. If women do not pay as much attention to job descriptions and overlook a stated female preference, then this could possibly lead to the differences we observe. However, if this were true, then one can reasonably expect men to display similar cognitive inability to perceive an explicit gender preference in a job ad description. However, we find that our results do not differ substantially for men when compared with China. A third potential explanation for the difference in female applicant behaviors is that Indian women, despite seeing an explicit female request in a job ad, either do not feel inclined enough or do not have necessary skills for a job ad predictive of a male preference. For instance, if a female were explicitly requested for a job of 'delivery personnel', many women may not feel comfortable applying or might not have the necessary skills (such as driving) to apply for such a job. Therefore, gendered associations with job ads may constrain Indian women when applying to jobs more than in China. In other words, Indian women seem to conform to gender stereotypes (as given in the job text) even in the presence of an explicit female request in the job ad.

3.3.4 Robustness checks

We further examine the robustness of our results to several modifications as discussed below:

Manual classification of occupations: First, we carry out all estimations using a more dis-aggregate manual occupational classification (with 747 occupation categories) derived from the job title of an ad as described in sub-section 3.2.2, and find that our results are largely robust. We continue to find that explicit gender preferences are less likely in high skill jobs with a higher education requirement (column (I), Appendix Table 3.12). Our results on male preferences when using the alternative occupation classification are very similar in sign and significance, with some differences in the size of the coefficients (column (IV), Appendix

Table 3.12). In wage regressions that use the sample of N jobs we find that the decrease in advertised wage associated with an increase in F_p continues to be far higher than the decrease associated with the same increase in M_p (column (I), Appendix Table 3.13). We also find a similar pattern of effects when we examine either the total number of applications or the share of female applicants as our dependent variables of interest upon using the alternative occupation classification (columns (I) and (IV), 3.14).

Firm fixed effects: We also carry out estimations with firm $\times$ state fixed effects rather than occupation $\times$ state fixed effects, and our most restrictive specification uses firm $\times$ occupation $\times$ state fixed effects.[36] A caveat is that we observe a few firms posting a large number of jobs across different sectors. Since we only observe a company ID and not firm names, we cannot rule out that some firms are actually HR recruiters. Nevertheless, we continue to find that our results are largely robust. We still find that higher education requirements result in a higher probability that a job ad has an explicit gender preference (columns (II) and (III), Appendix Table 3.12) and that higher F_p has a larger negative effect on the advertised log wage than higher M_p among N jobs (columns (II) and (III), Appendix Table 3.13). We also continue to find that an explicit female preference leads to a larger reduction in the number of applications than an explicit male preference while there is a substantial shift in the gender mix of the applicant pool in favor of women if there is an explicit female requirement in a job ad (columns (II)-(III) and (V)-(VI), 3.14).

Non-linear models: We also carried out estimations on binary and ordered outcomes using non-linear specifications with coarser fixed effects, and continue to find that our results are largely robust. We find similar marginal effects in probit regressions which use an explicit gender requirement as the outcome variable of interest (column (I), Appendix Table 3.15). These continue to support the negative skill targeting relationship. When looking at preferences for men we estimate an ordered probit which allows us to examine the marginal effect of different variables on either the male or female preference outcome; we continue to find similar results as before (columns (II) and (III), Appendix Table 3.15).

[36] In Appendix Tables 3.12-3.14 we report the number of observations as job ads for which the gender requirement or dependent variable varies within firms in a given state or within a firm and occupation in a given state (depending on the fixed effects used) since we are effectively only using these job ads in our estimations.

3.4 Deconstructing gendered word representations

Pioneering work by Akerlof and Kranton (2000) shows how association with a group identity matters for economic outcomes. Thereafter, several studies evaluate gender differences along various dimensions such as risk, overconfidence, competitive behavior, undertaking negotiations, and ambiguity aversion and show how these relate to economic returns in the labor market.[37] What remains under explored are the mechanisms through which such differences may be generated. In a recent study, Bordalo et al. (2019) use a cooperative game setup and find that gender associations or stereotypes contribute to gender gaps in self-confidence and consequently in behavior. The results in section 3.3.2 extend this literature by showing that *femaleness* or *maleness* associated with a job indeed matter for the gender wage gap and the gender mix of applicants for a job.

A natural question that follows is which words contribute to such gender associations. Previous studies in psychology show that wording of job ads might affect students' inclination to apply for a job. Born and Taris (2010) find that females respond more to feminine characteristics than men respond to masculine characteristics since men can be overconfident.[38] In the existing literature, the characteristics that attract women and men are drawn from small surveys. For instance, Taris and Bok (1998) compile 20 characteristics based on 512 job ads judged by 40 students as being typically male or female while Gaucher et al. (2011) use implicit association tests.[39] In a recent paper, Arceo-Gómez et al. (2020) use job ads to classify the most common words based on their relative frequency of occurrence in female or male requesting ads. We, on the other hand, uncover words associated with gender requests in job ads derived from the classification decisions of a state-of-the-art machine learning (ML) model itself. We then use the gendered word representation aggregated by categories such as hard skills, soft skills, personality and flexibility and look at which of these matter for female applicant share and advertised wages. This analysis enables us to understand which gendered job attributes matter for wages and whether variation in female applicant share along these attributes reflects the different willingness to pay for these by gender.

[37] A literature review of these studies is beyond the scope of this chapter.

[38] The study used the characteristics "solid business sense" and "decisiveness" (both masculine), and "communication skills" and "creativity" (both feminine) to describe desired candidate profile to 78 applicants.

[39] Abele and Wojciszke (2014) further divide these associations with words into largely communal and agentic types.

3.4.1 Methodology to construct word scores

We use the Layer-wise Relevance Propagation (LRP) technique proposed by Bach et al. (2015) to explain which words in job texts correspond to explicit gender preferences of employers. LRP overcomes the *black box* nature of complex non-linear ML models such as deep neural networks (DNN). It estimates the extent to which each input contributes towards making a specific classification decision. It uses backward propagation to redistribute the output "classification score" to the input space which, in our case, is a 300-dimensional vector representation of each word in a job ad. In doing so, it satisfies a layer-wise conservation principle which means that the sum of relevance values for each layer (including the input layer) are equal to the classification score. The relevance R_w of each word w is then computed from the input vectors by simply summing the input relevance scores $R_{w,d}$ over all the dimensions d of the vector. LRP decomposes the actual value of the classification score $f(x)$ returned by the ML model f using input x. Hence, it can identify words which have positive as well as negative relevance for a class. Another technique which is used to explain predictions of ML models is sensitivity analysis. It assigns relevance values to inputs by making a small perturbation in the input x and then measuring how the classification score f changes. This means that LRP directly answers what words make a particular job ad male targeted, whereas sensitivity analysis answers what *change in words* will make a job ad more or less male targeted. Therefore, LRP is more suitable for explaining the predictions of a ML model.

The existing literature uses LRP in many applications ranging from computer vision to biomedicine and its suitability over sensitivity analysis is now well established. We discuss many of these in the Technical Appendix Section 3.A.5. For us, the application of LRP to text data is more relevant. In the Natural Language Processing (NLP) literature, Arras et al. (2017b) evaluate the performance of DNN for sentiment analysis, and Ding et al. (2017) for neural machine translation. In particular, Arras et al. (2017a) use LRP to explain the predictions of a bag-of-words (BOW) Support Vector Machine (SVM) model and a Convolutional Neural Network (CNN) model for text classification.[40] They find that even though both the models performed comparably in terms of classification accuracy, the explanations from the CNN model were more human interpretable. There are a couple of reasons for this. First, the CNN model can take into account semantic similarity between words encoded in the word vector representations. Second, as opposed to a BOW model which only takes into account frequency of occurrence of word n-grams in each class, the CNN model takes into account the context in which words occur in a document. Consequently, the scores assigned

[40] CNN is a class of DNN that has achieved impressive results on various tasks related to computer vision and NLP.

to a specific word can be different across documents.

We introduce LRP to the domain of economics and demonstrate how labeled text data based on explicit gender requests in job ads can be used to explain decisions of the underlying ML model and systematically extract words that reflect gender associations. Explainability in itself might be desirable to assess the validity and the generalizability of the model, and hence to gain trust in its predictions.[41] Moreover, these words can potentially be used by researchers to demonstrate gender associations of ads in countries where no explicit gender requests are made. For instance, Burn et al. (2019) use word vectors to calculate cosine similarity of words from the existing literature in the field of industrial psychology with phrases in jobs ads to detect bias against older workers in the U.S.

We now elucidate the steps followed to implement the method to our text data. First, we train a CNN model (LeCun et al., 1990). Here, we restrict our sample to 15,272 job ads for which the employer mentions an explicit female or male preference.[42] We follow the same preprocessing steps as in Section 3.2.4. Then we tokenize the documents so that every word is encoded as a vector and the documents are represented as a matrix. To reduce training time, we truncate each sequence to a maximum length of 125 words ($\approx 90^{th}$ percentile document length) and zero pad, i.e. add a sequence of zeros at the end of shorter sequences until their length is equal to this maximum sequence length. We then perform stratified five-folds cross-validation. For each of the five folds, the model is trained on four folds or 80% of the data and its performance is assessed using the remaining 20% of the data as the test set. The model is able to correctly classify 89.3% and 85.7% jobs on the test set with explicit preference for females and males respectively. We discuss the details of the CNN model in the Technical Appendix Section 3.A.6.

In the second step, we map the classification scores into the input space using LRP over the test set documents and assign a relevance score to every word in each job ad text which indicates the importance of that word to each class (explicit gender request in our context).[43] Figure 3.8 shows a heat map visualization of words in distinctive job ads with explicit female and male preferences. Words indicating female preference are highlighted in red, while those predictive of male preferences are highlighted in blue. The color intensity reflects the strength of the attached gender association, with darker shades showing a higher

[41] A model can spuriously achieve a high accuracy on the test data set without learning anything meaningful. For instance, Lapuschkin et al. (2019) show that a Fisher Vectors based model which achieved high accuracy on an image classification task "misused" source tags that distinctively occurred in lower left corners of horse images instead of classifying on the basis of an actual horse image.

[42] Unlike our previous analyses, we do not place any restriction on these set of job ads and make use of all job ads out of 196,821 provided to us which have an explicit gender preference.

[43] We use the implementation of Ancona et al. (2017) available at https://github.com/marcoancona/DeepExplain.

strength. For example, the figure shows that for the "software trainee" job, employers prefer women when they are looking for *English* language skills, *grammar*, and *personality*. However, within the same job the knowledge of accounting software *Tally* and *GST* (goods and services tax) is associated with men. On the other hand, employers request males when they look for knowledge related to technology and specify working in *rotational shift*. However, *verbal* communication skills are still associated with women. In the next section we explain how the relevance scores are used to arrive at overall gender associations for every word.

3.4.2 Gender requests: which words matter?

We have seen in sub-section 3.3.3 that gender associations in job ad text matter. We now use 3,927 words that occur at least ten times in the 15,272 jobs ads with explicit male or female preferences and classify them manually into four stereotype categories: hard skills/other skills (361 words), soft skills (83), personality/appearance (109), and flexibility/benefits (56) associated with a job. We classify words under the category hard skills/skills if they are related to knowledge about a particular software, hardware or specific skills such as driving or typing. Soft skills include words that refer to communication or interpersonal skills. The third category personality/appearance refers to other personal attributes of a prospective candidate that a job requires. Lastly, job flexibility/benefits captures words related to job timings, travel requirements and other benefits such as retirement or insurance. The remaining words cannot be classified into any of these categories (most words are generic or reflect occupation or other job and candidate specific attributes) or fall under multiple categories; we classify these words as "others". Deming and Kahn (2018) also classify keywords in job ads into 10 skill categories to look at changing demand for skills over time in the U.S.[44] We use a more aggregate skill categorization into hard skills and soft skills, since we aim to look at the variation in gender associations within each of these broad categories and how that may affect the gender composition of job applicants. A highly disaggregate classification may give little variation in constructed scores (discussed in Section 3.4.3). For instance, disaggregating hard skills into financial and computer skills is not appropriate for our purpose since the two types of skills can possibly be attached to a different gender.

To obtain top words under each category, we take the mean of the relevance scores for each word over all its occurrences in every job ad. A more positive (negative) mean relevance score for a word reflects greater relevance of that word towards explaining explicit female (male) requests in job ads. We sort the words on the basis of their mean relevance scores

[44] The 10 categories are cognitive, social, character, writing, customer service, project management, people management, financial, computer and software skills.

and list the top 20 words that are associated with requests for females and males within each category in Table 3.5. The results are striking and show that many of the words that are typically associated with male and female job roles indeed show up on the list.

Within hard skills (columns (I) and (II), Panel A), skills associated with a beautician (*nailcare, pedicure, manicure, waxing, facial, makeup, threading*), human resource management (screening candidates), accounting tasks and software (*debtor, loss statements, tally*), and knowledge of tools used for communication and word processing (*email, word, ms, zoho*) appear for women. For men, skills related to jobs in IT/hardware (*rcm, mysql, oop, pcb, printer, api*) and finance (*demat, psa*) tend to dominate along with skills such as driving and repairing. Next we look at soft skills (columns (III) and (IV), Panel A) and again find a stark distinction within required soft skills across gender. While jobs requesting women focus on communication skills, interpersonal skills and coordination to maintain customer relations, those requesting men mention skills requiring assertiveness such as pitching to a client, negotiating, dealing, supervising and giving feedback. Interestingly, though language skills such as English, Hindi, and Urdu come under the female category, knowledge of Marathi comes under male category. This could be due to variations in gender stereotypes across different regions of the country.

The gender contrast is also evident in different personality traits across jobs that request women and men (columns (I) and (II), Panel B). Jobs requesting women require the candidate to be bold, open-minded, mature, have a pleasing and vibrant personality, fair complexion, no scars/tattoos on skin and other characteristics such as punctuality and sincerity. On the other hand, personality traits such as the ability to handle pressure and tight deadlines, being vigilant, sharp, creative, ethical/honest, hardworking, determined, curious, dedicated, go beyond the 'mile' and engage in teamwork are used when requesting a male candidate to apply for a position. Lastly, words indicating job flexibility and other benefits such as availability of cab for transportation (and pick/drop facility), dinner, accommodation, convenient timings with holidays and option to work from home (*skype*) are usually associated with jobs requesting a female (column (III), Panel B). On the other hand, night/evening/rotational shifts, possible relocation, travel (petrol/fuel, conveyance, Travel/Dearness Allowance (*ta, da*), other benefits such as gratuity, retirement benefits (*fpf* and *pf*) and insurance (*esic, esi*) are associated with male requests (column (IV), Panel B). Overall, we find that distinct skills and personality traits are associated with jobs that request men and women. Next, we use the generated relevance scores in the above four categories along with words in the "other" category to see how the gender mix of the applicant pool is affected by these measures.

3.4.3 Which gendered word representations matter for female applicant share?

We use the word classification in Section 3.4.2 to construct a score for each job ad under each gender-word classification category. We add up *absolute values* of relevance scores for all the words in a category in a job ad, separately for words having a positive relevance score (female) and for those having a negative relevance score (male). For instance, for the category hard skills/skills two separate variables are generated—Female (skills) and Male (skills). Female (skills) is the sum of relevance scores for all words classified under the category of hard skills/skills that get a positive relevance score in a job ad. Similarly, Male (skills) is the sum of the absolute value of relevance scores for all words classified under the category of hard skills/skills that get a negative relevance score in a job ad. If a job ad does not have any word in a given gender-category then it gets a zero score for it. This procedure is used to construct ten gender-category variables, two for each of the five categories (including "others"). We report the summary statistics for the gender-category variables in Table 3.16 separately for jobs with male, female and no gender preference.

The obtained score in each gender-category is then standardized for ease of interpretation. We regress the proportion of female applicants to a job on the standardized scores for the ten variables, along with controls for education and experience requirements mentioned in the job ad and month-year fixed effects (similar to equation (3.3.4), but using the standardized scores for the ten variables rather than explicit gender request variables). The specifications include either location fixed effects only or occupation $\times$ location fixed effects. As before, we weight the regressions by total applications and cluster the standard errors at occupation and state level. We estimate regressions separately for each type of job (N, F and M type) because the effect of word representations can be different across jobs that request a particular gender versus those that do not. The results reported in Table 3.6 are discussed below.

The results for N jobs show that an increase of one standard deviation in 'Female (skills)' increases the fraction of female applicants (column (III)). Also, increased word representations for 'Female (soft skills)' and 'Male (skills)' increases the female applicant share while increased word representations for 'Male (flexibility)' and 'Male (others)' reduces the female applicant share. Once occupation $\times$ location fixed effects are included, the effect of female soft skills on the female applicant share becomes insignificant while the other variables still remain significant (column (IV)). On the other hand, for jobs that explicitly request a female (F jobs), there is no additional effect of female word representations in any category when occupation $\times$ location fixed effects are included (column (II)); though the effect of female skills is comparable to N jobs, it is imprecisely estimated. Increased male word representations in the categories of skills, job flexibility and others reduce the female applicant share, thus reducing compliance with the employer's gender requirement. An increase in male word representation in M jobs for personality and others decreases the female applicant share or increases compliance with the explicit gender request while including other female attributes increases female applicant share or reduces compliance (column (VI)).

These results show that for N jobs, both female and male oriented hard skills matter the most in increasing the share of women applicants while male oriented job flexibility and other male attributes decrease the proportion of female applicants to a job, even when we use *within* occupation variation only. Compliance with gender requests also falls if words related to the opposite gender are present and increases if words related to the requested gender are included in a job ad. Overall, we find that the wording of a job ad matters for the gender mix of the applicant pool, especially for words related to hard skills and job flexibility even when using within occupation variation only. However, gendered word representations related to soft skills do not seem to matter once occupation controls are included.

3.4.4 Which gendered word representations matter for advertised wages?

The above analyses shows that gendered word representations matter for application rates. We have also seen in Table 3.3 that implicit femaleness associated with a job reduces advertised wage and this effect was larger for N jobs. Implicit femaleness can increase due to words related to females occurring more or due to words related to males occurring less in in such job advertisements. As a next step, we check which gendered words matter to the advertised wages. We regress the log of advertised wage in a job ad on the standardized scores for the ten variables, along with controls for education and experience requirements mentioned in the job ad and month-year fixed effects and cluster the standard errors at occupation and state level. The results are reported in Table 3.7.

The results for N jobs show that an increase of one standard deviation in 'Female (skills)' and 'Female (soft skills)' decreases the wage significantly by 1.3% and 0.6% respectively (column (III)). While, a one standard deviation increase in 'Male (skills)' and 'Male (flexibility)' increases the advertised wage by 1.2% and 2.4% respectively (column (III), Table 3.7). We then add controls for occupation fixed effects in every state and report the results in column (IV) of Table 3.7 to check if these results are driven by variation in occurrence across occupations. The negative effect of 'Female (skills)' on advertised wage and the positive effect of 'Male (skills)' and 'Male (flexibility)' persists, albeit the magnitudes decline. In addition, an increase in words related to 'Male (soft skills)' also have a positive effect on advertised wages.

Next, we discuss the results for F and M jobs, controlling for occupation and state fixed effects. The results for F jobs show that an increase in one standard deviation for words related to 'Female (skills)' and 'Female (soft skills)' decreases the wage significantly while words related to 'Female (personality)' are associated with higher advertised wages when an explicit request is made even within the same occupation (Column II). However, if the job ad contains words that indicate 'Male (flexibility)' it is associated with a higher advertised wage. Thus, if employers want a female for a position, they are willing to pay a higher wage premium for the likely case that the job requires longer working hours, travel, relocation or night shifts. Lastly, the results for M jobs show that none of the words related to men matter to advertised wages within a given occupation (column (VI)). Only the presence of words related to 'Female (skills)' decrease the advertised wage within an occupation for M jobs. Even without occupation fixed effects, only 'Male (flexibility)' increased the advertised wages (column (V)). These results show that both 'Female (skills)' and 'Male (flexibility)' matter to advertised wages. While words related to 'Female (skills)' decreases the advertised wages, those related to 'Male (flexibility)' increase the wages. This results align with female skills penalized in the market and also with the trade-off between flexibility and wages. Combined with the results in Table 3.6, that show that female applicant share increases as words related to 'Male (flexibility)' decrease, these show that women are willing to pay for flexible hours, lower travel requirements and other factors associated with higher male flexibility.

3.4.5 Discussion

These results have broader implications for the literature on gender wage gaps and labor market structure. In general, the gender wage gap reflects human capital differences between men's and women's productivity as well as differential treatment of men and women in the labor market. In most developed and developing countries the proportion of gender wage gap due to differences in human capital investments has fallen over time (Kunze, 2018; Deshpande et al., 2018). Commensurately, the unexplained or the residual wage gap has increased over time. This residual wage gap can be due to taste or statistical discrimination, occupational segregation, degree to which women negotiate, compete, accumulation of human capital. Goldin (2014) argues that variation in gender wage gap with age and children profile of women does not board well with the innate differences between men and women. She finds that in the U.S., within-occupation wage differentials account for a larger proportion of the gender wage gap than between-occupation wage differentials. Thus, a focus on attributes of jobs within occupations then becomes important. She argues that attributes related to 'job flexibility' can matter to within occupation differences.[45]

Goldin (2014) uses survey data on employed individuals from the U.S. and demonstrates this by showing that wages vary non-linearly with hours of work for employed individuals. A flexible schedule is hence less likely to be rewarded in the labor market. Recent studies by Mas and Pallais (2017) and Bustelo et al. (2020) using discrete choice experiments show that women workers are more likely to pay for flexible schedules and working from home in the U.S., Colombian and Chinese labor markets, albeit across different category of workers. The willingness to pay ranges from 8% to 20% depending on the nature of flexibility offered and the context. These estimates are higher for women and further higher for women with children (Bustelo et al., 2020). We use data from real job ads posted on an online platform and applications to these job ads to test the applicant behavior as job attributes change. Our findings show that job ads which offer less flexibility post higher advertised wages and the proportion of female applicants were higher for such jobs. These results are striking since they hold even within the same occupation and location and lend support to the hypothesis that the residual gender wage gap within an occupation can reflect the earnings that women are willing to pay for job attributes, most notably for greater job flexibility. Thus, the residual wage gap may not just reflect employer discrimination but sorting of workers across jobs based on the job attributes and worker preferences. Albeit, these preferences can be shaped by societal roles with women being the primary caregivers.

3.5 Conclusion

We examine explicit gender requests in job ads posted on an online job portal in India. We find that jobs with lower skill requirements are more likely to place a request and that ads requesting women offer lower wages. Applicant responses show high, but imperfect compliance, to these requests with women (men) applying proportionately more to jobs with an explicit female (male) request. We use detailed occupation level controls in our analyses to ensure our estimates are not capturing explicit requests being made in occupations regarded as traditionally male or female. Further, we use explicit gender preferences to derive implicit gender associations, and find that a job ad containing text predictive of an explicit female preference, i.e. *femaleness*, offers lower wages. On the applications side, an increase in *femaleness* associated with a job ad, leads to a higher female applicant share, even if an explicit female request is made. Lastly, uncovering the job ad wording, we find that ads

[45] Workplace flexibility can incorporate many features of the workplace like total hours, precise timings, flexibility to schedule one's day, possibility of work from home, extensive client meetings, travel etc.

requesting women and men differ along pre-existing gender stereotypes regarding skills, job flexibility and personality traits. These words have implications for compliance with gender requests in a job ad and affect applicant behavior for ads that do not make such requests.

Our results bear significant relevance, given the low female labor force participation rates in India, and in the absence of effective legal bans on gender requests in job ads (unlike the US or China). Our results indicate that placing restrictions on gender targeting in job ads can increase the share of job applications by women. We find that implicit gender associations matter for wage gaps and application rates, and a concerted policy to address these is important (Dhar et al., 2018). Lastly, our results using data from primarily entry level job ads are striking. We show that stereotypes matter at a stage when young people are entering the labor market, and can have important cumulative consequences for future labor market returns.

Figures and Tables

Figure 3.1: Word clouds of job titles

(a) Female preference (F jobs)

(b) Male preference (M jobs)

(c) No gender preference (N jobs)

Notes: The word clouds are constructed based on words contained in job titles of ads displaying an explicit female preference, an explicit male preference and no explicit gender preference. *Source:* Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890.

Figure 3.2: Predicted share of female (male) applicants with implicit femaleness (maleness)

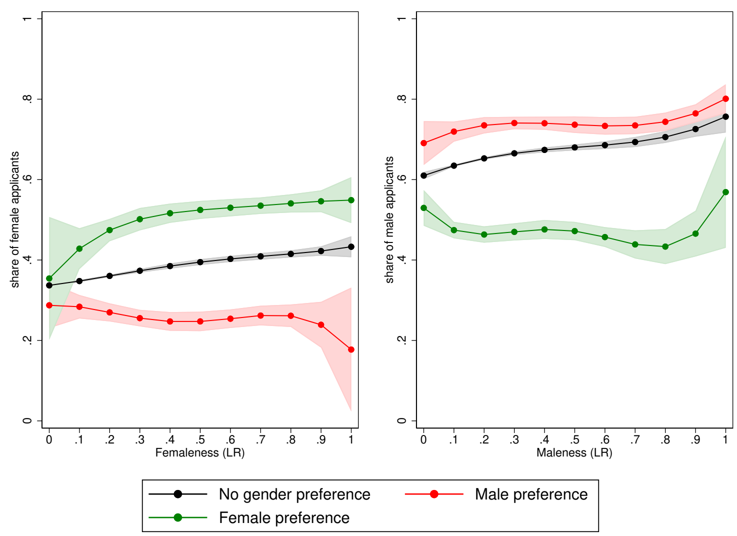

(a) occupation × state fixed effects

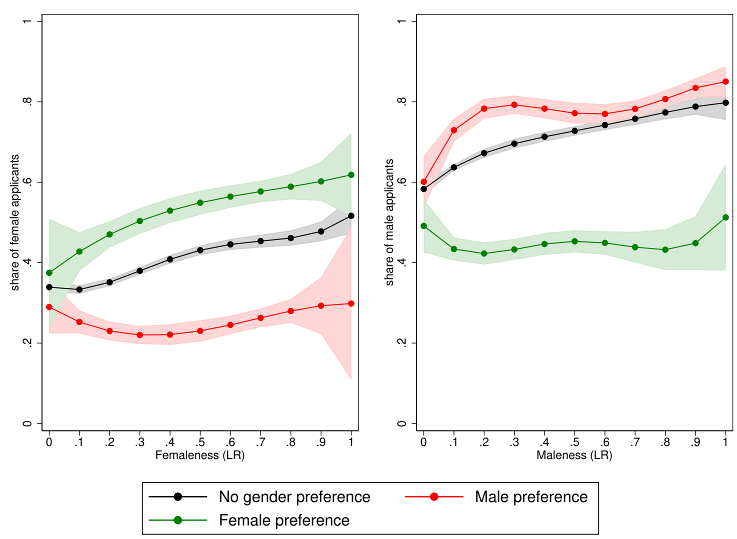

(b) state fixed effects

Notes: Shaded areas give the 95% confidence intervals around predicted values. The measure of implicit femaleness (maleness) is constructed using a logistic regression classifier as described in sub-section 3.2.4. Predictions are based on regressing the share of female (male) applicants on explicit gender preferences, quartics in implicit femaleness (maleness), their interactions and the set of controls specified in equation (3.3.4) and (month, year) fixed effects. Predictions used to construct the figures in (a) also include occupation × state fixed effects while those in (b) include state fixed effects only. These regressions are weighted by the total number of female and male applications, with standard errors clustered by occupation and state.

Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890.

Table 3.1: Explicit gender preference

	(I)	(II)	(III)	(IV)
Education requirements:				
Senior secondary	–0.0680***	–0.0363***	–0.0363***	–0.0346***
	(0.0106)	(0.0073)	(0.0072)	(0.0075)
Diploma	–0.0880***	–0.0456***	–0.0446***	–0.0428***
	(0.0134)	(0.0090)	(0.0090)	(0.0093)
Graduate degree, STEM	–0.1088***	–0.0548***	–0.0532***	–0.0441***
	(0.0136)	(0.0086)	(0.0085)	(0.0090)
Graduate degree, non-STEM	–0.0904***	–0.0453***	–0.0435***	–0.0377***
	(0.0132)	(0.0089)	(0.0088)	(0.0091)
Postgraduate degree, STEM	–0.0733***	–0.0668***	–0.0679***	–0.0638***
	(0.0176)	(0.0142)	(0.0142)	(0.0198)
Postgraduate degree, non-STEM	–0.0609***	–0.0437***	–0.0445***	–0.0193
	(0.0159)	(0.0131)	(0.0131)	(0.0182)
Experience requirements:				
$1 - 2$ years	0.0205***	0.0187***	0.0198***	0.0279***
	(0.0041)	(0.0032)	(0.0031)	(0.0035)
>2 years	–0.0148***	–0.0041	–0.0035	0.0107***
	(0.0028)	(0.0027)	(0.0027)	(0.0034)
Other job requirements:				
Age requirement present			0.0499***	0.0565***
			(0.0077)	(0.0090)
Beauty requirement present			0.0318***	0.0321***
			(0.0092)	(0.0097)
Working night shifts specified			0.0223**	0.0234**
			(0.0093)	(0.0097)
Advertised wage:				
ln(wage)				–0.0965
				(0.0583)
$\ln(\text{wage})^2$				0.0028
				(0.0023)
Fixed Effects	none	occ × state	occ × state	occ × state
N	156279	156279	156279	136729

Notes: The dependent variable takes the value one if a job ad shows a male or female preference and zero otherwise. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890. Each column reports the effective number of observations used for estimations after incorporating occ × state fixed effects. This excludes job ads for which there is no variation in the dependent variable within an occ × state cell.

Table 3.2: Explicit male preference

	(I)	(II)	(III)	(IV)
Education requirements:				
Senior secondary	–0.072***	–0.045***	–0.044***	–0.047***
	(0.012)	(0.008)	(0.008)	(0.008)
Diploma	–0.064***	–0.051***	–0.050***	–0.054***
	(0.016)	(0.010)	(0.010)	(0.010)
Graduate degree, STEM	–0.057***	–0.045***	–0.044***	–0.046***
	(0.016)	(0.010)	(0.009)	(0.010)
Graduate degree, non-STEM	–0.083***	–0.055***	–0.053***	–0.056***
	(0.015)	(0.010)	(0.010)	(0.010)
Postgraduate degree, STEM	–0.131***	–0.067***	–0.067***	–0.062***
	(0.019)	(0.014)	(0.014)	(0.019)
Postgraduate degree, non-STEM	–0.124***	–0.072***	–0.072***	–0.052***
	(0.019)	(0.014)	(0.014)	(0.019)
Experience requirements:				
$1-2$ years	–0.003	0.002	0.003	0.003
	(0.005)	(0.004)	(0.004)	(0.004)
>2 years	0.008***	0.009***	0.009***	0.004
	(0.003)	(0.003)	(0.003)	(0.003)
Other job requirements:				
Age requirement present			0.023***	0.037***
			(0.006)	(0.007)
Beauty requirement present			–0.051***	–0.055***
			(0.007)	(0.008)
Working night shifts specified			0.077***	0.078***
			(0.011)	(0.011)
Advertised wage:				
ln(wage)				0.259***
				(0.055)
ln(wage)^2				–0.010***
				(0.002)
Fixed Effects	none	occ × state	occ × state	occ × state
N	156279	156279	156279	136729

Notes: The dependent variable takes the value minus one if a job ad shows a female preference, zero if it does not show any gender preference and one if it shows a male preference. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890. Each column reports the effective number of observations used for estimations after incorporating occ × state fixed effects. This excludes job ads for which there is no variation in the dependent variable within an occ × state cell.

Table 3.3: Implicit *femaleness* (*maleness*) and the advertised wage

Sample:	*F* jobs		*N* jobs		*M* jobs	
	(I)	(II)	(III)	(IV)	(V)	(VI)
Femaleness	–0.162***	–0.192***	–0.393***	–0.286***	–0.340***	–0.189***
	(0.049)	(0.043)	(0.018)	(0.019)	(0.043)	(0.049)
Maleness	0.022	0.044	0.054***	0.013	–0.019	–0.030
	(0.041)	(0.035)	(0.019)	(0.016)	(0.032)	(0.036)
Education requirements:						
Senior secondary	0.053*	0.033	0.073***	0.045***	0.125***	0.029
	(0.026)	(0.017)	(0.008)	(0.007)	(0.031)	(0.023)
Diploma	0.112***	0.083***	0.017	0.027***	0.179***	0.109***
	(0.032)	(0.026)	(0.010)	(0.008)	(0.042)	(0.034)
Graduate degree, STEM	0.171***	0.114	0.207***	0.188***	0.206***	0.186***
	(0.062)	(0.069)	(0.013)	(0.012)	(0.062)	(0.064)
Graduate degree, non-STEM	0.099***	0.081***	0.077***	0.064***	0.209***	0.099***
	(0.022)	(0.020)	(0.010)	(0.007)	(0.032)	(0.025)
Postgraduate degree, STEM	1.312*	0.423	0.509***	0.392***	0.967*	1.110*
	(0.612)	(0.286)	(0.059)	(0.048)	(0.481)	(0.526)
Postgraduate degree, non-STEM	–0.057	0.112	0.278***	0.270***	0.043	–0.135
	(0.121)	(0.089)	(0.037)	(0.037)	(0.095)	(0.089)
Experience requirements:						
$1-2$ years	0.082***	0.105***	0.056***	0.072***	0.092***	0.089***
	(0.022)	(0.016)	(0.008)	(0.007)	(0.022)	(0.021)
>2 years	0.221***	0.227***	0.315***	0.308***	0.280***	0.265***
	(0.026)	(0.022)	(0.012)	(0.011)	(0.031)	(0.032)
Fixed Effects	state	occ × state	state	occ × state	state	occ × state
N	5963	5963	124313	124313	5208	5208

Notes: The dependent variable is the log wage advertised in a job ad. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890. Each column reports the effective number of observations used for estimations after incorporating occ × state fixed effects and dropping job ads which do not advertise a wage. This excludes job ads for which there is no variation in the dependent variable within an occ × state cell.

Table 3.4: Applications

Dependent variable:	total applications			share of female applications		
	(I)	(II)	(III)	(IV)	(V)	(VI)
Female preference	–8.938***	–8.496***	–5.791***	0.157***	0.157***	0.161***
	(0.960)	(0.942)	(0.860)	(0.007)	(0.007)	(0.007)
Male preference	–2.262	–2.566	–4.570*	–0.090***	–0.089***	–0.088***
	(3.104)	(3.100)	(2.283)	(0.007)	(0.007)	(0.008)
Education requirements:						
Senior secondary	1.753	1.803	1.347	0.028***	0.028***	0.029***
	(1.038)	(1.039)	(0.970)	(0.003)	(0.003)	(0.003)
Diploma	4.398**	4.357**	3.501	0.014**	0.014**	0.015**
	(1.906)	(1.910)	(1.919)	(0.006)	(0.006)	(0.006)
Graduate degree, STEM	59.613***	59.634***	57.606***	0.042***	0.041***	0.038***
	(10.142)	(10.161)	(9.889)	(0.005)	(0.005)	(0.006)
Graduate degree, non-STEM	11.430***	11.412***	9.036***	0.060***	0.060***	0.062***
	(2.039)	(2.050)	(1.796)	(0.004)	(0.004)	(0.004)
Postgraduate degree, STEM	–0.581	–0.179	–1.791	0.107***	0.106***	0.113***
	(8.509)	(8.537)	(11.634)	(0.013)	(0.013)	(0.015)
Postgraduate degree, non-STEM	0.950	1.240	–10.284***	0.090***	0.090***	0.096***
	(2.682)	(2.702)	(3.810)	(0.016)	(0.016)	(0.020)
Experience requirements:						
$1-2$ years	–24.135***	–24.267***	–17.852***	–0.019***	–0.019***	–0.018***
	(2.943)	(2.950)	(2.006)	(0.003)	(0.003)	(0.003)
>2 years	–45.458***	–45.503***	–36.164***	–0.042***	–0.042***	–0.039***
	(4.789)	(4.790)	(3.449)	(0.003)	(0.003)	(0.003)
Other job requirements:						
Age requirement present		–4.368***	–3.027***		–0.012**	–0.010*
		(1.042)	(1.096)		(0.005)	(0.005)
Beauty requirement present		–2.953**	–3.259***		–0.001	–0.000
		(1.276)	(1.251)		(0.004)	(0.003)
Working night shifts specified		17.530**	13.016		–0.031***	–0.028***
		(7.806)	(7.271)		(0.005)	(0.005)
Advertised wage:						
ln(wage)			–38.566			–0.066***
			(49.398)			(0.020)
$\ln(\text{wage})^2$			2.376			0.003***
			(2.088)			(0.001)
Fixed Effects	occ × state	occ × state	occ × state	occ × state	occ × state	occ × state
N	156279	156279	136729	156279	156279	136729

Notes: Dependent variable in columns (I)-(III) is number of applicants to a job ad and in columns (IV)-(VI) is fraction of female applicants. Omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category in experience requirement categories is none to less than a year of experience. All regressions include (month, year) of job posting FE and regressions in columns (IV)-(VI) weighted by total number of applications made to a job ad. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Population of all job ads and applicants on the portal, subject to restrictions described in section 3.2. Final number of job ads is 157890. Each column reports the effective observations used for estimations after incorporating occ × state fixed effects. This excludes job ads having no variation in the dependent variable within an occ × state cell.

Table 3.6: Gender word representations and female share of applicants

Sample:	F Jobs		N Jobs		M Jobs	
	(I)	(II)	(III)	(IV)	(V)	(VI)
Female (skills)	0.006	0.006	0.007***	0.003***	0.014*	0.011
	(0.005)	(0.005)	(0.001)	(0.001)	(0.007)	(0.006)
Female (soft skills)	–0.001	0.001	0.004**	–0.000	0.004**	0.001
	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)
Female (personality)	0.004	–0.001	0.001	0.001	0.002	0.000
	(0.003)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Female (flexibility)	0.003	–0.005	0.001	–0.000	0.009	–0.004
	(0.008)	(0.007)	(0.001)	(0.001)	(0.008)	(0.005)
Female (others)	0.008	–0.001	0.001	0.003	0.046***	0.028***
	(0.006)	(0.005)	(0.002)	(0.002)	(0.007)	(0.007)
Male (skills)	–0.033***	–0.020***	0.007***	0.002**	0.011***	0.007
	(0.008)	(0.007)	(0.002)	(0.001)	(0.003)	(0.004)
Male (soft skills)	–0.017***	–0.003	0.000	–0.000	0.005	0.008
	(0.005)	(0.005)	(0.001)	(0.001)	(0.005)	(0.005)
Male (personality)	0.002	0.003	0.002	0.001	–0.016***	–0.013***
	(0.006)	(0.005)	(0.001)	(0.001)	(0.005)	(0.004)
Male (flexibility)	–0.021***	–0.008***	–0.007***	–0.004***	0.002	–0.001
	(0.003)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Male (others)	–0.071***	–0.045***	–0.005**	–0.004*	–0.035***	–0.017***
	(0.009)	(0.006)	(0.002)	(0.002)	(0.004)	(0.005)
Fixed Effects	state	occ × state	state	occ × state	state	occ × state
N	6317	6317	143190	143190	5386	5386

Notes: The dependent variable is the fraction of female applicants in a job ad. All regressions control for a set of education and experience requirement categories given in a job ad. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects and are weighted by the the total number of applications made to a job ad. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890. Each column reports the effective number of observations used for estimations after incorporating occ × state fixed effects for the respective subsamples. This excludes job ads for which there is no variation in the dependent variable within an occ × state cell.

Table 3.5: Gendered word representations

(I)	(II)	(III)	(IV)
Panel A			
Hard skills/Skills		**Soft skills**	
Female	Male	Female	Male
nailcare	rcm	communicator	pitching
pedicure	mysql	accent	tackle
manicure	driving	fluency	verbally
waxing	corel	spoken	supervise
facial	repair	counsel	listening
stitching	oop	etiquette	probe
makeup	demat	english	pitch
eyebrow	psa	color	dealing
bioinformatic	scanner	speech	negotiation
debtor	machine	communicative	fluently
zoho	pcb	mti	negotiate
emailing	schematic	relations	spelling
word	troubleshooting	hindi	motivate
threading	mvc	editing	convince
coral	machinery	proofreading	write
ms	regulation	urdu	influence
loss	printer	ordination	edit
screen	install	articulate	marathi
shortcut	api	interpersonal	feedback
tally	metal	verbal	ar
Panel B			
Personality/Appearance		**Job Flexibility/Benefits**	
Female	Male	Female	Male
bold	pressure	mediclaim	rotationally
mature	vigilant	cab	fpf
minded (open)	multitask	epf	petrol
frank	chest	timing	esic
personality	bent	bond	fuel
pleasant	energetic	holidays	night
pleasing	ethical	pm	conveyance
scars (no)	sharp	holiday	esi
punctual	cm	convenient	pf
tattoos	tight	drop	da
tone	honesty	attractive	transport
vibrant	hardworking	coupon	shifts
smile	honest	reimburse	room
sincere	creatively	weekday	shift
talkative	listener	skype	gratuity
rejection	determination	accommodation	ta
complexion	curious	perk	evening
pleasuring	dedicate	dinner	relocate
charming	multitasking	vacation	lodging
patiently	teamwork	deduction	rotation

Notes: The table shows top 20 words in each of the four categories - Hard skills/Skills, Soft skills, Personality/Appearance, Job Flexibility/Benefits - for females (Column I and III) and males (column II and IV). Words are sorted in decreasing order of importance within each gender-category; *Source:* Data from the population of all job ads that specify a request for a gender, after dropping duplicates.

Table 3.7: Gender word representations and the advertised wage

Sample:	*F* Jobs		*N* Jobs		*M* Jobs	
	(I)	(II)	(III)	(IV)	(V)	(VI)
Female (skills)	–0.035***	–0.036***	–0.013***	–0.006*	–0.029**	–0.033***
	(0.007)	(0.007)	(0.003)	(0.003)	(0.012)	(0.011)
Female (soft skills)	–0.012***	–0.012***	–0.006**	0.001	0.000	–0.003
	(0.004)	(0.004)	(0.003)	(0.002)	(0.004)	(0.005)
Female (personality)	0.024***	0.010*	0.001	–0.002	0.006	0.000
	(0.006)	(0.005)	(0.002)	(0.002)	(0.003)	(0.003)
Female (flexibility)	–0.000	0.008	–0.003	0.003	0.005	0.003
	(0.018)	(0.012)	(0.003)	(0.002)	(0.011)	(0.011)
Female (others)	–0.030	–0.038***	0.041***	0.040***	0.044***	0.022
	(0.015)	(0.013)	(0.004)	(0.004)	(0.013)	(0.012)
Male (skills)	–0.019	0.010	0.012***	0.007**	–0.002	0.012
	(0.014)	(0.013)	(0.003)	(0.003)	(0.008)	(0.010)
Male (soft skills)	0.037*	0.021	0.001	0.003**	0.010	0.008
	(0.017)	(0.013)	(0.002)	(0.001)	(0.012)	(0.013)
Male (personality)	0.004	0.003	0.004	0.002	0.014	0.014
	(0.009)	(0.009)	(0.003)	(0.002)	(0.012)	(0.012)
Male (flexibility)	0.024***	0.013***	0.024***	0.013***	0.008**	–0.000
	(0.005)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
Male (others)	0.065***	0.057***	0.041***	0.018***	0.060***	0.029***
	(0.012)	(0.011)	(0.004)	(0.003)	(0.008)	(0.009)
Fixed Effects	state	occ × state	state	occ × state	state	occ × state
N	5963	5963	124313	124313	5208	5208

Notes: The dependent variable is the log of wage offered in a job. All regressions control for a set of education and experience requirement categories given in a job ad. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890 and those with non-missing wage data are 138218. Each column reports the effective number of observations used for estimations after incorporating occ × state fixed effects for the respective subsamples. This excludes job ads for which there is no variation in the dependent variable within an occ × state cell.

Appendix

3.A Technical Appendix

3.A.1 GSDMM, Preprocessing and Hyperparameter Choice

We use the following pre-processing steps on the text contained in job titles: (a) convert letters to lowercase; (b) remove non-latin characters, multiple occurrences of the same word in a job title, stop words, and words unrelated to job positions such as proper nouns; (c) remove words whose length is smaller than 2 or larger than 30 characters; (d) tokenize and lemmatize the job titles and (e) remove duplicate job titles as well as words that occur only once in the entire corpus. In our data, the resulting number of documents $D = 29316$ and the number of unique words $V = 3303$.

Note that tokenization splits a character sequence into tokens, which are meaningful semantic units for processing. Lemmatization reduces words to their base form or lemma. To implement these we use the small English model of *spaCy* trained on written text on the web such as blogs, news, comments etc. *spaCy* is an open source library used for advanced natural language processing in Python and Cython, and has pre-trained statistical models for over 60 languages. See https://spacy.io for more details.

The GSDMM algorithm first randomly assigns all documents K clusters where K is a pre-defined upper limit on the number of topics given as a human input to the algorithm. As long as K is larger than the 'true' number of clusters, the algorithm can automatically infer the appropriate number of clusters. In each subsequent iteration the algorithm probabilistically re-assigns each document one-by-one to a cluster based on two considerations: (a) sharing a more similar set of words, and (b) having more documents. As it proceeds, some clusters grow larger and others disappear until finally each cluster contains a similar set of documents. Mathematically, a document d is assigned to cluster z with probability:

$$p(z_d = z|\vec{z}_{\neg d}, \vec{d}) \propto \frac{m_{z,\neg d} + \alpha}{D - 1 + K\alpha} \frac{\prod_{w \in d}(n^w_{z,\neg d} + \beta)}{\prod_{i=1}^{N_d}(n_{z,\neg d} + V\beta + i - 1)}$$

where $\vec{z}$ is the cluster label of each document, m_z is the number of documents in cluster z, n_z is the number of words in cluster z and n_z^w represents the number of occurrences of word w in cluster z. $\neg d$ denotes that cluster label of document d is removed from $\vec{z}$. D refers to the total number of documents in the corpus, N_d is the number of words in document d and V is the total number of words in the vocabulary.

The parameter α is related to the prior probability of choosing an empty cluster. For example, when $\alpha = 0$, the probability of choosing an empty cluster is 0. The parameter β relates to homogeneity of clusters. If $\beta = 0$, a document will never be assigned to a cluster if any particular word in the document is not contained within any document in a cluster, even if the other words of the document may appear in multiple documents in that cluster. Therefore, a positive value of β should be chosen. We set the initial number of clusters $K = 750$, $\alpha = 0.1$, $\beta = .005$ and run the model for 60 iterations.

Yin and Wang (2014) use $\alpha = 0.1$, $\beta = 0.1$ and 30 iterations. However, we choose a smaller β to get more homogeneous clusters. The overall performance of the algorithm is not sensitive to α in range [0,1], and therefore, choose $\alpha = 0.1$ in line with the original work. We choose the number of iterations such that the number of clusters becomes stable and the number of documents transferred across clusters also become very small post that number. We tried up to 100 iterations and found that at approximately 60 iterations both these criteria are met. Lastly, the initial number of clusters (K) were chosen to be approximately equal to the number of clusters obtained in the manual classification using n-grams. Figure 3.6 shows that the number of clusters and the number of documents transferred across clusters initially falls sharply, and then tends to stabilize after a few iterations.

3.A.2 Preprocessing Bag-of-n-grams Logistic Regression

For preprocessing, we first remove all special characters, numbers as well as extra spaces, i.e. we retain only alphabets, and lowercase all the characters in the job text. We remove all the words indicating explicit gender preferences as mentioned in sub-section 3.2.1. If we retain these words, our algorithm's accuracy will be artificially inflated by classifying jobs largely on the basis of words which were originally used to code employers' gender preferences. We also filter out stop words such as "the", "are", "and" which are uninformative in representing the text. We use the Stopwords corpus of the Natural Language Toolkit (NLTK) version 3.5. NLTK is a python package used for NLP. For more details, see https://www.nltk.org/. We also remove words having length less than 2 or greater than 15 characters, and then lemmatize the job text using the small English model of *spaCy*.

In a bag-of-words (BOW) representation, each document is represented as a vector based on the occurrence of words in it, without taking into account their relative position in the document. This generates a matrix where each row represents a document and each column indexes a word or a set of words (also known as a token) that occurs in the corpus.

A discriminative classifier such as LR directly learns the mapping from inputs x to the class label y by fitting a hyperplane in the input feature space to separate the classes.[46] A generative model such as NB (McCallum et al., 1998), on the other hand, tries to solve a more general problem of modeling the joint probability Prob(x,y) as an intermediate step and then uses Bayes rule to calculate Prob(y|x). Consequently, LR has a lower asymptotic error, and is expected to outperform NB when the number of training examples is high enough, as in our case (Ng and Jordan, 2002).

3.A.3 TF-IDF Implementation

For a token t in document d, the $TF-IDF$ score is computed as follows:

$$TF-IDF(t,d) = TF(t,d) \times IDF(t)$$

such that,

$$TF(t,d) = \frac{N_{t,d}}{N_d} \quad \text{and} \quad IDF(t) = ln\frac{1+n}{1+DF(t)} + 1$$

where, $N_{t,d}$ is the number of occurrences of token t in document d; N_d refers to the length of document d; $DF(t)$ is the number of documents in which token t appears; and n is the total number of documents in the corpus. Additionally, the $TF-IDF$ vectors for each document are normalized to have Euclidean norm 1. Therefore, TF captures how important a token is to a document, whereas IDF scales down the weight of tokens that occur very frequently in the corpus, and hence are less informative for our classification.

3.A.4 Stratified k-folds Cross Validation

In stratified 10-folds cross-validation, for each of the 10 "folds", the model is trained on 9 folds (or 90% of the sample) and its performance is assessed using the remaining fold (or 10% of the sample) as the test set. If we use the same data for learning the parameters of the logistic regression model as well as evaluation, this will lead to overfitting, i.e. the model will perform exceptionally well on the training data, but will not generalize well. We also use $L2$

[46] The output y in our models is a binary variable indicating explicit gender preferences of employers. The input x is the bag-of-n-gram representation of text in job ads using $TF-IDF$ vectors for the LR model. In case of the NB classifier, the input x corresponds to binary-valued feature vectors indicating the presence or absence of n-grams in each job title.

regularization to prevent overfitting with regularization parameter (inverse of regularization strength) equal to 0.35 and 0.45 to calculate F_p and M_p respectively. To do this the sum of squared weights (i.e. coefficients) are multiplied by a constant C and added to the loss function. This adds a quadratic penalty to the weights as they move away from zero to prevent overfitting. A methodological issue may arise when two documents with exactly the same text are assigned different probabilities if they belong to different test sets for which slightly different training data is used. This, however, does not pose a significant challenge for us as 99.96% of the overall variance in the probabilities is explained between job texts, with the remainder explained within job texts.

3.A.5 Application of Layer-wise Relevance Propagation

Samek et al. (2016) and Srinivasan et al. (2017) show that LRP provides better explanations of DNN than sensitivity analysis for image and video classification respectively. Correspondingly, Becker et al. (2018) apply LRP to explain the predictions of DNN for audio classification. In neurosciences, LRP has been used to detect neural activity associated with motor functions using EEG data (Sturm et al., 2016) and to map cognitive state to brain activity using fMRI data (Thomas et al., 2019). In the medical domain, LRP has been applied to discover localized molecular features associated with cancer (Hägele et al., 2020; Binder et al., 2018). Others such as Horst et al. (2019) characterise gait patterns of individuals, and Lapuschkin et al. (2017) and Arbabzadah et al. (2016) identify which facial features are used by DNN for prediction of age, gender and psychological attributes using LRP. Chaturvedi and Chaturvedi (2020) apply LRP to a character based neural network model, and show how character patterns in South Asian names belonging to different religions are rooted in the distinct orthography of their language of origin.

3.A.6 Convolutional Neural Network Model Specifications

The CNN model is trained using Keras deep learning framework (Chollet et al., 2015). We initialize the embedding weights using 300 dimensional pretrained GloVe embeddings (Pennington et al., 2014) and use a dropout rate of 0.45 to prevent overfitting (Hinton et al., 2012). GloVe vectors are trained on common crawl comprising a vocabulary of 2.2 million words. For more details visit https://nlp.stanford.edu/projects/glove/. We use CNN kernel sizes k ranging from 1–5 with number of filters equal to $50 \times k$ for each kernel. We randomly initialize the kernel weights using He uniform distribution (He et al., 2015) and use ReLU activation function for the CNN layer. This is followed by a global max pooling layer. Finally, we get the probability of a document belonging to each class by converting the classification scores obtained by concatenating output from the global max pooling layer using the sigmoid function. We use the Nadam optimizer (Dozat, 2016) to train the model for 20 epochs using a minibatch size of 1024. We use balanced class weights and choose the model with the least binary cross-entropy loss on the test set.

Additional Figures and Tables

Figure 3.3: Distribution of age requirements by gender preference

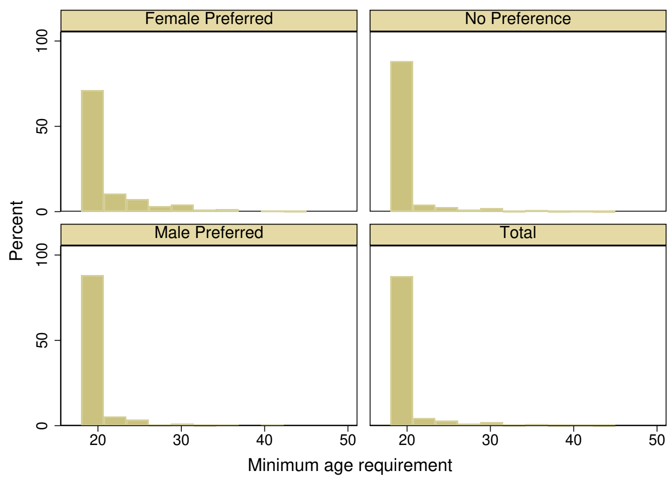

(a) Minimum age requirement by gender preference

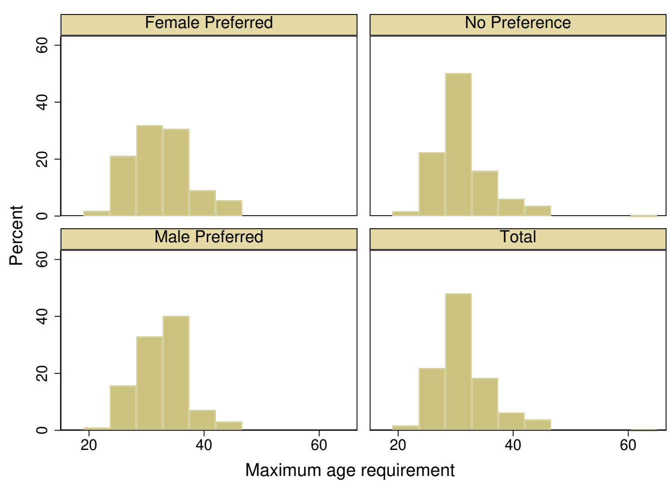

(b) Maximum age requirement by gender preference

Notes: Number of bins for construction of histograms restricted to 10.

Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2.

Figure 3.4: Distribution of ln(wage) specified in job ad by gender preference

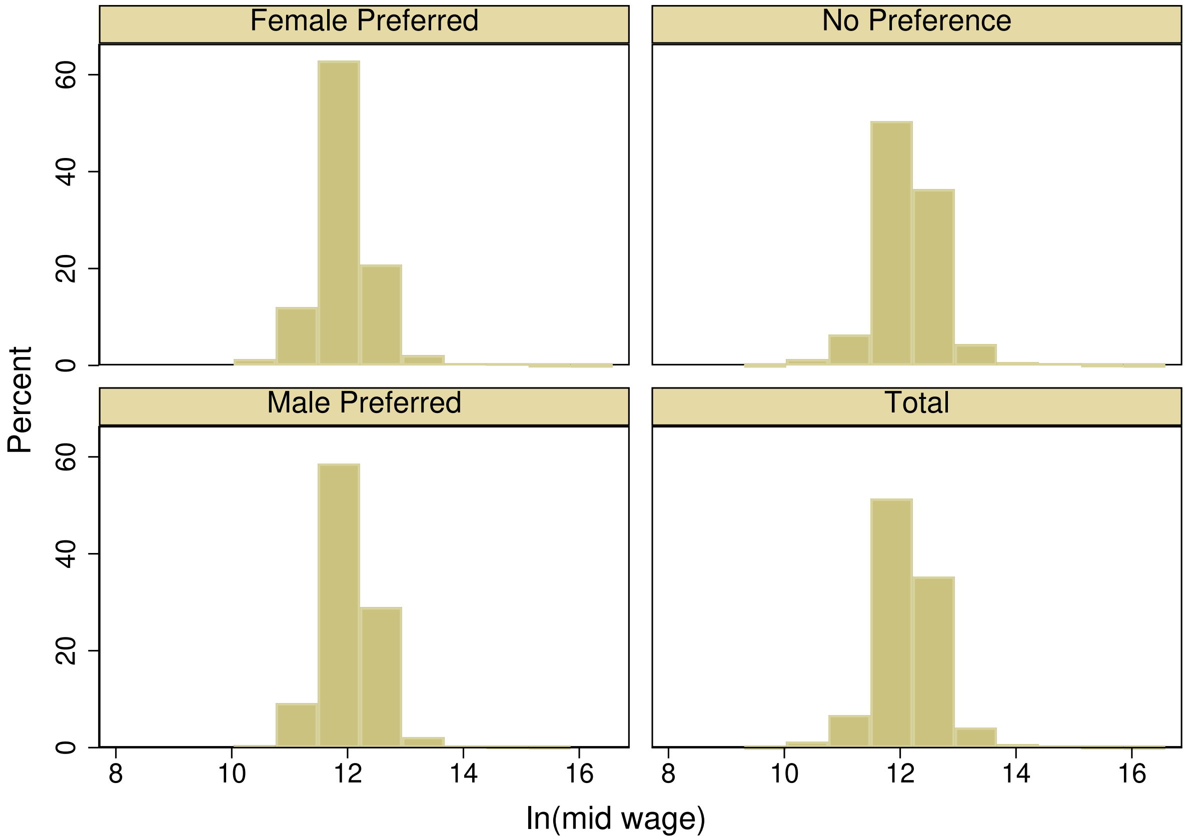

Notes: Number of bins for construction of histogram restricted to 10.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890.

Figure 3.5: Relative distribution of gender preferences by job role

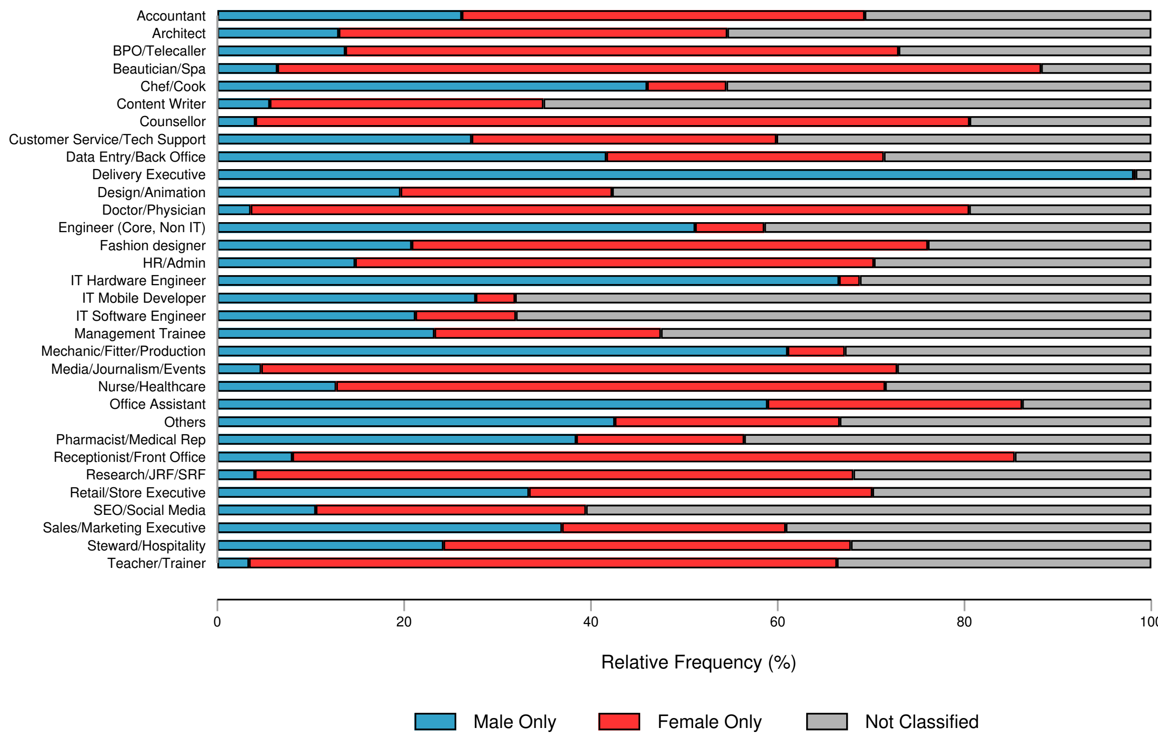

Notes: We calculate the relative frequencies by normalizing the total number of jobs in each gender preference category to one and then construct the stacked bar graph. The relative frequencies are then normalized to 100 within each job role. The stacked bars thus, provide the relative importance across job roles for each type of job ad - Male Only, Female Only and Not classified - and are *not* absolute proportions in each job role for each type of job.

Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890.

Figure 3.6: GSDMM Iterations and Clusters

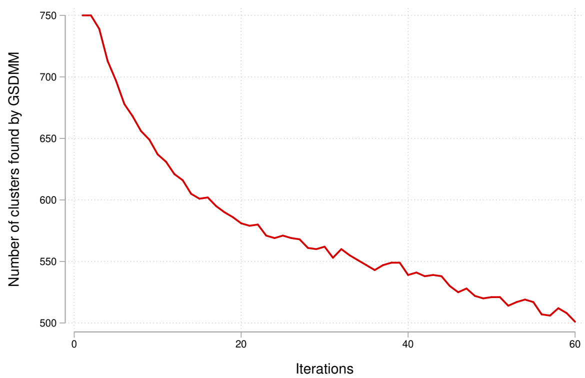

(a) Number of clusters

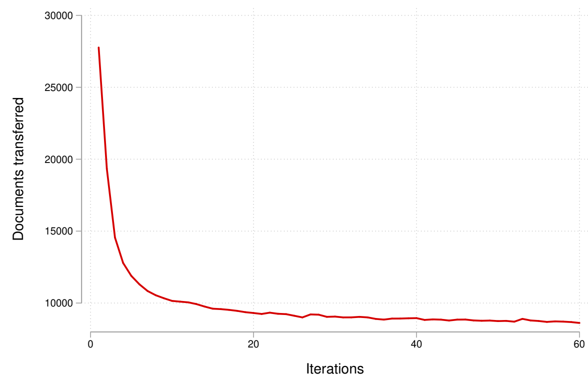

(b) Number of transferred documents

Notes: Number of clusters found by GSDMM in each iteration (subfigure a) and number of documents transferred across clusters in each iteration (subfigure b).

Figure 3.7: Predicted share of female (male) applicants with implicit femaleness (maleness)

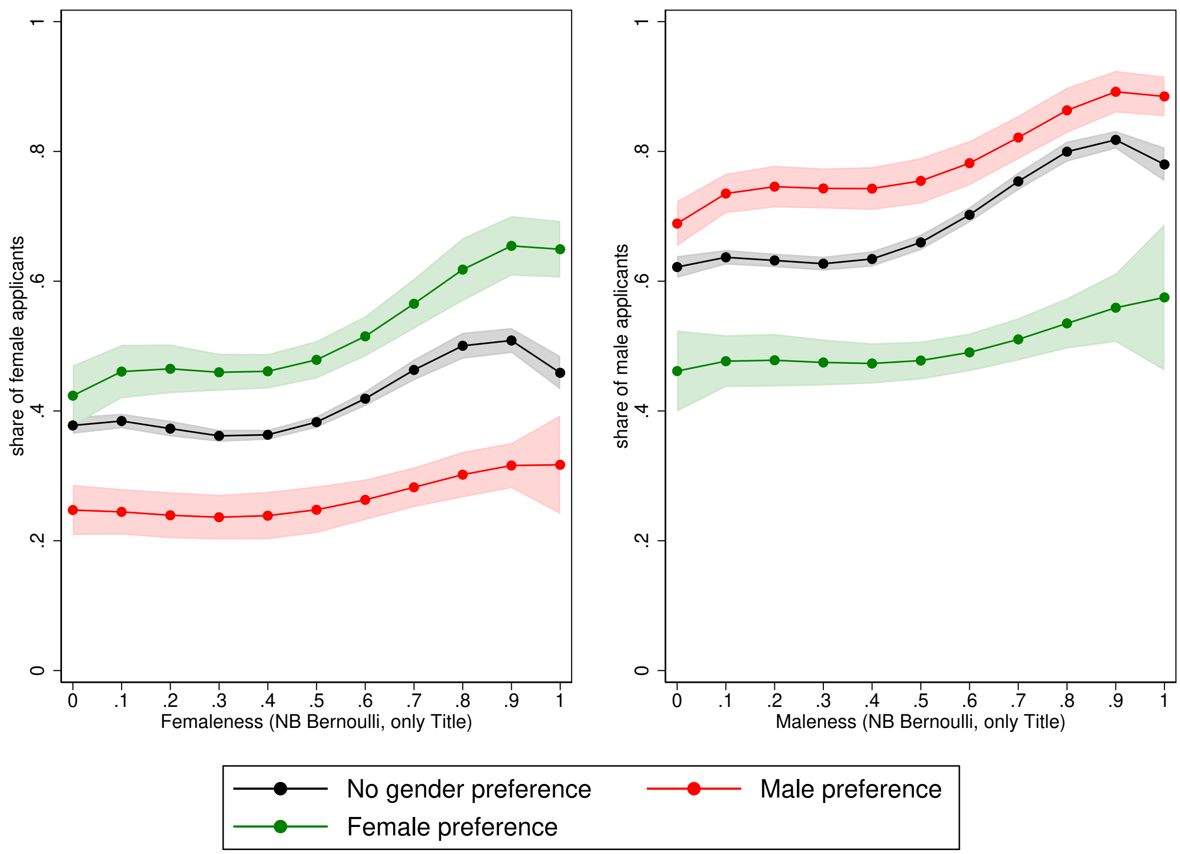

Notes: Shaded areas give the 95% confidence intervals. The measure of implicit femaleness (maleness) is constructed using a Bernoulli Naive Bayes classifier as described in sub-section 3.2.4. Predictions are based on regressing the share of female (male) applicants on explicit gender preferences, quartics in implicit femaleness (maleness), their interactions, the set of controls specified in equation (3.3.4) and (month, year) and state fixed effects. These regressions are weighted by the total number of female and male applications, and standard errors are clustered by occupation and state.

Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890.

Figure 3.8: Heat map visualization of words in distinctive job ads

- business development manager language bengali fluently speak english read write fluently speak hindi fluently speak groom must look like air hostess job role manager hr student counselling employee handle eod report share total office management bond applicable employee qualification preferable minimum graduate mba marketing master psychology candidate applicable good look smart candidate computer knowledge power point mail communication excel presentation skill age height weight proportionate height

- software trainee faculty follow subject basic computer complete knowledge ms office friendly internet advance english grammar personality development class good comunication skill basic accounting taly gst

(a) Female Preference

- sale market executive marketing executive contribute develop integrate marketing campaign task involve liaise networking range stakeholder include customer colleague supplier partner organisation communicate target audience manage customer relationship source advertising opportunity place advert press radio manage production marketing material include leaflet poster flyer newsletter newsletter dv ds write proofreading copy liaise designer printer organise photo shoot arrange effective distribution marketing material maintain update customer database organise attend event conference seminar reception exhibition source secure sponsorship conduct market research example use customer questionnaire focus group contribute develop marketing plan strategy manage budget evaluate marketing campaign monitor competitor activity support marketing manager colleague

- software trainee qualification tech sc bca mca sc fresher pass requirement candidate it computer science background prefer excellent verbal write communication skill basic knowledge it technologie quick learner able work rotational shift candidate prefer

(b) Male Preference

Notes: Panel (a) shows correctly classified job ads with an explicit female preference; panel (b) shows correctly classified job ads with an explicit male preference. Words highlighted in red reflect female associations, and those in blue correspond to male associations as captured by the CNN model based on explicit gender preferences of the employers. The color intensity reflects the strength of the attached gender association, with darker shades showing a higher association strength.

Table 3.8: Descriptive statistics, job ads with location in a single Indian state

	Prefer female	No pref.	Prefer male	Total
Education requirements:				
Other (education not specified)	0.006	0.004	0.004	0.004
None (illiterate)	0.019	0.014	0.040	0.015
Secondary education	0.114	0.099	0.317	0.108
Senior secondary education	0.304	0.263	0.267	0.265
Diploma	0.075	0.090	0.076	0.089
Graduate degree, STEM	0.040	0.089	0.055	0.086
Graduate degree, non-STEM	0.419	0.425	0.236	0.417
Postgraduate degree, STEM	0.010	0.006	0.000	0.006
Postgraduate degree, non-STEM	0.011	0.006	0.002	0.006
Experience requirements:				
$0-1$ years	0.681	0.664	0.680	0.665
$1-2$ years	0.220	0.176	0.208	0.179
>2 years	0.099	0.160	0.112	0.155
Other job requirements:				
Age requirement present	0.087	0.082	0.177	0.086
Minimum age requirement present	0.066	0.074	0.164	0.078
Maximum age requirement present	0.080	0.077	0.159	0.080
Beauty requirement present	0.112	0.057	0.063	0.059
Working night shifts specified	0.009	0.021	0.044	0.021
Advertised wage:				
Wage not specified	0.063	0.131	0.035	0.125
Annual wage, if wage specified in job ad	183231	224377	191366	221013
N (jobs with advertised wage)	6627	125763	5828	138218
Applications:				
Share of female applicants	0.507	0.319	0.140	0.321
Number of applications	17.593	42.373	31.665	40.853
N (all jobs)	7074	144777	6039	157890

Notes: Each cell gives average value of a variable in the respective sub-sample of job ads. Wages are annual wages in Rupees. Wages and experience are the mid-point of the range specified in a job ad.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2.

Table 3.9: Descriptive statistics, job applicants

	Female	Male	Total
Education:			
Other (education not specified)	0.002	0.002	0.002
None (illiterate)	0.000	0.000	0.000
Secondary education	0.004	0.016	0.012
Senior secondary education	0.030	0.068	0.054
Diploma	0.030	0.087	0.066
Graduate degree, STEM	0.535	0.545	0.541
Graduate degree, non-STEM	0.155	0.135	0.142
Postgraduate degree, STEM	0.122	0.067	0.087
Postgraduate degree, non-STEM	0.122	0.080	0.095
Experience:			
$0-1$ years	0.799	0.736	0.758
$1-2$ years	0.069	0.079	0.075
>2 years	0.132	0.185	0.166
Age:			
Age at registration	23.460	23.863	23.720
Applied wage:			
Mean annual wage	307222	294457	298955
Number of applications:			
Number of applications	6.148	6.047	6.083
N (Applicants)	374804	685931	1060735

Notes: Each cell gives average value of the variable in the respective sub-sample of job applications. Experience is given in years and is divided into four categories to correspond to the job advertisements sample.
Source: The applicant sample includes those who applied to at least one job in our job advertisement sample, and disclosed their gender.

Table 3.10: Applicant and match quality

Dependent variable:	Years of education			Matriculation score			% satisfying educ. req.	% satisfying exp. req.
	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)	(VIII)
Female preference	–0.001	–0.132***	–0.128***	0.438***	–0.558***	–0.391***	–0.003***	–0.011
	(0.015)	(0.017)	(0.018)	(0.129)	(0.113)	(0.102)	(0.001)	(0.006)
Male preference	–0.130***	–0.046	–0.045	–1.004***	–0.370	–0.285	–0.002***	0.005
	(0.032)	(0.029)	(0.030)	(0.187)	(0.194)	(0.191)	(0.001)	(0.013)
% female applicants		0.847***	0.894***		6.397***	6.433***		
		(0.036)	(0.041)		(0.245)	(0.251)		
ln(wage)			0.048***			0.666***		
			(0.008)			(0.062)		
Fixed Effects	month, occ × state	month, occ × state	month, occ × state	month, occ × state	month, occ × state	month, occ × state	month, occ × state	month, occ × state
N	156168	156168	136414	155335	155335	135697	154097	154097

Notes: The dependent variable in columns (I)-(III) is the average years of education of all applicants to a job ad. The dependent variable in columns (IV)-(VI) is the average matriculation score of all applicants to a job ad. The dependent variable in column (VII) is the fraction of applicants who satisfy an ad's education requirement and in column (VIII) who satisfy an ad's experience requirement. All regressions are weighted by the total number of applications made to a job ad, and include a set of education and experience requirement controls. The omitted category among education requirement categories includes other, illiterate, and secondary education. The omitted category among experience requirement categories is 0 to < 1 year of experience. Standard errors are clustered at the (state, occupation) level, and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01. *Source:* Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2.1. Years of education are reported in the applicant sample for almost 100% candidates while matriculation marks are observed for 95% candidates. The lower proportion of observed matriculation marks is due to some candidates not reporting these or when reporting using a CGPA scale with no information available on the conversion to percentage for such scores. All columns report the effective number of observations after incorporating occ × state fixed effects which exclude job ads for which there is no variation in the dependent variable within an occ × state cell.

Table 3.11: Descriptive statistics, PLFS Urban workers

	Female	Male	Total
Panel A: Age 16-60			
Education:			
None (illiterate)	0.159	0.075	0.094
Less than Secondary education	0.254	0.335	0.317
Secondary education	0.074	0.147	0.131
Senior secondary	0.075	0.117	0.108
Diploma	0.020	0.026	0.025
Graduate degree	0.263	0.216	0.226
Postgraduate degree	0.155	0.083	0.098
Age:			
Age	35.417	36.030	35.897
Salary:			
Annual Wage	167983	207824	199217
Observations	2954	10853	13807
LFPR	0.226	0.821	0.529
Panel B: Age 18-32			
Education:			
None (illiterate)	0.089	0.052	0.060
Less than Secondary education	0.170	0.321	0.288
Secondary education	0.075	0.140	0.125
Senior secondary	0.079	0.129	0.118
Diploma	0.028	0.035	0.033
Graduate degree	0.361	0.244	0.270
Postgraduate degree	0.196	0.079	0.105
Age:			
Age	26.417	26.436	26.432
Salary:			
Annual Wage	167490	178405	176001
Observations	1166	4382	5548
LFPR	0.242	0.774	0.518

Notes: The sample includes all urban workers in 63 majority urban districts (having at least 70% urban population) in India. Panel A includes all workers aged 16-60 while Panel B includes all workers aged 18-32. Each cell gives the average value of the variable in the respective sub-sample of workers. Age is given in years. The Labour force participation rate (LFPR) refers to proportion of individuals working majority of the year. This proportion is calculated for all individuals in the respective gender-age group.
Source: Periodic Labour Force Survey (PLFS) conduted in 2017-18.

Table 3.12: Explicit gender preference, robustness checks

Dependent variable:	any gender preference			male preference		
	(I)	(II)	(III)	(IV)	(V)	(VI)
Education requirements:						
Senior secondary	–0.016***	–0.061***	–0.020***	–0.021***	–0.065***	–0.014**
	(0.004)	(0.007)	(0.006)	(0.005)	(0.006)	(0.006)
Diploma	–0.016***	–0.078***	–0.023***	–0.024***	–0.066***	–0.028***
	(0.005)	(0.009)	(0.008)	(0.006)	(0.008)	(0.008)
Graduate degree, STEM	–0.023***	–0.094***	–0.032***	–0.016***	–0.064***	–0.019***
	(0.005)	(0.009)	(0.007)	(0.005)	(0.010)	(0.007)
Graduate degree, non-STEM	–0.017***	–0.080***	–0.023***	–0.024***	–0.086***	–0.031***
	(0.004)	(0.009)	(0.006)	(0.005)	(0.008)	(0.007)
Postgrad degree, STEM	–0.039***	–0.067***	–0.033	–0.031***	–0.098***	–0.025
	(0.011)	(0.016)	(0.017)	(0.012)	(0.014)	(0.017)
Postgrad degree, non-STEM	–0.012	–0.060***	–0.012	–0.030**	–0.098***	–0.047***
	(0.012)	(0.011)	(0.017)	(0.012)	(0.011)	(0.016)
Experience requirements:						
$1-2$ years	0.015***	0.015**	0.014***	–0.002	–0.007*	–0.007
	(0.002)	(0.007)	(0.005)	(0.002)	(0.003)	(0.005)
>2 years	–0.004	–0.010**	–0.000	0.007***	0.003	0.000
	(0.002)	(0.004)	(0.005)	(0.002)	(0.005)	(0.004)
Other job requirements:						
Age requirement present	0.053***	0.044*	0.060***	0.022***	0.031***	0.035***
	(0.006)	(0.021)	(0.011)	(0.006)	(0.010)	(0.009)
Beauty requirement present	0.031***	0.011	0.009	–0.047***	–0.030***	–0.018**
	(0.007)	(0.009)	(0.007)	(0.006)	(0.007)	(0.007)
Working night shifts specified	0.018**	0.036	0.057***	0.067***	0.087***	0.090***
	(0.007)	(0.025)	(0.016)	(0.008)	(0.031)	(0.017)
Fixed Effects	alt occ × state	firm × state	firm × occ × state	alt occ × state	firm × state	firm × occ × state
N	152402	102203	62359	152402	102203	62359

Notes: The dependent variable in columns (I)-(III) takes the value one if a job ad shows a male or female preference and zero otherwise while the dependent variable in columns (IV)-(VI) takes the value minus one if a job ad shows a female preference, zero if it does not show any gender preference and one if it shows a male preference. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects. Standard errors are clustered at the (state, job title) level and reported in parentheses; * p-value < 0.05, ** p-value $<$ 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890. Each column reports the effective number of observations used for estimations after incorporating the respective fixed effects, which vary across columns. This excludes job ads for which there is no variation in the dependent variable within the fixed effect cell.

Table 3.13: Implicit *femaleness* (*maleness*) and the advertised wage, robustness checks

	(I)	(II)	(III)
Femaleness	–0.217***	–0.310***	–0.165***
	(0.014)	(0.021)	(0.019)
Maleness	–0.015	–0.004	–0.044**
	(0.013)	(0.014)	(0.019)
Education requirements:			
Senior secondary	0.033***	–0.017	–0.026**
	(0.006)	(0.013)	(0.010)
Diploma	0.017*	0.038**	0.009
	(0.008)	(0.015)	(0.015)
Graduate degree, STEM	0.151***	0.144***	0.116***
	(0.011)	(0.028)	(0.024)
Graduate degree, non-STEM	0.058***	0.022	0.003
	(0.006)	(0.011)	(0.011)
Postgrad degree, STEM	0.368***	0.187***	0.113
	(0.046)	(0.068)	(0.072)
Postgrad degree, non-STEM	0.227***	0.225***	0.232***
	(0.036)	(0.054)	(0.088)
Experience requirements:			
$1-2$ years	0.063***	0.047**	0.012
	(0.005)	(0.019)	(0.010)
>2 years	0.288***	0.265***	0.185***
	(0.010)	(0.026)	(0.013)
Fixed Effects	alt occ × state	firm × state	firm × occ × state
N	121428	74412	41968

Notes: The dependent variable is the log wage advertised in a job ad. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value <0.05, ** p-value <0.025, *** p-value <0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of such job ads is 157890. From these, only job ads without an explicit gender preference (or N jobs) and which advertise a wage are retained. Each column reports the effective number of observations used for estimations after incorporating the respective fixed effects, which vary across columns. This excludes job ads for which there is no variation in the dependent variable within the fixed effect cell.

Table 3.14: Applications, robustness checks

Dependent variable:	total applications			share of female applications		
	(I)	(II)	(III)	(IV)	(V)	(VI)
Female preference	−6.001***	−8.913***	−6.193***	0.137***	0.171***	0.119***
	(0.701)	(1.073)	(1.216)	(0.006)	(0.010)	(0.009)
Male preference	0.717	−6.722***	2.200	−0.079***	−0.112***	−0.072***
	(3.326)	(2.548)	(3.246)	(0.005)	(0.008)	(0.008)
Education requirements:						
Senior secondary	2.006***	−0.203	1.082	0.025***	0.024***	0.015***
	(0.731)	(0.889)	(0.834)	(0.002)	(0.004)	(0.004)
Diploma	1.588	12.485***	3.232*	0.020***	−0.005	0.028***
	(1.565)	(1.612)	(1.633)	(0.003)	(0.010)	(0.007)
Graduate degree, STEM	42.872***	35.105***	14.461***	0.040***	0.041***	0.052***
	(5.252)	(3.774)	(3.098)	(0.003)	(0.012)	(0.006)
Graduate degree, non-STEM	7.652***	2.771*	−0.135	0.048***	0.081***	0.063***
	(1.385)	(1.246)	(0.976)	(0.003)	(0.009)	(0.005)
Postgrad degree, STEM	−2.561	1.929	−10.959	0.111***	0.119***	0.099***
	(6.967)	(7.198)	(16.839)	(0.016)	(0.022)	(0.027)
Postgrad degree, non-STEM	−3.574	−3.595	−3.142	0.086***	0.110***	0.095***
	(2.150)	(7.322)	(4.088)	(0.018)	(0.015)	(0.015)
Experience requirements:						
$1-2$ years	−23.396***	−10.575***	−11.317***	−0.013***	−0.016***	−0.014***
	(3.246)	(1.635)	(1.391)	(0.002)	(0.004)	(0.003)
>2 years	−42.748***	−19.256***	−20.323***	−0.033***	−0.043***	−0.032***
	(5.302)	(2.775)	(1.480)	(0.002)	(0.006)	(0.005)
Other job requirements:						
Age requirement present	−4.672***	0.515	−0.740	−0.005*	−0.014	−0.012
	(1.263)	(1.709)	(1.766)	(0.002)	(0.011)	(0.008)
Beauty requirement present	−2.853	−4.194***	−0.624	0.001	−0.003	0.001
	(2.062)	(1.185)	(0.680)	(0.003)	(0.004)	(0.004)
Working night shifts specified	17.805*	−3.487	−6.275**	−0.017***	−0.027***	−0.021*
	(8.190)	(4.176)	(2.639)	(0.006)	(0.007)	(0.010)
Fixed Effects	alt occ × state	firm × state	firm × occ × state	alt occ × state	firm × state	firm × occ × state
N	152402	102203	62359	152402	102203	62359

Notes: The dependent variable in columns (I)-(III) is the number of applicants to a job ad and in columns (IV)-(VI) is the share of female applicants. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects and regressions in columns (IV)-(VI) are weighted by the total number of applications made to a job ad. Standard errors are clustered at the (state, occupation) level and reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890. Each column reports the effective number of observations used for estimations after incorporating the respective fixed effects, which vary across columns. This excludes job ads for which there is no variation in the dependent variable within the fixed effect cell.

Table 3.15: Explicit gender preference, robustness to non-linear specifications

Dependent variable:	any gender pref	male pref	female pref
	(1)	(2)	(3)
Education requirements:			
Senior secondary	−0.028***	−0.018***	0.023***
	(0.002)	(0.001)	(0.001)
Diploma	−0.029***	−0.017***	0.022***
	(0.003)	(0.001)	(0.002)
Graduate degree, STEM	−0.045***	−0.015***	0.019***
	(0.004)	(0.002)	(0.002)
Graduate degree, non-STEM	−0.031***	−0.021***	0.026***
	(0.002)	(0.001)	(0.001)
Postgrad degree, STEM	−0.048***	−0.024***	0.030***
	(0.010)	(0.004)	(0.005)
Postgrad degree, non-STEM	−0.025***	−0.024***	0.030***
	(0.008)	(0.004)	(0.005)
Experience requirements:			
$1 - 2$ years	0.015***	0.000	−0.000
	(0.002)	(0.001)	(0.001)
>2 years	−0.012***	0.004***	−0.005***
	(0.002)	(0.001)	(0.001)
Other job requirements:			
Age requirement present	0.034***	0.018***	−0.022***
	(0.003)	(0.001)	(0.002)
Beauty requirement present	0.023***	−0.019***	0.024***
	(0.003)	(0.001)	(0.002)
Working night shifts specified	0.004	0.036***	−0.045***
	(0.004)	(0.002)	(0.003)
Fixed Effects	job role, state	job role, state	job role, state
N	157069	157097	157097

Notes: Column (I) reports average marginal effects from a probit where the dependent variable takes the value one if a job ad shows a male or female preference and zero otherwise. Columns (II) and (III) report average marginal effects from an ordered probit where the dependent variable takes the value minus one if a job ad shows a female preference, zero if it does not show any gender preference and one if it shows a male preference. The omitted category among education requirement categories is other (education not specified), illiterate, and secondary education. The omitted category among experience requirement categories is none to less than a year of experience. All regressions include (month,year) of job posting fixed effects. Standard errors are reported in parentheses; * p-value < 0.05, ** p-value < 0.025, *** p-value < 0.01.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2. The final number of job ads is 157890. Each column reports the effective number of observations used for estimations after incorporating job role and state fixed effects. This excludes job ads for which there is no variation in the dependent variable within a job role or state cell.

Table 3.16: Descriptive statistics: Stereotypes

	F Jobs	N Jobs	M Jobs	All jobs
Female (skills)	0.315	0.421	0.185	0.407
Male (skills)	0.160	0.389	0.304	0.375
Female (soft skills)	0.322	0.112	0.202	0.125
Male (soft skills)	0.039	0.078	0.026	0.074
Female (personality)	0.136	0.073	0.087	0.077
Male (personality)	0.053	0.070	0.037	0.068
Female (flexibility)	0.107	0.177	0.134	0.172
Male (flexibility)	0.155	0.047	0.306	0.062
Female (others)	6.396	5.416	3.800	5.398
Male (others)	3.816	5.452	7.456	5.455
Observations	7074	144777	6039	157890

Notes: Each cell gives the average (non-standardized) value of a variable in the respective sub-sample of job ads. The gender association scores for each word are obtained by applying layer-wise relevance propagation technique to the convolutional neural network model based on explicit preferences of employers. The score for each gender $\times$ stereotype category is then obtained for each job ad as described in Sections 3.4.2 and 3.4.3.
Source: Data from the population of all job ads and applicants on the portal, subject to the restrictions described in section 3.2.

Chapter 4

Group Size and Political Representation Under Alternate Electoral Systems

4.1 Introduction

Representation of ethnic groups in democratic governments is an important determinant of their welfare. This is especially true for smaller groups or "minorities" who might be more vulnerable to exclusion.[1] Political representation ensures that they can voice their interests and desires to the government more easily which, in turn, gives them greater access to public resources and makes public policy more inclusive.[2]

However, the minorities within a country often vary widely in their population sizes. We ask: do larger minorities necessarily enjoy greater representation and access to resources? The answer to this question seems far from obvious. Consider the case of South Africa, where the ethnic minorities vary substantially in their population shares—ranging from 1.5% (the Ndebele) to 18% (the Xhosa).[3] However, all of them share similar level of representation in the government.[4] On the other hand, the Asante in Ghana (16%), despite being much smaller than the northern groups Mole-Dagbani,

[1]For ease of exposition, we refer to all the groups that are not the largest group of a country to be "minorities" and define the largest group to be the "majority." In more than 80% countries in our sample the "majority" group indeed have absolute majority in the country's population. Moreover, our data show that on average, only a third of minority groups get any representation in the national executive of democracies during the post-World War II period. In contrast, the "majority" group is almost always represented.

[2]Previous works show that representation fosters trust and approval in government decision-making (Banducci et al. 2004), engenders greater political participation among minority group's members (Bobo and Gilliam 1990), and consequently, improves allocation of public resources towards them (see Cascio and Washington (2013) for the case of African Americans in the US and Besley et al. (2004, 2007) etc., for the case of minority castes and tribes in Indian villages).

[3]The Zulu are the majority group in South Africa comprising more than 23% of population.

[4]Same is true of minorities in Indonesia. The Papuans, the Acehnes (about 1% population share each), the Malays (3%), and the Sundanese (13%) wield similar level of power in the Indonesian cabinet.

Gurma and Grusi (25%), have significantly *higher* representation in the government.[5]

What accounts for such differences in the patterns of representation? Since these countries adopted different electoral systems,[6] this motivates us to ask whether different electoral systems provide differential incentives for political parties to represent minorities of varying population sizes. We examine this issue first theoretically, and then empirically across a large number of democracies in the post-war period. We focus on two broad categories of electoral systems—majoritarian (MR), where elections are typically contested over single member districts, and proportional representation (PR), where seats are allocated to parties in proportion to their vote share in large multimember districts.

To contrast PR and MR systems, we propose a model with three groups (i.e., one majority and two minorities) and two parties in a probabilistic voting setup. In our model, political parties compete for votes from all groups and promise representation in the government to each group as platforms. Representation determines the per capita private transfer of government resources targeted towards group members.[7] We look at equilibrium representation (and consequently, per capita resource allocation) for various population share compositions of the two minorities (keeping the majority group's size fixed). We show that under PR, group size of minorities has *no effect* on their representation in the national government, whereas it has an *inverted-U shaped* effect under MR. Therefore, under MR, increase in group size initially improves representation. However, there is an "optimal" minority size above which its representation and per capita resource allocation begin to fall.[8]

The result is in contrast to the theoretical predictions of Trebbi et al. (2008) who find that, in the context of Whites and Blacks in southern US municipalities, access to power for the Black minority never falls with its population share within any electoral system, and rises eventually with population share in PR.[9] We show that this result gets modified when we allow political parties to compete for all groups in a multiple minority context.[10] Since we are concerned with representation

[5] Similarly, in Nepal the Adivasi group (31%) is larger than the Madhesis (12%), though both were historically marginalized by the Hill Hindu Elites (the majority group). Yet, following Nepal's democratization in 2006, the Madhesis gained representation in the national government, while the Adivasis continue to remain excluded from power.

[6] South Africa and Indonesia follow the proportional representation system, while Ghana and Nepal adopted the majoritarian system.

[7] This is motivated by the idea that many of the government provided goods are private in nature or local public goods (such as street lamps) that could be targeted to ethnically homogenous areas. This is consistent with findings in Ejdemyr et al. (2018) who show that politicians strategically target local public goods towards coethnics in Malawi. This is further substantiated in Franck and Rainer (2012) and Dickens (2018) who find evidence of ethnic favoritism by political leaders across 35 sub-Saharan countries and provide evidence that government appointments are linked with ethnicity. Our results are consistent as long as the private/local public goods form a sizeable chunk of government transfers to minorities whose cost of provision increases with group size. This assumption allows us to show that our empirical results are driven by private rather than pure public goods.

[8] The majority group under both systems gets higher representation and larger per capita resource allocation compared to both the minorities.

[9] The paper explains the *choice* of electoral system by incumbent whites after the effective enfranchisement of black population in the southern US municipalities following the Voting Rights Act, 1965. We, on the other hand, examine how minorities of differing sizes fare under a given electoral system.

[10] Trebbi et al. (2008) abstract away from parties and assume that the voters' ethnic identity and candidate support are aligned—all whites vote for one candidate and all blacks for the other. In our model we relax

in national governments, the assumption of multiple minorities seems reasonable. Further, in most countries, the major national parties do attempt to court multiple groups, even though certain groups may have inclinations towards specific parties.[11]

The result we get for the PR system is a straight forward implication of the standard probabilistic voting model with multiple groups. In PR system, parties essentially maximize votes. There are two opposing forces that result in a group's representation being unresponsive to its population share. Consider two minority groups with one being larger than the other. Though offering higher representation (and hence, per capita transfers) to the larger group gets a party more total votes, it is cheaper for a party to attract a higher *share* of voters from the smaller group. When representations are equal, these two forces balance each other out across groups.

In MR, on the other hand, parties want to win constituencies, and hence, they have to consider settlement patterns of groups across constituencies, i.e., over space. Representation of groups, therefore, depends on their settlement patterns. However, for any given population shares of groups at the national level, there are many possible ways they can be distributed across constituencies. Hence, to comment on how group size affects representation, we must first understand how changes in group size map into changes in the group's spatial distribution. This is understandably a hard problem to solve. Previously researchers have simplified the problem by either assuming only two groups –one majority and one minority (Trebbi et al., 2008) or assuming that the population composition of groups are identical across all constituencies (Milesi-Ferretti et al., 2002).

We, on the other hand, propose a novel and parsimonious framework that models the spatial distribution of three groups across an arbitrary number of constituencies to address this issue, and therefore, manage to relax both kinds of restrictions imposed by the earlier methods. We use insights from the urban geography literature, specifically the *settlement scaling theory* proposed by Bettencourt (2013), to map population share of a group at the national level to its settlement area over space. We postulate that the area occupied by a minority group has a *concave* relationship with its population share. Intuitively, if the benefit of living in an area is increasing in the density of own group members living in the area (due to positive network effects), then we should observe that larger groups live more densely, giving rise to the concave relationship. Even though the settlement scaling theory is proposed in the context of cities, we argue and then show empirically that the forces at play are general enough to explain settlement patterns of population groups in countries as well.[12] We take the concave relation as exogenously given. Majority group is assumed to be present in all parts of the country. This allows us to characterize the groups' distribution across constituencies by imposing minimal structure on the problem. The concave relationship implies that larger minorities live more densely compared to smaller ones. Therefore, if one minority group becomes "too large" in size, it can suffer a geographical disadvantage in MR by becoming "too concentrated" in only few constituencies. This leads to wastage of votes and parties respond to it by diminishing the promised representation (and allocation) to that group. This is at the core of the inverted-U shaped result that emerges as the equilibrium in our model.

this assumption.

[11]We allow the parties to have differential incentives to attract votes across different groups.

[12]The theoretical model in Bettencourt (2013) predicts that the elasticity of the relationship between area of settlement and population should be 0.67. We estimate this elasticity for ethnic groups within countries and, interestingly, find an estimate of 0.63 which is statistically indistinguishable from 0.67. We elaborate on this in section 4.3.3 and 4.6.1, respectively.

In the empirical section, we show evidence in favor of the concavity assumption and test the comparative static results of the model using measures of political representation as well as per capita resource allocation. We compile an ethnicity level panel dataset comprising 421 minorities across 87 countries during 1946–2013 by triangulating various sources described in section 4.4.1. Our main measure of political representation is an indicator that takes value one if a group has *any* representation in the national government and zero if it is either powerless or discriminated by the state. The indicator, therefore, captures the extensive margin of political representation. We use nightlight luminosity per unit area for each group (calculated using GIS maps of settlement areas of ethnic groups) as a proxy for allocation of public resources by the government. As we argue in section 4.6.2, existing evidence shows that nightlight luminosity is highly correlated with provision of electricity—a publicly provided good which is subject to political influence—as well as provision of other public goods. Additionally, we compile a cross-sectional data on existing length of road per unit area at the ethnicity-country level for 54 countries to show robustness of our results to alternate measures of public resource allocation.

Consistent with our theoretical predictions, the result shows a statistically significant inverted-U shaped relationship between population share and political inclusion under MR and no relationship under PR. The predicted "optimum" population share for minorities in MR countries is estimated to be 0.26. The result is robust to a number of alternate specifications and sample restrictions. Importantly, the result is replicated using logarithm of nightlight emissions per unit area in the settlement area of a group as the dependent variable. Further, using cross-sectional data on road length per unit area for ethnic groups we show that same patterns emerge with this measure as well.

We are, however, cognizant that the electoral system of a country might not be exogenous. Political actors in positions of power may strategically choose electoral systems that maximize their chances of winning, as Boix (1999) and Trebbi et al. (2008) show. This means that the electoral system at the time of democratization of a country, and even changes in it later may depend on existing power distribution across groups (Colomer 2004; Persson and Tabellini 2005). We address this endogeneity problem by looking at a subsample of erstwhile colonies. Consistent with Reynolds et al. (2008), we show that electoral systems of former colonial rulers systematically predict electoral systems of the colonies post-independence. We, therefore, use this as an instrument for the electoral system of a colony.[13] The exclusion restriction would hold in our specification even if the colonial ruler's electoral system is correlated with the average representation of minorities. As long as it doesn't have any direct *differential* effect on minorities of different population sizes, our identification strategy remains valid.[14] The two-stage-least-squares estimates replicate our results for both political representation and measures of resource allocation.

Remarkably, the result also holds when we compare *same* group present in more than one country within a continent and exploit the plausibly exogenous variation in its population share

[13]We restrict our sample to colonies which democratized not too long after independence. We use a maximum lag of 30 and 50 years between independence and democratization for our analysis. See Sections 4.5.2 and 4.6.3 for a detailed discussion.

[14]For example, British colonies could have had a more liberal regime for the minorities as compared to the Spanish or Portuguese colonies. This may have implications for the representation of minorities in those colonies once they became independent. This, however, would not threaten our analysis as long as the British were not *differentially* liberal with respect to minorities of different sizes. This is likely to be the case as the legal codes of colonizers didn't differentiate across minorities of varying population shares.

across countries. In this strategy, the variation comes primarily from a group falling unequally on two sides of the national boundary.[15] This strategy heavily restricts our sample and consequently, our sample size falls by more than 80%. Even in the reduced sample, we find an inverted-U shaped relation in MR and no relation in PR for both political representation and nightlight. The coefficients for the nightlight regression are, however, imprecise, presumably due to small sample size.

Our work is related to a large literature examining the effect of electoral systems on public policy and other political outcomes. Myerson (1999) and Persson and Tabellini (2002) discuss and extensively review the literature on theoretical aspects of electoral systems. Some of the outcome variables that have been studied with regard to effects of electoral systems are party polarization (Dow 2011), public goods provision and redistribution (Lizzeri and Persico 2001; Milesi-Ferretti et al. 2002; Persson and Tabellini 2004), trade policy (Rickard 2012), corruption (Kunicova and Rose-Ackerman 2005), public attitude towards democracy (Banducci et al. 1999), voter turnout (Herrera et al. 2014; Kartal 2015), and incentive to engage in conflict (Fjelde and Höglund 2016).[16]

Some papers such as Moser (2008) and Wagner (2014) have compared differences in the level of minority representation across the two systems by exploiting the variation in electoral systems over space and time in specific countries (Russia and Macedonia, respectively). In both cases the authors argue that settlement pattern of minorities are important when analyzing change in their representation across electoral systems. Our analysis accounts for this and points out the exact nature of this influence—both theoretically and empirically. Moreover, while these papers look at the level of power enjoyed by minorities in aggregate, we focus on differential access to power received by minorities of differing sizes *within* a system. Our result has important implications for power inequality between minorities. It suggests that PR distributes power more equally across minority groups, and hence, their (per capita) resource inequality is also minimal. The implication for inequality in the MR system is more nuanced. Our result suggests that small and large minorities might enjoy similar level of power and material well-being in MR countries while the mid-sized groups enjoy a greater access to and benefit from the government.

4.2 Brief Description of Electoral Systems

The decline of colonialism and autocratic rule, and a transition towards democracy has characterized the world in the post-World War II period. An interesting aspect of this wave of democratization is the adoption of different electoral systems by the newly emerging democracies. On one hand, we have MR in which elections are typically contested over single member districts. The candidate or party with a plurality or an absolute majority in a district wins the district and parties generally attempt to win as many districts as possible. Among MR systems, single member district plurality (SMDP)—where individuals cast vote for one candidate in single member district and the candidate with the most votes is elected—is the most common. SMDP system is currently followed for legislative elections in countries such as India, Nigeria and United Kingdom among others. Around 63% of country-year

[15] Dimico (2017) uses a similar identification strategy to identify the effect of group size on its level of economic performance in the African continent.

[16] In a related work, Ghosh and Mitra (2019) show that the level of discrimination against minorities across democracies and autocracies depends on the size of the majority group.

observations in our dataset that follow MR have this system.[17]

In contrast, in PR, parties typically present list of candidates and seats are allocated to parties in proportion to their vote share in multimember districts. This reduces disparity in vote share at the national level and seat share of a party in the legislature. Examples of countries that currently have PR system are Argentina, Belgium, South Africa and Turkey among others. Figure 4.1 depicts the countries with MR and PR systems in 2013. The PR countries vary in terms of institutional details such as district magnitude, quotas and remainders—which may have implications for proportionality (Gallagher 1991; Gallagher 1992; Blais and Lago 2009). We, however, focus only on the broad distinction between MR and PR to sharply capture the differences in incentives faced by political parties.[18]

We discuss the trends in choice of MR and PR by countries over the decades in appendix section 4.A. However, one aspect of the choice is worth highlighting here—namely the role played by colonial history in shaping the electoral systems of the colonies. Most of the countries that were once British and French colonies adopted MR while those that had been colonized by Belgium, Netherlands, Portugal and Spain adopted PR. We use this aspect of the choice of electoral systems in the empirical analysis to address causality.

4.3 Model

4.3.1 Set Up

In this section we develop a probabilistic voting model with two parties *à la* Persson and Tabellini (2002). There are three groups of voters. Each group has a continuum of voters of mass n_j with $\sum_{j=1}^{3} n_j = 1$. We will treat group 3 as the majority group and groups 1 and 2 as the minorities. Therefore, $n_3 \in (0.33, 1)$. Voters have preferences over private transfers made by the government. These transfers can be targeted at the level of groups but not at the individual level. We represent individual preference of any voter in group j by the utility function:

$$U(f_j) = \ln f_j$$

where f_j denotes per capita private transfers to group j. Our assumption of the specific utility function ensures that the utility function is strictly increasing, strictly concave, and satisfies the Inada conditions—requirements which are necessary for our results to hold. This is a simplifying assumption and the results can be generalized to cases when utility function is not the logarithm of the per-capita transfer. f_j is completely determined by the political processes of a country. Before election takes place, the two political parties A and B simultaneously announce the group composition of the government that they will form in the event of an election win. In principle, both parties can

[17] Another variant of MR systems is a two-round system (TRS). In TRS candidates or parties are elected in the first round if their proportion of votes exceeds a specified threshold. Otherwise, a second round of elections takes place—typically one or two weeks later—among the top candidates. France and Mali currently employ TRS for parliamentary elections.

[18] Some countries also use mixed systems which are a combination of both MR and PR. However, we do not include them in our empirical analysis.

offer representation to all the groups. We define group j's representation in the government promised by party h, G_j^h, as simply the total number of government positions announced by party h in favor of group j. G_j^h, determines how much per capita transfer voters of group j will get if party h comes to power. We denote this as follows:

$$f_j^h = f(G_j^h) \quad \text{or} \quad G_j^h = f^{-1}(f_j^h).$$

More representation in government is always beneficial for group members, i.e., $f'(G_j^h) > 0$. Recent research, for example, by De Luca et al. (2018) shows that ethnic favoritism is a global phenomenon and not just restricted to poor countries with weak institutions. This implies that the assumption of monotonic relation between group's political representation and per capita transfers is valid more broadly. Since representation in government determines the individual level payoff of the voters, political parties commit to allocation of government positions as their platforms during the election. In the following analysis, we use f_j^h directly as a choice variable of the parties, since representation in government (G_j^h) and per capita transfer (f_j^h) are synonymous in our model. Any voter i belonging to group j votes for party A if:

$$U(f_j^A) > U(f_j^B) + \delta + \sigma_{i,j}$$

where $\delta \sim U[\frac{-1}{2\psi}, \frac{1}{2\psi}]$ and $\sigma_{i,j} \sim U[\frac{-1}{2\phi_j}, \frac{1}{2\phi_j}]$ are preference shocks to the voter.

This is a standard probabilistic voting set up where δ can be interpreted as population wide wave of support in favor of party B (relative to A). $\sigma_{i,j}$ represents (relative) ideological bias of a member i of group j towards party B. ϕ_j is the height of the p.d.f. of the $\sigma_{i,j}$ distribution. It measures the responsiveness of group j voters to private transfers by parties. A larger value of ϕ_j would imply that for the same increase in promised per capita transfer by any party, a greater proportion of group j voters would sway in favor of that party. We assume that minority groups 1 and 2 are identical in their political responsiveness to transfers, i.e., $\phi_1 = \phi_2 = \phi$. Group 3 (the majority group) is more responsive to transfers compared to the minorities, i.e., $\phi_3 > \phi$. This assumption is motivated by the observation that minorities often have stronger attachments to specific parties owing to historical factors. Consequently, this makes them less pliable compared to the majority group from the parties' point of view. Values of ψ and ϕ_j are known to both the parties. The total government budget is exogenously fixed at S. Each party h maximizes the probability of forming government p_h by choosing f_j^h subject to the budget constraint:

$$\sum_{j=1}^{3} n_j f_j^h \leq S$$

In proportional system p_h is the probability that vote share is larger than 0.5, while in the majoritarian system it is the probability of winning more than half of the electoral districts. We assume that in majoritarian system there are K equally populated electoral districts of population $\frac{1}{K}$ each. We denote by n_j^k the population share of group j relative to population in district k. Therefore,

$$\sum_{j=1}^{3} n_j^k \;=\; 1 \quad \text{for all } k = 1, 2, \ldots, K$$

$$\text{and} \;\; \frac{1}{K}\sum_{k=1}^{K} n_j^k \;=\; n_j \quad \text{for } j = 1, 2, 3.$$

We compare equilibrium political representation in single district PR system with that in K district MR voting system.

4.3.2 Equilibrium Characterization

Since the parties are symmetric, we have policy convergence in equilibrium, i.e., both parties choose the same equilibrium policy in any system. The following two propositions characterize the equilibrium allocation of resources (and hence, equilibrium representation) under the two systems.

Proposition 1. *Under a single district proportional representation voting system, group size n_j of a minority has no effect on equilibrium representation G_j^* and equilibrium transfer f_j^*. In equilibrium:*

$$\phi_j U'(f_j^*) = \phi_l U'(f_l^*) \quad \forall\, j \neq l. \tag{4.3.1}$$

We relegate all proofs to appendix section 4.I. Proposition 1 implies that under PR, minority groups 1 and 2 would receive identical per capita transfers irrespective of their population shares, i.e., $f_1^* = f_2^*$ for all n_1 and n_2. To understand the result intuitively, consider the case where group 1 is the larger minority, i.e., $n_1 > n_2$. Suppose that f_1 and f_2 are the initial transfers promised by any party. Further, consider the party taking away $\epsilon > 0$ per capita transfer from group 1 and reallocating it to group 2. The per capita transfer of group 2, therefore, would increase by $\frac{n_1\epsilon}{n_2} > \epsilon$. This highlights the fact that it is always cheaper to increase per capita transfer of the smaller group. This reallocation, for a small ϵ, would cost the party $n_1\phi U^{'}(f_1)$ votes from group 1 and would increase votes from group 2 by $n_2\phi U^{'}(f_2)\frac{n_1}{n_2}$. Since in PR the political parties maximize votes, the party would prefer to reallocate as long as the gain and the loss from reallocation are different. It is obvious that when $f_1 = f_2$, they equalize. Therefore, even though vote *shares* of the smaller group are cheaper to buy, the return to a party for doing this (in terms of *total* votes) is lower, precisely because the group is small. These two opposing forces balance each other out in equilibrium, giving us the result.

Moreover, we get that the majority group gets higher per capita transfer compared to minorities, i.e., $f_3^* > f_1^* = f_2^*$. This is a direct result of our assumption that majority group voters are easier to sway through electoral commitments and hence, parties compete more fiercely for their votes.

The following result characterizes the equilibrium transfers in MR:

Proposition 2. *Under the majoritarian voting system with K districts, the following set of equations characterizes the equilibrium transfers (f_1^*, f_2^*, f_3^*) announced by both parties:*

$$\phi_j U'(f_j^*) \sum_{k=1}^{K} \frac{n_j^k/n_j}{\sum\limits_{j'=1}^{3} \phi_{j'} n_{j'}^k} = \phi_l U'(f_l^*) \sum_{k=1}^{K} \frac{n_l^k/n_l}{\sum\limits_{j'=1}^{3} \phi_{j'} n_{j'}^k} \quad \forall\ j \neq l \tag{4.3.2}$$

This is a general characterization result for any arbitrary distribution of groups across constituencies. We emphasize two aspects of the result above. Firstly, the characterization implies that equilibrium representation and transfer to groups under MR depends on the population shares. Importantly, the transfer also depends on distribution of groups across electoral districts, suggesting that *settlement patterns* of groups across districts or over space are important in determining the exact nature of the relation between group size and transfers. Moreover, if all groups have the same responsiveness to transfers, i.e., if $\phi_1 = \phi_2 = \phi_3$, then equation (4.3.2) collapses to equation (4.3.1). Therefore, heterogeneity in responsiveness across groups—especially majority and minorities is critical for group size to matter in MR system.

To explore this issue a little further we rewrite equation (4.3.2) as the following:

$$\phi_j U'(f_j^*) \frac{\sum_{k=1}^{K} \omega^k n_j^k}{n_j} = \phi_l U'(f_l^*) \frac{\sum_{k=1}^{K} \omega^k n_l^k}{n_l} \quad \text{where } \omega^k = [\sum_{j'=1}^{3} \phi_{j'} n_{j'}^k]^{-1}.$$

ω^k is therefore the inverse of the average responsiveness of district k, and $\sum_{k=1}^{K} \omega^k n_j^k$ is the weighted average of the group j's shares across districts with ω^k as the weights. Thus, the proposition above states that in MR, a group will get higher political representation and private transfers relative to another group if it is concentrated more in districts having a less responsive mass of voters, i.e., if the group has a higher correlation between n_j^k and ω^k. Since the majority group is more responsive, it therefore follows that a minority group would gain if it is concentrated more in districts with low majority group population. This happens because parties in MR wish to win electoral districts (as opposed to votes). Therefore, if a minority is settled in districts where the majority group is relatively scarce, the group becomes attractive to the political parties for the purposes of winning those districts. This logic plays an important role in determining the nature of the comparative static exercise we perform in the following section.

4.3.3 Spatial Distribution of Groups and Comparative Statics

In this section, we study equilibrium representation and transfers in MR for minorities of differing group sizes. Specifically, we see how equilibrium outcomes change when we vary n_1 and n_2, keeping the majority population share n_3 fixed. Our comparative static exercise, therefore, looks at the effect of changing n_1, holding n_3 constant. Now, any change in the composition of population shares of minorities at the national level changes their distribution across districts, i.e., the values of n_1^k and n_2^k for all k. Therefore, even though proposition 2 characterizes the equilibrium for any given profile of population shares of groups, it is hard to comment on the nature of comparative static result without specifying how changes in the population shares of groups lead to consequent change in their spatial distribution across electoral districts. Below we provide a framework to incorporate this concern in our model.

We first normalize the total area of the country to 1. We denote by A_j the measure of the area where group j has presence and postulate that $A_j = n_j^{\alpha_j}$ for some $\alpha_j \geq 0$.[19] We assume that for group 3 (i.e., majority group) $\alpha_3 = 0$, or $A_3 = 1$, i.e., the majority group is dispersed over all the space in the country. This is borne out in our data wherein the majority group is dispersed all over the country's area in 44% of the cases as opposed to 22% for the minorities. In the remaining cases, the majority group's settlement area has a significant overlap with the minorities. For groups 1 and 2, we consider two possibilities. In one case, we assume $\alpha_1 = \alpha_2 = \alpha > 0$, i.e., both minorities are geographically concentrated in some region of the country. In the alternative scenario we allow group 2 to be dispersed and group 1 to be concentrated, i.e., $\alpha_1 = \alpha$ and $\alpha_2 = 0$.[20]

Importantly, we take $\alpha < 1$ for groups that are geographically concentrated, i.e., the settlement area of a group has *concave* relationship with its population share. This assumption is motivated by the insight from literature on urban geography. Specifically, Bettencourt (2013) provides a parsimonious theoretical framework to predict the relationship between population and area of settlement (and other characteristics of the population, such as network length, interactions per capita etc.) in the context of cities. He argues that the benefit of living in a city is increasing in the population density of the area. This would be true because for the same distance travelled, an individual will have larger number of productive interactions with people. On the other hand, the travel cost is increasing in the diameter of the city, i.e., it is proportional to the square root of the area. City size is in equilibrium when the benefit and cost are equalized. The equilibrium relationship is therefore given by $A = c_0 n^{\frac{2}{3}}$, for some constant c_0. Therefore, he provides a theoretical prediction of the elasticity of the relationship. He further shows that for a sample of cities in the USA, the prediction is indeed valid. We estimate the value of α in the context of ethnic groups in our data and, surprisingly, find the same result (i.e. $\alpha = 2/3$). We discuss this in Section 4.6.1. This assumption will turn out to be important for the result we derive below.

Now we consider dividing the country in K equally populated electoral districts. Note that in the case where both minorities are geographically concentrated, we have three types of districts: (i) group 3 is present with only one minority group in the district, (ii) all the three groups are present, and (iii) only group 3 is present. The last type of district will not be there if group 2 is also dispersed. For us, the most important type of district is the one where all groups are present. Since the majority group is present everywhere, the proportion of this type of district is determined by the overlap region of the settlement areas of the two minorities. We denote by $A_{1\cap 2}$ the measure of area where groups 1 and 2 overlap and correspondingly we define the overlap coefficient (also known as

[19] Note that the same space can have presence of multiple groups, and therefore, $\sum_{j=1}^{3} A_j$ need not be one. If groups overlap over space, $\sum_{j=1}^{3} A_j$ could, in fact, be larger than one.

[20] If all groups are dispersed then the population distribution of groups in the country is replicated in each of the districts individually and consequently, the result for MR collapses again to the PR case.

the Szymkiewicz-Simpson coefficient) as:

$$O = \frac{A_{1\cap 2}}{\min\{n_1^\alpha, n_2^\alpha\}}$$

We, therefore, have $O \in [0,1]$. With these objects defined, we state the main result that establishes the relationship between group size and political representation for minorities in MR systems.

Proposition 3. *We state the results separately for the two cases that we consider:*

1. *If group 2 is also concentrated, then G_1^* follows an inverted-U shaped relation with n_1 with the peak of political representation at $n_1^* = \frac{(1-n_3)}{2}$ if and only if $O > O^*$ for some $O^* \in (0,1)$.*

2. *If group 2 is geographically dispersed, equilibrium political representation of group 1, G_1^*, follows an inverted-U shaped relation with n_1 with the peak of political representation at $n_1^* = (1-\alpha)^{\frac{1}{\alpha}}$.*

The result states that when both groups are concentrated, the equilibrium representation of (and consequently, transfers to) both groups have an inverted-U shaped relationship with group size. The intuition behind this result follows from the discussion of proposition 2. Our assumption about concave relationship between group population share and area implies a concave relation between population share and population density. This means that aggregate minority population density in the overlap region would be highest if the minorities are equal sized (i.e., $n_1 = n_2 = \frac{(1-n_3)}{2}$). As their population shares diverge from each other, i.e., as one becomes larger and the other smaller, their aggregate population density in the overlap region decreases. Since the entire overlap region is in type (ii) districts, divergence in the population shares of minorities away from the "mid-size" would imply that in those districts the share of majority group increases. This, according to our earlier discussion, harms both minorities, as they become concentrated in districts with larger majority share. The minority group which is getting smaller, therefore, loses out in both type (i) and (ii) districts. The group which is getting larger faces opposing forces on its representation. It becomes more important in type (i) districts, but less important in type (ii) districts. Therefore, overall increase in population share would harm the group if most of its population is settled in the type (ii) districts, i.e., if the overlap coefficient is high enough.

An alternative way to think about the result is to notice that the concave relationship between population share and area occupied implies that larger minorities, on average, have higher population density than smaller ones. For minorities which are not dispersed through out the country, there is an "optimal" density that maximizes their presence across districts. If a minority is too dispersed, they become less important everywhere. If they are too concentrated, their importance remains clustered around a few districts only. Our model shows that the larger minorities can suffer from the second problem. Parts 1 and 2 of Proposition 3 refer to two distinct cases when the other minority group is geographically concentrated vs when it is dispersed. In the former case the peak of political representation depends on the size of the two minorities relative to the majority group. In the latter case, the population density of the concentrated minority will be maximized (or the share of the majority group minimized in the constituencies occupied by this group) on the basis of its elasticity of expansion relative to group size. Therefore, the peak doesn't depend on the size of the majority group in this case.

4.4 Data Description

4.4.1 Data Sources

In this section, we briefly describe the various data sources that we have put together for this project. To conserve space, the full description of each dataset is provided in appendix section 4.C.

EPR: The information on political representation and demographic details at ethnic group level comes from Ethnic Power Relations (EPR) core dataset 2014 (Vogt et al. 2015). The EPR dataset provides a measure of political representation in the national government for every ethnic group in a country for all years from 1946–2013. The measure, called the "power rank," can belong to one of six categories signifying the degree of representation. These are, in descending order of power, *monopoly*, *dominant*, *senior partner*, *junior partner*, *powerless*, and *discriminated by the state*. The first two categories refer to cases where a group has substantial representation in the government, for the next two, some representation, and the final two categories refer to cases where the group has no representation. The categorization is created by scholars in the field after taking inputs from over one hundred country experts. It is nonetheless a subjective measure of representation, and therefore, could potentially be biased. We address this concern in different ways, including validating the measure against a more objective measure of representation for a subset of countries. We discuss this in detail in appendix section 4.D.

For our analysis, we coarsen the information and define an indicator of *political inclusion* of a group in the national government, which takes value 1 if the group in a given country and year is neither powerless nor discriminated by the state, as coded by the dataset, and 0 otherwise. We take this indicator to be our main political variable. We say that a group is politically included in the government if the indicator takes value 1 for the group. This variable therefore captures an extensive margin of political representation.

Apart from the power rank measure, EPR also provides annual group-country level data on population shares, settlement patterns and other characteristics of groups.

Nightlight Luminosity: The EPR dataset is complemented with GeoEPR (Wucherpfennig et al. 2011) which consists of GIS maps of settlement areas of a subsample of ethnic groups in the EPR dataset which are geographically concentrated in a region. These maps are overlaid with DMSP-OLS Nighttime Lights Time Series to measure average nightlight luminosity in an ethnic group's settlement area.

Electoral Systems Data: The data for electoral rules come from two sources —the Democratic Electoral Systems (Bormann and Golder 2013) and the IDEA Electoral System Design Database. For any given year, the electoral system in a country is the electoral system used in the most recent election. We restrict our analysis to Majoritarian and Proportional systems.

Polity IV: Polity IV Project allows us to identify periods of autocratic and democratic rule for a country. We classify a country as democracy in a particular year if the position of the chief executive is chosen through competitive elections and include only those observations in the sample (Marshall et al. 2016).[21]

Colonial History: The ICOW Colonial History Dataset 1.0 (Paul 2014) is used to identify the primary colonial ruler and the year of independence for each country that was colonized. The primary colonial ruler is typically the state that ruled the largest area of the colony or ruled it for the longest time. We use this dataset to find the electoral rule followed by the primary colonial ruler in the colony's year of independence for our identification strategy.

4.4.2 Summary Statistics

Appendix table 4.6 reports summary statistics for both the ethnicity level (Panel A) and the country level (Panel B) variables. In our data, 43.87% of country-year observations have MR, whereas 56.13% have PR. The countries with MR are more fractionalized, have greater number of relevant groups, but allow lesser political competition and place fewer constraints on decision making powers of the chief executive compared to PR. However, these differences are not statistically significant at 10% level. On an average, the largest group comprises 73.5% of the politically relevant population and in 84.9% of country-year observations the largest group has an absolute majority in the country (i.e., population share over 50%).

Figure 4.2 shows the distribution of minority population shares in our data. Overall 36.6% of minorities are politically included and 78.4% are geographically concentrated. The ethnicity level characteristics are also not significantly different between countries with MR and PR.

4.5 Empirical Methodology

4.5.1 Main Specification

We use the linear probability model to estimate the effect of group size on political inclusion under MR and PR. In the baseline specification we first check if population share of a group has any relationship with its probability of inclusion in the national executive and whether the relationship is different across the two electoral systems. The following is our preferred specification:

$$\mathbb{P}[I_{ict} = 1] = \delta_{ct} + \beta_1 n_{ict} + \beta_2 n_{ict}^2 + \beta_3 P_{ct} * n_{ict} + \beta_4 P_{ct} * n_{ict}^2 + \gamma X_{ict} + \epsilon_c \tag{4.5.1}$$

where I_{ict} is a dummy indicating whether group i is politically included in country c in year t, δ_{ct} denotes fixed effects at the level of country-year pairs, n_{ict} is the population share of the group, P_{ct} is a dummy indicating whether PR system has been used in the latest national elections in country c in year t; X_{ict} is a vector of ethnicity level controls (which include years of peace, settlement patterns, trans-ethnic kin inclusion/exclusion and fraction of the group associated with the largest language and religion in the group, see Bormann et al. 2017; Cederman et al. 2013 for more details). The error term ϵ_c is clustered at the country level. We include a square term for the population share of the

[21] Our results are robust to using the more conventional definition of democracy based on the polity score.

group to check for non-linearity in the relation. Our sample is restricted only to minority groups in each country.

Given this specification, we compare groups *within* a country-year. Therefore, we only consider countries with 2 or more minorities. The specification controls for a variety of observables and unobservables that vary at the country-year level and may affect the relation we estimate. We argue that two groups of same size across two different countries or in same country but in two different years may wield different political power. This is because a group's access to state power may depend on the number and size composition of all the groups, including the majority, their explicit or implicit political alliances, electoral strategies of political parties, voters' attitudes towards the groups and any political, economic or social contingency that may affect all these factors in complex and unpredictable ways. It may depend on other historical and cultural factors as well, which may depend on time varying characteristics of the country which are often hard to observe. By comparing groups within a country-year, we are able to cut through all these issues which may affect a group's political representation and focus sharply on group specific features only. Our analysis, therefore, avoids any "cross-country" analysis in the sense that the coefficients are not estimated by comparing groups across countries (or by comparing the same group over time).

An alternative, though imperfect, way of estimating the relationship would be to use the panel nature of our data and compare the same minority over time, by exploiting its temporal change in population share and political inclusion status. However, the estimation strategy suffers from a major drawback. There are unobservable political factors in a country, some of which we have listed above, that can change over the years which may affect the relationship we wish to estimate. A panel regression would not be able to absorb such changes. For this reason this is not our preferred empirical specification. We therefore relegate the discussion on the panel analysis to appendix section 4.G.

4.5.2 Identification

Instrumental Variable Strategy

The baseline specification treats the electoral system of a country as exogenous. However, scholars have argued that the choice of electoral system is endogenous to the existing power structure of the country (Boix 1999; Lijphart 1992; Trebbi et al. 2008). In the presence of such concerns our interaction terms in specification (4.5.1) are likely to be misidentified. One potential solution to the issue could have been to focus on the few countries that switch from one electoral system to the other during the sample period. However, such switches themselves might be endogenous as they could be precipitated by the discontent of some groups with the current distribution of power.

Reynolds et al. (2008) argue that many colonies adopted the electoral system of their colonial ruler. We, therefore, look at a subset of countries which had once been colonies. We use the primary colonial ruler's electoral system in the independence year of a colony as an instrument for electoral system of the colony. The exclusion restriction for this specification requires that the colonialists' electoral system did not have a direct *differential* effect on the political power of minorities of different sizes. This would hold even if the electoral system of the colonial ruler is correlated with the power of minority groups on average as long as it is uncorrelated with the power inequality among minorities. For example, one concern with the IV strategy could have been that the British might have had more egalitarian and permissive legal codes in their colonies as compared to the Spanish or Portuguese. This might result in differences in political representation of minorities across MR and PR countries today. However, as long as the liberal legal codes of British colonies gave similar kind of advantages to minorities of all sizes—which is likely to be the case—our identification strategy would remain valid.[22]

For the IV strategy, we keep only those colonies in the sample which democratized not too long after gaining independence from their colonial ruler. Some countries, such as Indonesia and Brazil, became dictatorships after independence and remained so for many decades before democratizing. In such cases the colonial ruler's electoral system matters much less for a country. For example, there are 7 countries which democratized at least 50 years after becoming independent.[23] Only one of them have MR even though all except one were colonized by countries with the MR system. We use two thresholds for sample selection—countries which democratized within 30 and 50 years of independence.[24] We first run the following first stage regressions:

$$P_{ct} * n_{ict} = d_{ct} + a_1 n_{ict} + a_2 n_{ict}^2 + a_3 H_c * n_{ict} + a_4 H_c n_{ict}^2 + \pi X_{ict} + u_c$$

$$P_{ct} * n_{ict}^2 = e_{ct} + b_1 n_{ict} + b_2 n_{ict}^2 + b_3 H_c * n_{ict} + b_4 H_c n_{ict}^2 + \omega X_{ict} + v_c$$

where $H_c = 1$, if colonialist of country c had proportional system in the colony's year of independence. We then get the estimates of β_1–β_4 from specification (4.5.1) in the second stage regression.

Comparing Same Group Across Countries

Sometimes a group is present in more than one country in the same region.[25] Examples include the Kurds who are present in both Turkey and Iran (figure 4.3, panel A), the Basques in France and Spain (panel B) and the San in Botswana and Namibia (panel C). Therefore, we exploit the differences in the sizes of the *same* group across those countries to identify the effect of group size. When the countries have different electoral systems (as in the case of France and Spain), the differential effect of electoral systems could also be estimated by comparing the group across those countries. The idea is that the variation in population shares of the same group across countries within a region comes from the group being unequally divided into multiple national jurisdictions, and therefore, can be

[22] There could be a further threat to the IV strategy if, for example, the colonial rulers with different electoral systems happened to colonize countries having different group size compositions. Appendix table 4.7 reports the results of regressing the indicator that colonialist's electoral system is PR on various population composition measures (fractionalization of minorities, number of minorities etc). We use the population figures of the groups and number of groups for the earliest period in the sample when the country was independent. The coefficients show that colonialist's electoral system is not correlated with the population composition of groups at or near the time of the colony's independence.

[23] These are Bhutan, Brazil, El Salvador, Honduras, Indonesia, Nicaragua and Panama.

[24] There are 18 countries which democratized over 30 years after independence. Of them 10 have PR, though only 2 were colonized by countries with a PR system.

[25] The countries belong to one of five regions: Africa, Asia, Americas, Europe and Oceania.

treated as exogenous.[26] We estimate the following model:

$$\mathbb{P}[I_{ict} = 1] = \delta_{irt} + \theta P_{ct} + \beta_1 n_{ict} + \beta_2 n_{ict}^2 + \beta_3 P_{ct} * n_{ict} + \beta_4 P_{ct} * n_{ict}^2 + \gamma X_{ict} + \epsilon_{ic} \tag{4.5.2}$$

where δ_{irt} denotes ethnicity-region-year fixed effects, error term ϵ_{ic} is double clustered at group and country level to adjust standard errors against potential auto-correlation within group and country. The coefficient θ is the intercept of the relationship and β_1–β_4 are the coefficients of interest, as before.

4.6 Results

4.6.1 Verifying the Main Assumption of the Model

Before we discuss the empirical results, we verify one key parameter restriction of the model that we need for our main theoretical result to hold. Proposition 3 requires the minority groups' settlement areas to be inelastically related to their population shares. Moreover, Bettencourt (2013) argues that the value of α should be 0.67. We run the following specification to test this:

$$\ln S_{ict} = \alpha \ln n_{ict} + \gamma X_{ict} + \delta_{ct} + \epsilon_c \tag{4.6.1}$$

where S_{ict} is the settlement area of a group i in country c in year t and n_{ict} is the population share that group. α measures the elasticity of settlement area with respect to population share of a group, and therefore, the estimate $\hat{\alpha}$ is a direct estimate of the parameter α in the theoretical model. The EPR dataset provides information about the settlement area of groups which are geographically concentrated. This allows us to estimate equation (4.6.1). We report the results in table 4.1. Column (1) reports the main estimate of α to be 0.625. It is statistically significant at 1% level and significantly lower than one, also at 1% level. Further, the coefficient is statistically indistinguishable from 0.67, confirming the prediction of Bettencourt (2013). Moreover, we also estimate this parameter in two sub-samples—where the minority groups' population shares are smaller than 0.25 (column (2)) and smaller than 0.1 (column (3)). Both estimates are close to each other and are similar to the main estimate. This shows that the elasticity of settlement area with respect to population share of a group is indeed stable, further confirming our model's assumption. It is important to mention here that this result is in line with papers that also verify the theoretical claim of Bettencourt (2013) in various contexts (Ortman et al. 2014, 2015; Ortman et al. 2016; Cesaretti et al. 2016).

4.6.2 Baseline Results

Political Representation: Table 4.2 column (2) shows the results for political representation from our baseline specification. The coefficient of population share is positive and significant at 1%

[26] The method is similar to Dimico (2017) who shows in the context of Africa that the partition of an ethnicity in two countries adversely affects their political representation when the resulting groups are small. We show that the effect of how an ethnic group is divided in two democracies on the group's political representation and economic development depends on the electoral system.

level and that of population share-squared is negative and significant at 5% level. Their magnitudes imply that for MR countries there is an inverted-U shaped relation between population share of a group and its probability of political inclusion. Probability of political inclusion peaks at population share of 0.26. The interactions of population share and its square with the PR dummy are statistically significant (at 5% level) and have opposite signs. F-tests for the hypotheses $\beta_1 + \beta_3 = 0$ and $\beta_2 + \beta_4 = 0$ give p-values of 0.33 and 0.96, respectively. This indicates that there is no relation between population share and political inclusion under PR. Column (1) reports the results with a weaker specification—having country and year fixed effects separately. We observe that the coefficients remain similar in magnitude. Also, the PR dummy has a positive and marginally significant coefficient. This suggests that very small minority groups presumably enjoy higher political representation under PR compared to MR.

The aforementioned result is unlikely to be driven by a systematic bias in coding of the power rank variable. Since we compare the groups within a country-year pair, we effectively control for the researcher(s) who were responsible for the power ranking of these groups. For the result to be driven by biased coding, it must be the case that the sets of researchers coding the MR and PR countries are systematically biased against different subsets of minorities having different population shares. Further, the coefficients of ethnicity level controls as reported in appendix table 4.8 are of the expected sign. The coefficient of peace years is positive and statistically significant at 1% level. An additional decade without any conflict incidence experienced by an ethnicity is associated with 4.15% more likelihood of its political inclusion. The coefficient of transethnic-kin exclusion dummy is positive and significant. This might be due to the fact that politically excluded ethnic groups sometimes migrate to countries where they might get political representation. An indicator of an ethnic group's cohesiveness is the fraction of its members associated with the largest language spoken by the group. Groups that are linguistically more cohesive find it easier to organize themselves and put forth their demands. Therefore, they are more likely to be politically included. This is supported by the result that a 10 percentage points increase in fraction of group members associated with the largest language for the group is related with a 2.10% increase in likelihood of political inclusion for the group.

Robustness: Appendix table 4.9 reports the results of various robustness exercises we carry out to ensure that the result is not driven by either sample selection or the chosen specification. For example, the result remains same if we restrict the sample to country-year observations having polity score of at least 7 (out of the maximum polity score of 10), or parliamentary democracies, or having a group with absolute majority in population. It is also robust to using relative population share of group (i.e., ratio of population share of a group to the population share of the majority group) as explanatory variable, or using the power rank of groups as the dependent variable. We discuss the robustness checks in detail in appendix section 4.E. Further, our model generates some additional predictions. For example, our model shows that the inverted-U shape under MR is driven by groups that are geographically concentrated; under PR, on the other hand, geographic concentration is not important. Moreover, Proposition 3 predicts that the "optimum" share of minorities under MR would be smaller the larger is the size of the majority group. We empirically validate these two predictions in appendix tables 4.11 and 4.12, respectively. We discuss the results in appendix section 4.F.

Nightlights: Our model predicts that per capita allocation of public resources to ethnic groups follows the same pattern as their political representation. We test this using the same specification. We use nightlight intensity as a proxy for public resource allocation towards groups which are settled in a geographically well demarcated region within a country. Electricity in most countries is publicly provided and is an essential public good for any region within a country. Therefore, nightlight luminosity acts as a direct proxy for government allocation of resources—in the form of electricity access—in an area. In fact, this has been shown to be the case in Senegal and Mali (Min et al. 2013), and Vietnam (Min and Gaba 2014). We use (logarithm of) nightlight intensity per unit area as our dependent variable to test specification (4.5.1).[27] Michalopoulos and Papaioannou (2013a) use the same measure to proxy for economic development of ethnic groups in the African continent. They further use micro-data from Afrobarometer surveys to confirm that the measure is a good proxy for various public goods such as "access to electrification, presence of a sewage system, access to piped water, and education" within settlement areas of ethnic groups.[28] Given the volume of evidence coming from a wide range of countries, we feel confident that our measure is a good proxy for allocation of government resources, and more generally, for the level of development of an ethnic group.[29]

The use of nightlight luminosity imposes two restrictions in the data—it is available only from 1992 onwards and can be used only for groups which have a well-demarcated and contiguous settlement area as specified by the EPR dataset. Table 4.2 column (4) reports the results. It shows that the result for political inclusion is replicated with nightlight as outcome variable. The estimated population share with peak nightlight intensity in MR countries is 0.21 which is similar to what we estimated for inclusion. Moreover, we see in column (3) that the dummy for PR system again has a positive (though imprecisely estimated) coefficient, consistent with the column (1) result. This suggests that the patterns of political inclusion indeed have implications for resources allocated

[27] We add 0.01 as a constant to nightlight intensity per area measure before taking the logarithm.

[28] Nightlight luminosity is also a well-documented and widely used proxy for the level of economic development in a geographic region. For a discussion about using nightlight luminosity as a measure of economic activity see Doll (2008) and Henderson et al. (2012). The papers using nightlight data as a proxy for economic development in various contexts are too numerous to cite here. The papers that use nightlight data to answer political economy related questions include among others, Michalopoulos and Papaioannou (2013b,a), Prakash et al. (2019), and Baskaran et al. (2015). Alesina et al. (2016) show that nightlights can indeed be used as a proxy to measure ethnic inequality globally using a sample of 173 countries, i.e. both in developing as well as developed country context. Similarly, Hodler and Raschky (2014) use nightlights data from 126 countries to show evidence of regional favoritism.

[29] Henderson et al. (2012) have raised important issues with using nightlight luminosity as proxy for economic outcome. Many of these concerns are however addressed in our empirical analysis. Firstly, Henderson et al. (2012) point out that the nightlight data is captured using different satellite sensors and therefore, the luminosity data is not comparable across the years. This is addressed in our analysis since we use country-year fixed effects. Henderson et al. (2012) similarly use year fixed effects to address the issue. The other concern is that nightlight data is not captured in countries with high latitudes during summer time. Thus, Henderson et al. (2012) remove the regions above the Arctic Circle from their analysis. All the countries in the Arctic Circle, barring Russia, are not in our sample as well, since they have only one minority group. The third concern with nightlight data is the phenomenon of blurring, i.e., tendency of light to be captured beyond the exact source (due to coarse light sensors). However this is more of an issue in using nightlight data in smaller areas. The extent of blurring ranges from 4.5 km to 9 km depending on the radiance of the light source (Abrahams et al., 2018). Since the median area of ethnic groups in our sample is about 23,500 square km, we do not think this to be a major source of measurement error.

towards groups.

4.6.3 IV Results

The IV results are reported in table 4.3. Panel B (column 1) of the table shows that the presence of PR in a country is 47% more likely in countries that democratized within 30 years of independence if the electoral system of its primary colonial ruler was also PR in the colony's year of independence. The coefficient is statistically significant at 1% level.

Panel A reports the second stage results using political inclusion dummy and log of nightlight intensity per unit area as dependent variables. The first two columns report the results for countries which democratized within 30 years of being independent and the next two columns report the same with a 50 year threshold. In all four columns we find the same pattern. For MR countries we get a strong inverted-U shaped relationship. The peak is achieved at population shares 0.22 and 0.24 for political inclusion, and 0.22 and 0.26 for nightlight intensity, for the 30 and 50 year threshold regressions respectively. Moreover, the relationships are indeed flat for PR, as both the tests $\beta_1 + \beta_3 = 0$ and $\beta_2 + \beta_4 = 0$ fail to reject the null hypothesis in all four columns. The coefficients for political inclusion across columns (1) and (3) are similar in magnitudes and comparable to the coefficients estimated in the baseline specification (table 4.2, column (2)). Importantly, the Kleibergen-Paap rk LM statistic for the first stage regressions are high in all specifications, alleviating concerns related to under-identification. The F statistics for the each first stage regression are also very large in magnitude.

Finally, for the sake of transparency, we report in appendix table 4.13 the IV strategy results when we do not put any restrictions on the sample. Both political inclusion (column 1) and nightlight (column 2) regressions show an inverted-U shaped relationship for MR countries. We get a flat relationship for political inclusion in PR countries. For the nightlight regressions, however, the coefficients have imprecise estimates. Importantly, the regressions don't pass the under-identification tests as the Kleibergen-Paap rk LM statistics are low. This suggests that our sample restrictions are indeed useful in making our specification stronger.

4.6.4 Results for Comparison of Same Group across Countries

Table 4.4 reports the results for specification (4.5.2) having political inclusion (column 1) and log nightlight intensity per area (column 2) as dependent variables. The within group comparison reaffirms the inverted-U shaped effect of population share on political representation under MR and no relation under PR. The coefficients reported in column (1) are a bit larger compared to those estimated in the IV regression (table 4.3). The peak of political representation under MR is achieved at population shares of 0.20 in this identification strategy, which is similar to what we estimated before.

We also find that nightlight intensity indeed follows same pattern with the peak achieved at population share of 0.19 for MR countries. The coefficients estimated however have large standard errors, presumably due to small sample size. Also, coefficient of the PR dummy is positive and significant for political inclusion, suggesting that very small minorities get better represented in PR relative to MR. We plot the marginal effect of population share on political inclusion for the two methods in the figure 4.4. It suggests that mid-sized groups enjoy higher level of political inclusion under MR.

4.6.5 Using Road Network Data as Alternate Outcome

Road construction is widely believed to be an important activity of governments and often constitutes an important item in their annual budget. Burgess et al. (2015), for example, use road building in Kenya to show how democracy affects allocation of public resources across ethnic groups. We construct cross-sectional data on road construction across ethnicity-country pairs and use it as a proxy for allocation of public resources to test the robustness of our result. We use geo-spatial data from the Global Roads Inventory Project (GRIP) (see Meijer et al. 2018) who provide GIS locations of various kinds of roads across several countries as they exist currently.[30] The dataset therefore is cross-sectional in nature. The dataset distinguishes among roads of five categories: national highways, primary roads, secondary roads, tertiary roads and local roads. We overlay the road network map on the maps of the settlement areas of ethnic groups and national boundaries to get the section of road network that lies within an ethnicity-country pair. We then aggregate the road length of the first four types of roads falling within the area of each pair. We don't consider local roads in our analysis because they are unlikely to be allocated by the national government. We use total road length per square kilometer of the settlement area of an ethnicity as our measure of public resource allocation towards that group.

We match the road data with our main dataset for the latest year (i.e., for 2013). We then run the cross-sectional version of specification 4.5.1 for the year 2013 (with only country fixed effects). Since we have a cross section of groups for a subset of democracies, the number of observations in the regression is small. So we run the specification with and without group level controls. The results are reported in table 4.5, columns (1) and (2). The coefficients indicate that the pattern mirrors our main result—an inverted-U shaped relationship in MR, and no relationship in PR. However, when we include group level controls, the coefficients expectedly become noisier.

We then run our IV strategy specification on the sample of erstwhile colonies for the year 2013. We use the 30 year democracy lag as the sample restriction. The results without and with group level controls are reported in columns (3) and (4), respectively. The F-stats of the first stage regressions are above the commonly used threshold of 10. The second stage estimates show that the pattern is replicated even in the small cross-sectional sample. The optimal group size in MR in the baseline specification is around 19% which is similar to the one estimated in the second identification strategy. The optimal group size for the IV specification is smaller at 15%. However, given that different (small) sub-samples of countries are used in some of the regressions, getting different estimates of the optimal group size is not unlikely.

4.7 Concluding Remarks

This chapter examines how electoral systems influence the relation between population share of a minority group and its access to power in the national government. First, we develop a model with

[30]The dataset is freely available in the project website: http://www.globio.info/download-grip-dataset.

multiple minority groups which predicts that in PR countries, population share of a minority has no effect on its political representation, while in MR countries the relation is inverted-U shaped. We then compile a large panel dataset at the ethnicity level for 87 countries for the post-war period to test the predictions of our model. The empirical analysis remarkably exhibits the same pattern for both—political representation as well as per capita resource allocation. Our results imply that electoral systems can have stark effect on power (and welfare) inequality. We get that under PR, group size inequality does not translate into inequality in the political representation of minorities and consequently, the inequality in material well-being would also be minimal. On the other hand, power inequality among minorities in MR countries depends on the size distribution of the groups. It is the mid-sized minority groups that enjoy maximum access to power in MR, while the small and large minorities enjoy similar levels of representation. Our work further highlights the importance of settlement patterns of groups in determining their representation in the government under MR. We, however, take settlement patterns as exogenous. One interesting line of future enquiry can be to consider the settlement patterns of mobile minorities to be endogenous and explore if electoral system influences the settlement decisions of such minorities. We wish to take up this issue in our future work.

Figures and Tables

Figure 4.1: Electoral system distribution in 2013

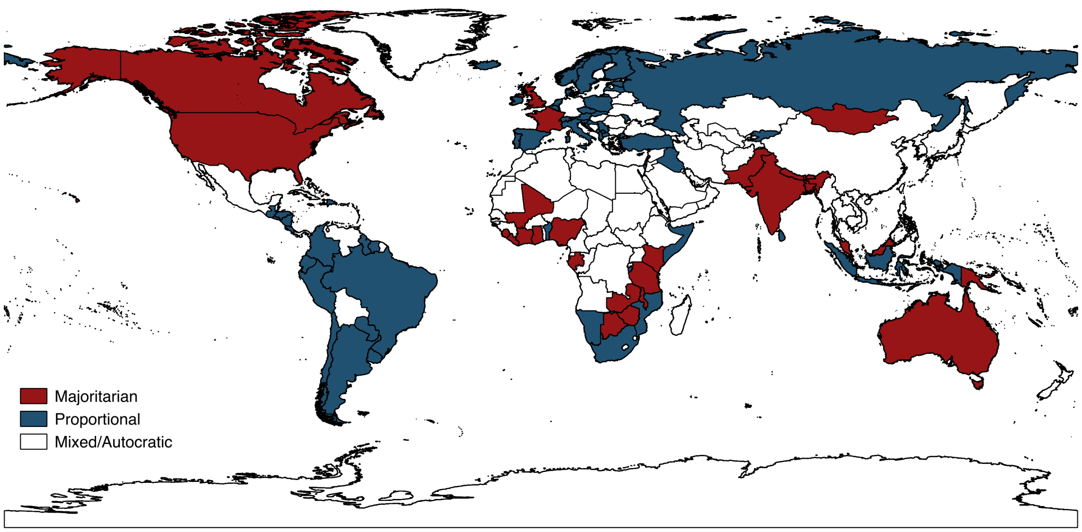

Figure 4.2: Distribution of minority population shares

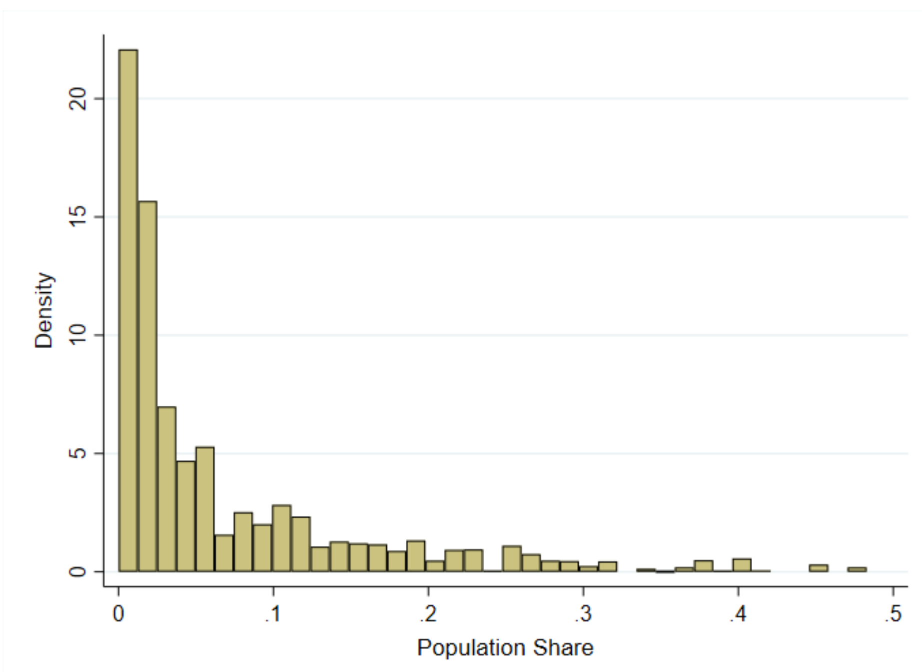

Figure 4.3: Examples of groups with settlement areas across national boundaries

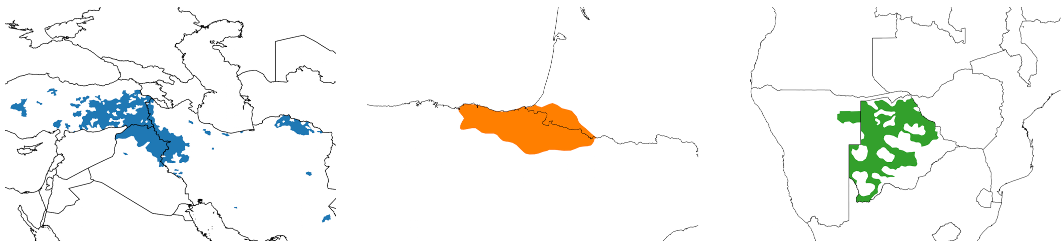

(a) Kurds (b) Basques (c) San

Panel (a) shows *Kurds* in Iran, Iraq and Turkey; panel (b) shows *Basques* in France and Spain; and panel (c) shows *San* in Botswana and Namibia.

Figure 4.4: Marginal Effect of Group Size on Political Inclusion

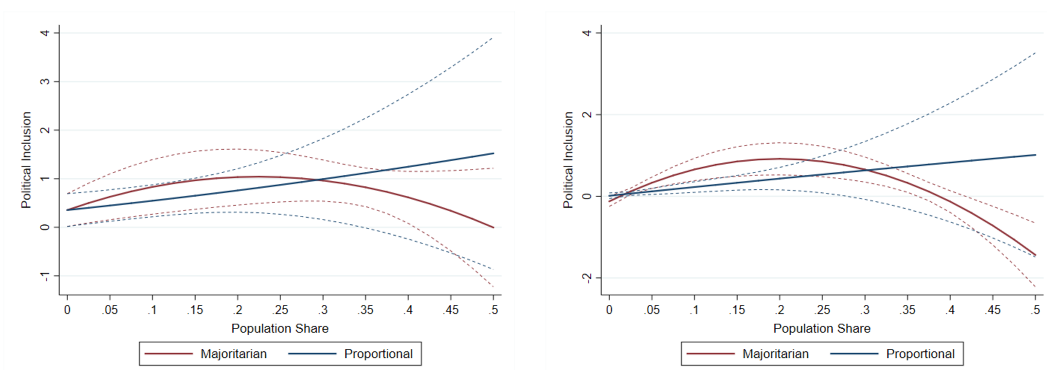

(a) IV Estimation (b) Same Group Across Countries

Table 4.1: Settlement Area Expands Inelastically and $\alpha = 0.67$

	ln(Settlement area)		
	(1)	(2)	(3)
α: ln(Population share)	0.625***	0.661***	0.668***
	(0.122)	(0.134)	(0.124)
$H_0 : \alpha \geq 1$ (one tailed p-value)	0.001	0.007	0.005
$H_0 : \alpha = 0.67$ (p-value)	0.736	0.968	0.992
Mean dependent	10.140	10.006	9.783
Observations	6,665	5,946	4,357
R-squared	0.792	0.779	0.742
Ethnicity-year controls	YES	YES	YES
Country-year FE	YES	YES	YES

Notes: Column (1) reports results for all concentrated minorities. Minority population share in column (2) ≤ 0.25 and that in column(3) ≤ 0.10. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.2: Inverted-U shaped relation under MR and no relation under PR

	Political inclusion		Nightlight	
	(1)	(2)	(3)	(4)
β_1: Population share	4.405***	4.825***	10.77***	11.02***
	(1.239)	(1.227)	(3.741)	(3.840)
β_2: Population share - squared	-7.884**	-9.276**	-24.48**	-24.49**
	(3.883)	(3.955)	(10.67)	(11.30)
β_3: Proportional x Population share	-3.011*	-3.661**	-9.729*	-10.11
	(1.687)	(1.721)	(5.840)	(6.103)
β_4: Proportional x Population share - squared	6.903	9.106*	24.03	24.29
	(5.159)	(5.313)	(15.27)	(16.17)
Proportional	0.247*		0.328	
	(0.144)		(0.382)	
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.23	0.33	0.84	0.86
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.77	0.96	0.97	0.99
Predicted optimal size	0.279	0.260	0.22	0.21
Mean dep. var.	0.366	0.366	-0.23	-0.26
Observations	9,294	8,706	3,756	3,469
R-squared	0.652	0.687	0.821	0.816
Ethnicity-year controls	YES	YES	YES	YES
Country-year controls	YES	NO	YES	NO
Country FE	YES	NO	YES	NO
Year FE	YES	NO	YES	NO
Country-year FE	NO	YES	NO	YES

Notes: The dependent variable for columns (1) and (2)—political inclusion—is a dummy variable that takes value one if the group in a country in a given year is neither powerless nor discriminated by the state. Sample restricted only to minorities in each country. Column (1) includes 438 ethno-country groups in 102 countries, and column (2) includes 421 ethno-country groups in 87 countries during 1946–2013. The dependent variable for columns (3) and (4) is logarithm of nightlight luminosity per unit area in the settlement area of groups. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.3: IV strategy replicates main results

Panel A: Second stage				
	Lag < 30 years		Lag < 50 years	
	Political inclusion	Nightlight	Political inclusion	Nightlight
	(1)	(2)	(3)	(4)
β_1: Population share	6.142***	26.87**	5.832***	21.95**
	(1.999)	(10.23)	(1.541)	(8.936)
β_2: Population share - squared	-13.72**	-68.13**	-12.17**	-46.23**
	(6.695)	(25.77)	(4.629)	(21.89)
β_3: Proportional x Population share	-4.332*	-41.69**	-4.049**	-36.53**
	(2.421)	(17.67)	(1.962)	(16.37)
β_4: Proportional x Population share - squared	14.77*	92.90*	13.35*	69.35
	(8.698)	(49.71)	(7.158)	(44.86)
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.102	0.174	0.102	0.174
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.859	0.488	0.844	0.504
Predicted optimal size	0.223	0.224	0.239	0.263
Observations	4,361	1,720	4,632	1,926
R-squared	0.700	0.773	0.711	0.766
Ethnicity-year controls	YES	YES	YES	YES
Country-year FE	YES	YES	YES	YES
Kleibergen-Paap rk LM stat	5.06	3.10	5.12	3.12
Cragg-Donald Wald F stat	172.18	43.01	188.47	50.96
F stat (Proportional*Population share)	119.51	40.47	260.33	125.57
F stat (Proportional*Population share - squared)	312.74	72.36	919.01	516.56

Panel B: First Stage (Country level)				
	Proportional		Proportional	
Colonialist proportional	0.470***		0.522***	
	(0.162)		(0.143)	
Observations	508		818	
R-squared	0.653		0.561	
Region-year FE	YES		YES	

Notes: The first two columns in Panel A and the first column in Panel B include countries which were once colonies and democratized within 30 years of gaining independence ("Lag < 30 years"). The last two columns in Panel A and the second column in Panel B has countries with independence-democracy lag less than 50 years ("Lag < 50 years"). Standard errors clustered at the country level are reported in parentheses. Sample restricted only to minorities in each country. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.4: Comparing same group across countries replicate main results

	Political inclusion	ln(Nightlight per area)
	(1)	(2)
β_1: Population share	10.44***	58.54
	(2.424)	(35.90)
β_2: Population share - squared	-26.13***	-156.4
	(6.091)	(92.29)
β_3: Proportional x Population share	-8.269***	-58.72
	(2.686)	(35.96)
β_4: Proportional x Population share - squared	25.79**	147.7
	(10.96)	(96.88)
Proportional	0.138**	0.991
	(0.0513)	(1.352)
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.17	0.99
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.96	0.83
Predicted optimal size	0.200	0.187
Observations	1,370	417
R-squared	0.836	0.887
Group-year controls	YES	YES
Country-year controls	YES	YES
Group-region-year FE	YES	YES

Notes: Column (1) compares 21 minorities in 40 countries and column (2) compares 12 minorities in 30 countries. Standard errors double clustered at the group and country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.5: Road Construction and Electoral Systems

	Road length per unit area			
	Baseline		IV	
	(1)	(2)	(3)	(4)
β_1: Population share	1.492*	1.236	3.456**	2.430
	(0.762)	(0.863)	(1.540)	(1.690)
β_2: Population share - squared	-4.071*	-3.229	-11.35**	-8.246
	(2.267)	(2.396)	(5.248)	(5.499)
β_3: Proportional x Population share	-1.500	-1.180	-3.597**	-2.563
	(0.981)	(0.981)	(1.699)	(2.035)
β_4: Proportional x Population share - squared	4.480*	3.571	11.22*	7.424
	(2.657)	(2.797)	(5.662)	(6.932)
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.99	0.89	0.79	0.85
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.77	0.76	0.94	0.76
Predicted optimal size	0.18	0.19	0.15	0.15
F stat (Proportional x Population share)	–	–	16.36	11.65
F stat (Proportional x Population share - squared)	–	–	16.23	12.37
Observations	227	227	105	105
R-squared	0.750	0.777	0.754	0.768
Group Controls	NO	YES	NO	YES
Country FE	YES	YES	YES	YES

Notes: Dependent variable is kilometers of non-local roads in the settlement area of a group per square kilometer of the area in 2013. The data is cross-sectional. Columns (1) and (3) have no group level controls while columns (2) and (4) have the same set of group level control as the previous regressions. The baseline regressions (columns (1) and (2)) have 54 countries and IV regressions (columns (3) and (4)) have 24 countries. Standard errors clustered at the country level are reported in parentheses. Sample restricted only to minorities in each country. *** p<0.01, ** p<0.05, * p<0.1.

Appendix

Appendix Sections:

A. Trend in Electoral System and Minority Representation

B. Electoral Systems and Government Formation

C. Data Compilation

D. Subjectivity in Political Representation Measurement

E. Robustness of the Baseline Result

F. Validating Additional Comparative Static Results of the Model

G. Panel Analysis

H. Figures and Tables

I. List of Countries

J. Proofs of Propositions

4.A Trend in Electoral System and Minority Representation

From 1950s to the 1970s, a larger fraction of countries followed MR. However, the past few decades have seen a trend towards the adoption of PR. This can be observed in figure 4.5, where we plot the number of country-year observations by electoral system for each decade from 1950s through 2000s.

This is largely driven by adoption of the PR system by the new democracies in Latin America, Africa, and Mediterranean, Central and Eastern Europe in the 1970s and 1980s.[31] Several countries have also changed their existing systems to more proportional electoral formulas. For example, Japan and New Zealand switched from MR and held their first general elections under a mixed system in 1996. Another case in point is Russia, which changed its mixed electoral system and employed PR for the 2007 legislative election.[32]

Figure 4.6 plots the proportion of minority groups in democracies in each power status category during 1946–2013. As it shows, there has been a gradual decline in state administered discrimination against minorities over the years. However, the share of groups in the powerless category has correspondingly increased. There is also no clear pattern in the proportion of groups in power sharing arrangements with other groups (i.e., junior and senior partner) and of those who rule virtually alone (dominant and monopoly groups). While this proportion was increasing during 1990s, it has remained virtually stable afterwards and was in fact declining during some of the earlier decades.

4.B Electoral Systems and Government Formation

The electoral system pertains to the legislature while we look at representation of minorities in the national government (or the executive). Our analysis includes countries with both parliamentary and presidential systems. The fact that in parliamentary systems representation in the legislature has a bearing on the executive is understandable, since the executive is selected from the legislature itself. The case for non-parliamentary systems, however, is less obvious and needs an explanation. The first thing to note is that a significant proportion of such countries have a semi-presidential system where the cabinet is either formed by the legislature, or faces the threat of no confidence vote from the legislature, or both. France, Poland, Sri Lanka, Peru, and Senegal are examples of such countries. The difference in the strategic incentives of parties across MR and PR, therefore, would be relevant in such countries. Among the countries with a presidential system, some still need formal approval of the legislature to form the cabinet. In fact, even in countries where the president can appoint and dismiss the cabinet freely without any legislative approval, there is a high correlation between seat share of parties in the legislature and seat share in the cabinet.[33] Therefore, the electoral strategies

[31] The possible reasons for adoption of PR system by these countries are discussed in Farrell (2011).

[32] Other examples include Argentina, Sri Lanka and Moldova which switched directly from MR to PR for their parliamentary elections held in 1963, 1989 and 1998 respectively. There have also been a few instances of changes in the opposite direction—i.e. towards less proportionality. These include Venezuela, Madagascar and Bulgaria where PR was replaced in favor of mixed system in 1993, 1998 and 2009 legislative elections respectively.

[33] Silva (2016), for example, shows for Brazil that even though the party of the president gets an advantage in the cabinet, the cabinet portfolio share increases by 0.9% for every percentage point increase in legislative

of the parties to form the government seem to be similar to the strategies for legislative elections even in presidential systems. This is understandable given that legislative and executive elections are often held simultaneously and consequently, political parties have consistent platforms (in terms of group representation) for both elections.[34]

4.C Data Compilation

EPR: Our primary source of data is the Ethnic Power Relations (EPR) core dataset 2014 (Vogt et al. 2015). The dataset contains various characteristics of well-identified groups ("ethnicities") within countries for about 155 countries across the world at an annual level for the period 1946–2013. All sovereign states with a total population of at least 500,000 in 1990 are included in the dataset. The dataset defines a group "as any subjectively experienced sense of commonality based on the belief in common ancestry and shared culture."[35] The dataset is concerned with groups that are politically relevant; a group is politically relevant if at least one political organization or a political party has at least once claimed to represent it at the national level or the group has been explicitly discriminated against by the state during any time in the period 1946–2013. This aligns with our interest as well. As long as there is some marker of identity which is salient in the society and is also politically meaningful, we should consider them in our analysis.

The demarcation of groups is intuitive and meaningful. India, a large and diverse country, for example, has 20 groups—the second highest in our sample.[36] These groups are based on religion (Kashmiri Muslims and Other Muslims), caste (SC/STs, OBCs) as well as language or ethnicity (Non SC/ST Bengalis, Non SC/ST Marathis, Mizo, Naga etc). United States, on the other hand, has 6 groups—Whites, African Americans, American Indians, Asian Americans, Arab Americans and Latinos. All the countries in our sample, barring India and Russia, have number of groups ranging from 2 to 14, with the average number of groups in the total sample being 4.6. We list in Appendix 4.H the samples of countries used in our empirical exercises along with the respective number of minority groups and number of years in the sample, i.e., having a democratic regime.[37]

The dataset provides annual group-country level data on population shares, settlement patterns, trans-border ethnic kinship, as well as religious and linguistic affiliations for the period 1946–2013.

seat share even for non-presidential parties.

[34] All our empirical results remain the same if we do not consider countries with the presidential system where the president doesn't require any approval from the legislature for cabinet formation.

[35] Cederman et al. (2010) further point out that in different countries different "markers may be used to indicate such shared ancestry and culture: common language, similar phenotypical features, adherence to the same faith, and so on." Further, in some societies there may be multiple dimensions of identity along which such "sense of commonality" may be experienced.

[36] Russia with 39 groups has the highest number.

[37] It is important to note that politically relevant ethnic divisions in a country may change over time. New cleavages may emerge increasing the number of groups or some existing group may cease to be politically relevant as well. In case of South Africa, for example, racial divisions primarily between the Whites and the Blacks marked the political climate during the Apartheid era, while divisions within the black South African population along ethno-linguistic lines (such as between *Xhosa* and *Zulu*) has become more prominent in the subsequent period. The dataset recognize this fact. The number of groups in some countries, therefore, changes a little bit over the years. The number of minorities specified in Appendix 4.H is the maximum number in the sample.

However, most importantly for us, it also codes a group's access to national executive. A group's access to absolute power in the national government is coded based on whether the group rules alone (power status = *monopoly*, *dominant*), shares power with other groups (power status = *senior partner*, *junior partner*) or is excluded from executive power (power status = *powerless*, *discriminated by the state*). We rank these six categories in a separate variable called "power rank"; they range from 6 to 1 in decreasing order of power (i.e., from *monopoly* to *discriminated*).[38] The power ranking of the groups is evidently a subjective exercise. The researchers, however, are fairly transparent in the method that they follow in assigning power ranks. They look at the degree and nature of presence of members of a particular group in the most important political positions in the national government in determining its power rank. The details about group demarcation of the countries and the justification of the power rankings of each group is fully described on the official website of the EPR project: https://growup.ethz.ch.

The EPR dataset also provides information about the settlement patterns of the groups. Specifically, it categorizes the groups as being *dispersed*, i.e., those who do not inhabit any particular region within a country and, *concentrated*, i.e., settled in a particular region of the country which is easily distinguishable on a map. For concentrated groups, it further gives information about the country's land area (km^2) that they occupy.[39]

The EPR dataset was created by scholars who work on group based conflict. The first version of the dataset was created as part of a research project between scholars at ETH Zurich and University of California, Los Angeles (UCLA), which was then updated and released by Vogt et al. (2015). The information about the attributes of groups, including their power status is coded by the researchers based on inputs from about one hundred country experts. This dataset has certain advantages for this chapter over other existing datasets about political outcomes of groups. Some of the prominent datasets used by scholars of conflict are the Minorities At Risk (MAR) dataset, the All Minorities at Risk (A-MAR) dataset and the dataset used by Fearon (2003). Though most of these datasets give information about group sizes, none of these provide any detail about the settlement patterns of the groups. This is critical for us since we demonstrate that the pattern observed in our data is driven by groups which are geographically concentrated. Also, the EPR dataset provides information about the power status of all groups; this is in contrast to the MAR dataset which systematically excludes the groups that are in the government.

Night Light: The EPR dataset is complemented with GeoEPR dataset (Wucherpfennig et al. 2011) which consists of GIS maps of the settlement areas of a subsample of ethnic groups in the EPR dataset which are geographically concentrated in a region. These maps are overlaid with DMSP-OLS Nighttime Lights Time Series to measure average nightlight luminosity in an ethnic group's settlement area. We use this to create a proxy of public resource allocation at the level of ethnic groups.

Electoral Systems Data: The data for electoral rules used for national elections come from merging two datasets. The first of these is the Democratic Electoral Systems (DES) data compiled

[38] There is an additional categorization in the data, known as *self-exclusion*. This applies to groups which have declared independence from the central state. They constitute only 0.7% of our sample and we do not consider them in our analysis.

[39] The GIS shape file of their area of settlement is also provided on the EPR website.

by Bormann and Golder (2013). It contains details about electoral systems used for about 1200 national elections for the period 1946–2011. We complement this with a second source of data—the IDEA Electoral System Design Database, which gives us information about the electoral systems for some additional countries. The classification into broad electoral systems is based on the DES dataset. For any given year, the electoral system in a country is the electoral system used in the most recent election. We restrict our analysis to Majoritarian and Proportional systems.

Polity IV: Polity IV Project allows us to identify periods of autocratic and democratic rule in a country. We define democracy as country-year pairs where the position of the chief executive is chosen through competitive elections and include only those observations in the sample. We prefer this definition over the standard categorization based on the Polity IV score because we wish to look at all the countries that have competitive elections and have one of the two electoral systems of our interest. However, there are other aspects of a regime such as extent of checks and balances on the executive that affect the Polity IV score as well, which are of less relevance to our specific analysis. We, of course, show robustness of our result using a different definition of democracy based on the polity score.

Colonial History: The ICOW Colonial History Dataset 1.0 compiled by Paul (2014) recognizes the primary colonial ruler and the year of independence for each country that was colonized. The dataset marks the start of colonial rule with the establishment of first permanent outpost or settlement by a state-sponsored company. *Primary colonial ruler* is defined as the one "most responsible for shaping the development" of the colony. This is typically the state that ruled the largest area of the colony or ruled it for the longest time. To obtain the electoral systems of the colonial rulers we use the data on electoral systems provided in The Handbook of Electoral System Choice (HESC) (Colomer 2004). The HESC provides information about electoral systems of democracies since 1800. We use this to find the electoral rule followed by the primary colonial ruler in the colony's year of independence. We use this information for our identification.

4.D Subjectivity in Political Representation Measurement

One concern with the power rank variable is that it is a subjective measure and therefore, could potentially be biased. We address the concern in three different ways. First, we use the dataset created by Francois et al. (2015) for 15 countries in Africa. It contains share of cabinet positions held by ethnic groups within each country for every year during 1960–2004. This could be considered to be a more objective measure of representation. We, therefore, match the ethnic groups from that dataset to the EPR data. We are able to match 90% of groups. We then correlate the power ranking (from the EPR data) with cabinet shares. Figure 4.7 graphically shows this correlation using a binned scatter plot. We observe that the two variables are highly positively correlated ($r = 0.56$) and also, the nature of the relationship is linear. This suggests that the power rank variable is indeed informative about the real power held by groups within governments. Further, we notice that out of the six categories of power rank, the last two categories (*powerless* and *discriminated by the* *state*) refer to cases where the group is either has no representation in the government or is actively discriminated by the state. We consider these to be sharply contrasting cases where the problem of subjectivity is presumably minimal. We, therefore, coarsen the power ranking to create an indicator of *political inclusion* which takes value one if a group is neither powerless nor discriminated by the state, and zero otherwise. The political inclusion indicator therefore measures the extensive margin of representation of a group, i.e., whether the group has *any* representation in the government or not. We take the indicator of political inclusion as our main political variable.[40] Moreover, as we argue in section 4.6 that given our empirical specification, our results are unlikely to be driven by biased measurement. Finally, we provide evidence that, consistent with model's prediction, the pattern is replicated with measures of material well-being of groups.

4.E Robustness of the Baseline Result

In table 4.9 we run specification (4.5.1) on various sub-samples of the data. Columns (1) and (2) show results for two time periods 1946–1979 and 1980–2013, respectively. The broad patterns depicted in our baseline specification continue to hold over time, though the coefficients are larger for the earlier period, indicating a more pronounced inverted-U relationship for MR countries in the first half of the post–war period. Column (3) shows the cross-sectional result for the latest year in our sample, i.e., for 2013. The coefficients here are quite similar to the column (1) coefficients. In column (4) we replace the main explanatory variable by the relative population share, i.e., the ratio of population share to the population share of the majority group in the country-year observation. Columns (5) restricts the sample to countries with an absolute majority and column (6) restricts the sample to parliamentary democracies only. In column (7) we only include election years in the sample and column (8) includes countries which are full democracies according to the Polity IV dataset (i.e., countries with a polity score of at least 7). Finally, in column (9) we use the power rank variable as our dependent variable. The variable takes value 1 through 6 with 1 being discriminated, 2 powerless and so on. In all specifications we fail to reject that $\beta_1 + \beta_3 = 0$ and $\beta_2 + \beta_4 = 0$. Therefore, in all specifications we get that there is no relation between population share and political inclusion in a PR system. Similarly, in all specifications we get that the relationship is inverted-U shaped in the MR system, though the coefficient β_2 is noisily estimated in some specifications. The consistency of the pattern across various sub-samples of the data strongly suggests that the result is a general phenomenon observed across democracies.

Further, we rerun the baseline regressions for political inclusion and log nightlight per area by reweighing the observations by the (inverse of) the number of minority groups in the country-year observations. We do this to ensure that our results are not driven by countries with large number of groups. We report the results in appendix table 4.14. Coefficients in both columns suggest that our result remains the same with this specification.

[40] We of course show robustness of our results to using the main power rank variable as the outcome variable.

4.F Validating Additional Comparative Static Results of the Model

The primary aim of the model is to justify the empirical pattern established in the Section 4.6 of the chapter. The model, however, generates some additional predictions regarding the exact nature of the relationship between group size and access to political power. It is, therefore, important to test if these additional comparative static results hold in order to verify if the proposed model is indeed valid. We now turn to that discussion in the following paragraphs.

Proposition 3 states that we should observe the inverted-U shaped relationship between group size and power status under the MR system only for groups which are geographically concentrated. Also, a group's geographic concentration should not matter for the result of the PR system. We verify this by running the following specification for the samples of MR and PR country-year observations separately:

$$Y_{ict} = \delta_{ct} + \eta_1 n_{ict} + \eta_2 n_{ict}^2 + \eta_3 C_{ict} * n_{ict} + \eta_4 C_{ict} * n_{ict}^2 + \gamma X_{ict} + \epsilon_c \tag{4.F.1}$$

where C_{ict} is a dummy indicating whether the group i is geographically concentrated in country c in year t. Proposition 3 implies that for the sample of MR countries, η_1 and η_2 should be zero and we should have $\eta_3 > 0$ and $\eta_4 < 0$. For the set of PR countries all the coefficients η_1–η_4 should be zero. Table 4.11 reports the results and the predictions are verified. Column (1) reproduces the main result, and columns (2) and (3) provides the estimates of η_1–η_4 for MR and PR countries, respectively. As is evident, for the MR countries the relationship is only true for geographically concentrated groups. For PR countries, none of the coefficients are statistically significant.

Proposition 3 further specifies that under the MR system, the peak political representation is achieved when the population share of the group equals $\frac{1-n_3}{2}$ when the group is geographically concentrated, where n_3 is the population share of the majority group. Therefore, for larger values of the majority group's share, the peak is achieved at lower values of the minority group's size. We test this prediction by running specification (4.5.1) on various sub-samples of the data where we vary the size of the majority group. The results are reported in table 4.12. Columns (1)–(3) report the results for sub-samples where the majority group's population share is larger than 0.3, 0.5, and 0.7, respectively. The table also reports the population shares at which the peak inclusion is achieved. We see that the population share at which the peak inclusion is achieved declines as we move to countries with larger majority groups.

4.G Panel Analysis

An alternative, though imperfect, way of estimating our main relationship would be to compare the same minority over time, by exploiting its temporal change in population share and political inclusion status. However, such a specification would not be able to account for all the time varying unobservable factors that we mention in the main chapter. In this section we discuss the panel regression. The panel specification could be written as:

$$\mathbb{P}[I_{ict} = 1] = \delta_{ic} + \phi_t + \beta_1 n_{ict} + \beta_2 n_{ict}^2 + \beta_3 P_{ct} * n_{ict} + \beta_4 P_{ct} * n_{ict}^2 + \gamma_1 X_{1ict} + \gamma_2 X_{2ct} + \epsilon_{ict} \tag{4.G.1}$$

where δ_{ic} is a group-country fixed effect, ϕ_t is a year fixed effect, X_{1ict} is a vector of ethnicity characteristics and X_{2ct} is a vector of country characteristics. However, there are two important drawbacks in this estimation strategy. Importantly, there are unobservable political factors in the country, some of which we have listed above, that can change over the years which may affect the likelihood of political inclusion of the group. The direction of this effect is uncertain as it would depend on the nature of the change in the political climate of the country. Therefore, the coefficients $\beta_1 - \beta_4$ are likely to have noisier estimates. Also, the size composition of other groups, including the majority group would change over time which may affect the relationship as well. To partly account for the last factor we run specification (4.G.1) using relative population share in place of population share. The relative population share is defined as the ratio of the population share of the group in country-year observation to the population share of the majority group in the same country-year pair.

We report the results in table 4.15. Columns (1) and (4) report the results for our two main dependent variables using the full sample. We see that the coefficients β_3 and β_4 for column (1) do not have the expected signs and all the coefficients are noisily estimated. The coefficients for the nightlight regression (column 4) do have the expected signs. The magnitudes of β_1 and β_2 imply that group size has an inverted-U shaped relationship with nightlight intensity in MR countries, though the standard errors of the coefficients are high. The coefficients β_3 and β_4 have the opposite signs, implying that the relationship is flatter for PR. Since annual variations in population share would not immediately translate to changes in representation or material welfare, we keep in sample every third (columns (2) and (5)) and fifth (columns (3) and (6)) year that a group is present in the data. We see that the all coefficients for political inclusion have the expected signs in column (3), though the magnitude of β_3 is smaller than β_1. The coefficients for the nightlight regressions in column (5) and (6) maintain their correct signs. The coefficients for the interaction terms are, however, smaller in magnitudes. The panel results indicate that the relationship observed for minorities within a country-year becomes less precise when we follow the same minority over the years. This is expected given our discussion above.

Additional Figures and Tables

Figure 4.5: Electoral systems by decade

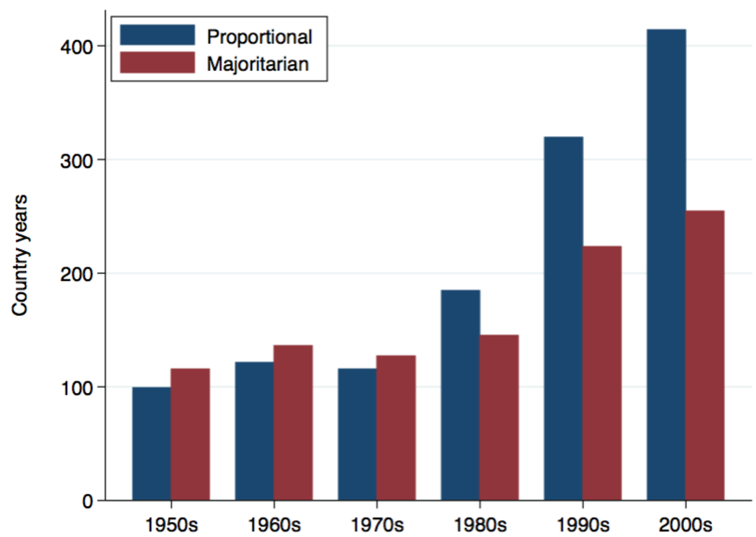

Figure 4.6: Minority power status over time

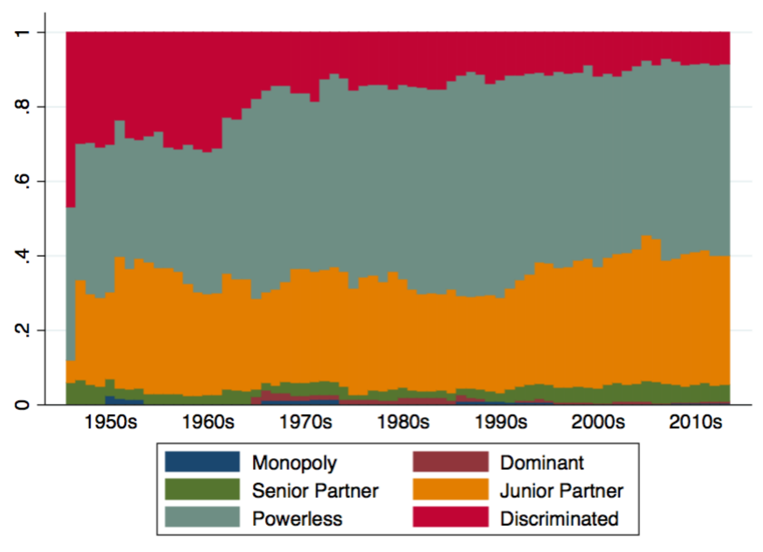

Figure 4.7: Comparing Power Rank with Cabinet Shares

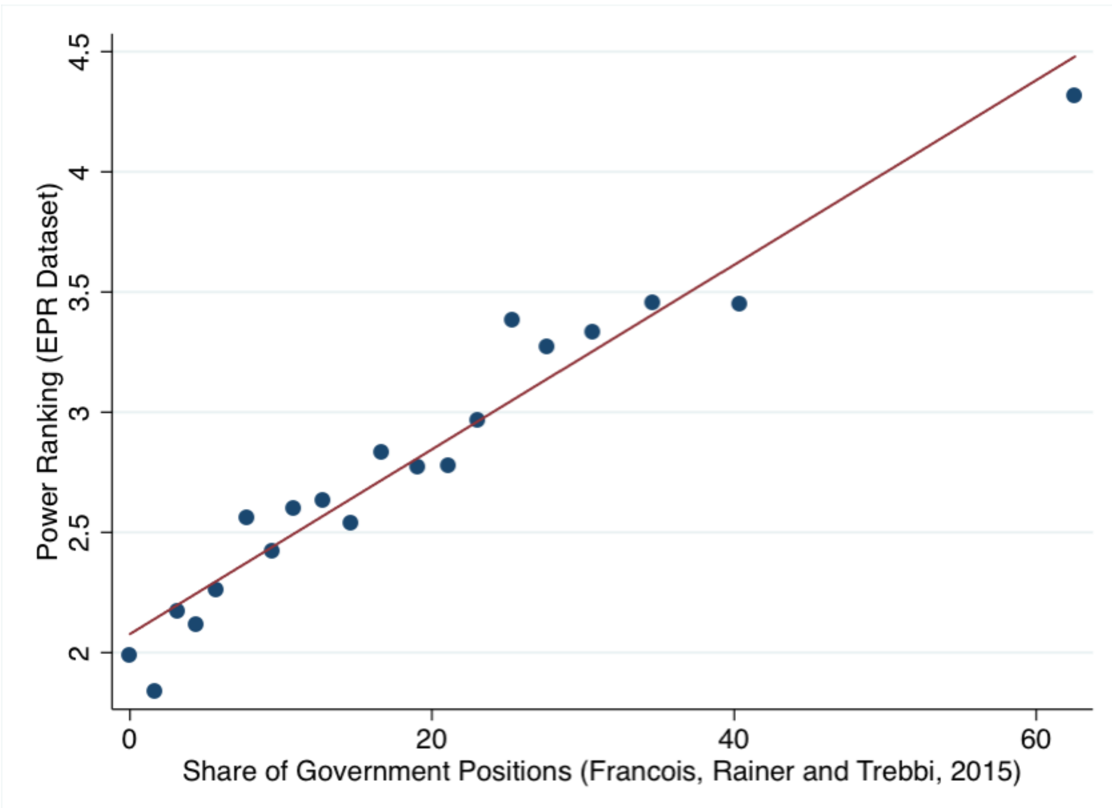

Table 4.6: Descriptive statistics

	All data	Majoritarian system	Proportional system	Difference
	(1)	(2)	(3)	(4)
Panel A: Ethnicity level				
Political inclusion	0.366	0.444	0.275	0.169
	(0.482)	(0.497)	(0.446)	(0.112)
Power rank	2.294	2.391	2.180	0.211
	(0.793)	(0.770)	(0.806)	(0.188)
Population share	0.074	0.070	0.079	-0.009
	(0.099)	(0.090)	(0.108)	(0.024)
Years peace	31.418	29.223	34.029	-4.806
	(20.285)	(19.178)	(21.236)	(4.162)
Aggregate settlement	0.002	0.001	0.004	-0.003
	(0.046)	(0.031)	(0.059)	(0.005)
Statewide settlement	0.032	0.026	0.040	-0.014
	(0.176)	(0.158)	(0.195)	(0.045)
Regional and urban settlement	0.381	0.416	0.339	0.077
	(0.486)	(0.493)	(0.474)	(0.114)
Urban settlement	0.087	0.103	0.067	0.036
	(0.282)	(0.305)	(0.251)	(0.061)
Regional settlement	0.369	0.325	0.421	-0.096
	(0.482)	(0.468)	(0.494)	(0.106)
Dispersed settlement	0.109	0.118	0.098	0.020
	(0.312)	(0.323)	(0.298)	(0.074)
Migrant settlement	0.020	0.011	0.031	-0.020
	(0.140)	(0.103)	(0.174)	(0.028)
Transethnic-kin inclusion	0.417	0.402	0.435	-0.033
	(0.493)	(0.490)	(0.496)	(0.103)
Transethnic-kin exclusion	0.521	0.460	0.594	-0.135
	(0.500)	(0.498)	(0.491)	(0.105)
Fraction largest religion	0.719	0.750	0.682	0.069
	(0.209)	(0.222)	(0.186)	(0.053)
Fraction largest language	0.879	0.889	0.867	0.023
	(0.223)	(0.214)	(0.232)	(0.045)
Observations	9,294	5,049	4,245	9,294
Panel B: Country level				
Ethnic fractionalization	2.433	2.885	2.079	0.806
	(1.989)	(2.201)	(1.723)	(0.494)
Number of relevant groups	4.596	5.470	3.913	1.557
	(3.772)	(4.221)	(3.221)	(0.944)
Largest group size	0.735	0.687	0.772	-0.086
	(0.219)	(0.238)	(0.195)	(0.054)
Absolute majority	0.849	0.753	0.923	-0.170*
	(0.359)	(0.432)	(0.266)	(0.086)
Competitiveness of participation	3.989	3.873	4.079	-0.207
	(1.056)	(1.252)	(0.962)	(0.232)
Constraints chief executive	6.121	5.978	6.233	-0.256
	(1.291)	(1.370)	(1.497)	(0.270)
Observations	2,601	1,141	1,460	2,601

Notes: The data is at the ethnicity-country-year level for 438 ethno-country groups in Panel A and country-year level for 102 countries in Panel B for the period 1946–2013. Standard deviation in parenthesis in columns (1), (2) and (3). Standard errors clustered at the country level in parenthesis in the last column. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.7: Group size distribution is not correlated with colonialist's system

	Colonialist Proportional			
	(1)	(2)	(3)	(4)
Minority Fractionalization	0.0220			
	(0.0215)			
Number of relevant minorities		0.00772		
		(0.0125)		
Largest group size			0.0789	
			(0.187)	
Absolute majority				0.0607
				(0.0908)
Observations	95	95	95	95
R-squared	0.220	0.214	0.212	0.215
Region-year FE	YES	YES	YES	YES

Notes: Country level data for 95 countries. Earliest year for which group size data is available is taken for each country. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.8: Inverted-U shaped relation under MR and no relation under PR

	Political inclusion			
	(1)	(2)	(3)	(4)
β_1: Population share	4.405***	4.825***	10.77***	11.02***
	(1.239)	(1.227)	(3.741)	(3.840)
β_2: Population share - squared	-7.884**	-9.276**	-24.48**	-24.49**
	(3.883)	(3.955)	(10.67)	(11.30)
β_3: Proportional*Population share	-3.011*	-3.661**	-9.729*	-10.11
	(1.687)	(1.721)	(5.840)	(6.103)
β_4: Proportional*Population share - squared	6.903	9.106*	24.03	24.29
	(5.159)	(5.313)	(15.27)	(16.17)
Proportional	0.247*		0.328	
	(0.144)		(0.382)	
Years peace	0.00409***	0.00415***	0.00378	0.00361
	(0.00135)	(0.00130)	(0.00360)	(0.00357)
Aggregate settlement	0.549***	0.541***	-3.023***	-3.003***
	(0.110)	(0.114)	(0.432)	(0.430)
Statewide settlement	0.294	0.139		
	(0.375)	(0.352)		
Regional and urban settlement	0.174**	0.170**	0.0107	0.0264
	(0.0784)	(0.0789)	(0.419)	(0.410)
Urban settlement	0.0180	0.00905		
	(0.0663)	(0.0650)		
Regional settlement	-0.0105	-0.00942	-1.029**	-1.004**
	(0.0488)	(0.0483)	(0.399)	(0.392)
Migrant settlement	-0.140	-0.150		
	(0.195)	(0.195)		
Transethnic-kin inclusion	0.00421	0.000118	0.0502	0.0355
	(0.0446)	(0.0477)	(0.152)	(0.159)
Transethnic-kin exclusion	0.0897**	0.103***	0.114	0.134
	(0.0347)	(0.0348)	(0.132)	(0.142)
Fraction largest religion	-0.125	-0.108	-0.0609	-0.0201
	(0.109)	(0.105)	(0.537)	(0.530)
Fraction largest language	0.193**	0.210***	1.437***	1.441***
	(0.0737)	(0.0748)	(0.387)	(0.388)
Ethnic fractionalization	0.0203		-0.0680	
	(0.0251)		(0.0806)	
Number of relevant groups	0.0123		0.0134	
	(0.0197)		(0.0776)	
Competitiveness of participation	0.00848		0.0201	
	(0.0166)		(0.0310)	
Constraints chief executive	-0.0169		0.0107	
	(0.0104)		(0.0298)	
Observations	9,294	8,706	3,756	3,469
R-squared	0.652	0.687	0.821	0.816
Country FE	YES	NO	YES	NO
Year FE	YES	NO	YES	NO
Country-year FE	NO	YES	NO	YES

Notes: Data is at the level of ethnicity-country-year. The dependent variable for columns (1) and (2)–political inclusion—is a dummy variable that takes value one if the group in a country in a given year is neither powerless nor discriminated by the state. The sample for column (1) includes 438 ethno-country groups in 102 countries, and for column (2) includes 421 ethno-country groups in 87 countries the period 1946–2013. The dependent variable for columns (3) and (4) is logarithm of nightlight luminosity per unit area of groups which have well-demarcated settlement areas. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.9: Main results are robust

	Political inclusion								Power rank
	1946-1979	1980-2013	2013						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
β_1: Population share	6.259***	4.470***	6.333***		5.130***	6.816***	3.732***	6.903***	5.543***
	(2.016)	(1.027)	(0.928)		(1.814)	(1.902)	(1.367)	(1.979)	(1.784)
β_2: Population share - squared	-14.14*	-8.141***	-12.31***		-7.732	-17.18**	-5.714	-17.21**	-7.910
	(7.334)	(2.919)	(2.643)		(5.362)	(6.820)	(4.473)	(7.771)	(5.773)
β_3: Proportional x Population share	-6.674**	-2.745*	-4.838***		-4.385*	-7.080***	-3.949**	-6.577***	-4.972*
	(2.453)	(1.526)	(1.586)		(2.220)	(2.362)	(1.915)	(2.429)	(2.797)
β_4: Proportional x Population share - squared	17.33**	6.322	10.79**		9.334	22.30**	9.625	20.29**	11.81
	(8.306)	(4.472)	(4.404)		(6.619)	(8.043)	(5.958)	(8.821)	(9.201)
(β_1): Relative population share				2.381***					
				(0.402)					
(β_2): Relative population share-squared				-2.108***					
				(0.459)					
(β_3): Proportional x relative population share				-1.574***					
				(0.582)					
(β_4): Proportional x relative population share-squared				1.815***					
				(0.675)					
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.717	0.161	0.259	0.087	0.559	0.862	0.853	0.770	0.808
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.277	0.611	0.663	0.600	0.640	0.277	0.240	0.298	0.605
Predicted optimal size	0.221	0.275	0.257	–	0.332	0.198	0.327	0.201	0.350
Mean dependent	0.332	0.378	0.403	0.366	0.214	0.428	0.320	0.363	2.276
Observations	2,295	6,411	303	8,706	5,750	4,854	1,773	5,832	8,706
R-squared	0.669	0.704	0.735	0.693	0.675	0.681	0.702	0.728	0.675
Ethnicity-year Controls	YES	YES	YES	YES	YES	YES	YES	YES	YES
Country FE	NO	NO	YES	NO	NO	NO	NO	NO	NO
Country-year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Data is at the level of ethnicity-country-year. Only minorities are part of the sample. Political inclusion is a dummy variable that takes value one if the group in a country in a given year is neither powerless nor discriminated by the state. Columns (1) and (2) have sample for the periods 1946–1979 and 1980–2013, respectively. Column (3) runs the specification for the year 2013 only. Column (4) uses relative population share as the main explanatory variable. Relative population share is the ratio of population share of the group and the population share of the largest group in the country-year. Column (5) restricts the sample only to countries where the largest group is absolute majority. Column (6) restricts the sample to parliamentary democracies. Column (7) restricts the sample to only election years. Column (8) restricts the sample only to full democracies i.e. countries with a polity score ≥ 7. Column (9) uses power rank of a group as the dependent variable. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.10: Results not driven by number of parties

	Political inclusion	
	(1)	(2)
β_1: Population share	4.923***	0.852
	(1.462)	(2.320)
β_2: Population share - squared	-10.25**	-2.847
	(4.583)	(5.407)
enpp x Population share	-0.133	0.145
	(0.0954)	(0.402)
enpp x Population share - squared	0.539	0.0457
	(0.377)	(0.904)
Mean inclusion	0.460	0.212
Observations	3,932	2,950
R-squared	0.677	0.795
Ethnicity-year controls	YES	YES
Country-year FE	YES	YES

Notes: Data is at the level of ethnicity-country-year. Only minorities are part of the sample. Column (1) uses only MR countries and column (2) uses only PR countries. $enpp = \frac{1}{s_i^2}$, where s_i is the share legislative seats won by party i. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.11: The pattern in MR is explained by geographical concentration

		Political inclusion	
	(1)	(2)	(3)
Population share	4.825***	1.910	3.324
	(1.227)	(1.609)	(3.122)
Population share - squared	-9.276**	-1.864	-4.437
	(3.955)	(5.917)	(6.917)
Proportional x Population share	-3.661**		
	(1.721)		
Proportional x Population share - squared	9.106*		
	(5.313)		
Concentrated x population share		4.811***	-0.987
		(1.610)	(3.290)
Concentrated x population share - squared		-11.67**	1.054
		(5.589)	(7.651)
Mean inclusion	0.366	0.447	0.265
Observations	8,706	4,830	3,876
R-squared	0.687	0.648	0.734
Ethnicity-year controls	YES	YES	YES
Country-year FE	YES	YES	YES

Notes: Data is at the level of ethnicity-country-year. Only minorities are part of the sample. Column (1) shows the baseline result of column (4) in table 4.2. Column (2) uses only MR countries and column (3) uses only PR countries. Concentrated is a dummy variable that takes value one if the group has a well-demarcated settlement area in a country. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.12: Optimal minority size smaller in countries with larger majority

	Political inclusion		
	(1)	(2)	(3)
β_1: Population share	3.741***	5.130***	7.531***
	(1.297)	(1.814)	(2.159)
β_2: Population share - squared	-5.365	-7.732	-17.93***
	(3.650)	(5.362)	(5.977)
β_3: Proportional x Population share	-2.607	-4.385*	-7.838***
	(1.787)	(2.220)	(2.553)
β_4: Proportional x Population share - squared	5.324	9.334	21.95***
	(5.160)	(6.619)	(7.421)
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.377	0.559	0.857
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.991	0.640	0.540
Predicted optimal size	0.349	0.332	0.210
Mean inclusion	0.286	0.214	0.156
Observations	6,917	5,750	3,871
R-squared	0.685	0.675	0.732
Ethnicity-year controls	YES	YES	YES
Country-year FE	YES	YES	YES

Notes: Data is at the level of ethnicity-country-year. Largest group size in column (1) ≥ 0.3, in column (2) ≥ 0.5, and in column (3) ≥ 0.7. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.13: IV results: Full sample

Panel A: Second stage	Political inclusion	ln(Nightlight per area)
	(1)	(2)
β_1: Population share	5.823***	5.307
	(1.660)	(8.759)
β_2: Population share - squared	-11.79**	-8.388
	(4.994)	(20.48)
β_3: Proportional x Population share	-6.262**	19.23
	(2.482)	(16.39)
β_4: Proportional x Population share - squared	18.88*	-73.32
	(9.990)	(47.51)
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.76	0.02
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.37	0.03
Predicted optimal size	0.247	0.316
Observations	5,047	2,226
R-squared	0.702	0.765
Ethnicity-year controls	YES	YES
Country-year FE	YES	YES
Kleibergen-Paap rk LM stat	2.42	1.89
Cragg-Donald Wald F stat	432.12	183.47
F stat (Proportional*Population share)	193.93	106.45
F stat (Proportional*Population share - squared)	543.95	325.80

Panel B: Country level	Proportional
Colonialist proportional	0.463***
	(0.118)
Mean dependent	.450
Observations	1,309
R-squared	0.388
Region-year FE	YES

Notes: Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.14: Weighting Replicates Main Results

	Political Inclusion	ln(Nightlight per area)
	(1)	(2)
β_1: Population share	3.756***	9.300**
	(1.143)	(3.741)
β_2: Population share - squared	-5.087	-18.75**
	(3.161)	(8.889)
β_3: Proportional x Population share	-3.474**	-13.42*
	(1.584)	(7.062)
β_4: Proportional x Population share - squared	7.032	31.17*
	(4.717)	(17.39)
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.80	0.51
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.55	0.41
Predicted optimal size	0.369	0.248
Observations	8,706	3,469
R-squared	0.737	0.863
Country-year FE	YES	YES
Ethnicity-year controls	YES	YES

Notes: Data is at the level of ethnicity-country-year. Only minorities are part of the sample. All the observations are weighted by the inverse of the number of relevant minorities used in each regression in the given country-year. Standard errors clustered at the country level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4.15: Panel Analysis Produces Similar Patterns

	Political Inclusion			ln(Nightlight per area)		
	(1)	(2)	(3)	(4)	(5)	(6)
(β_1): Relative population share	1.547	1.513	2.127	9.835	31.62**	18.96
	(1.749)	(1.290)	(1.741)	(9.883)	(15.07)	(12.96)
(β_2): Relative population share-squared	-1.020	-1.068	-1.487	-6.674	-21.80**	-11.14
	(1.271)	(0.963)	(1.252)	(6.865)	(10.15)	(9.354)
(β_3): Proportional x relative population share	0.420	0.0847	-0.199	-3.760**	-4.417***	-4.998**
	(0.925)	(0.725)	(0.762)	(1.650)	(1.527)	(2.091)
(β_4): Proportional x relative population share-squared	-0.207	0.367	0.933	6.624	9.091**	9.362
	(0.825)	(0.651)	(0.781)	(4.671)	(3.856)	(5.922)
$H_0 : \beta_1 + \beta_3 = 0$ (p-value)	0.19	0.19	0.28	0.52	0.06	0.27
$H_0 : \beta_2 + \beta_4 = 0$ (p-value)	0.17	0.37	0.66	0.99	0.13	0.85
Observations	9,289	2,979	1,695	3,748	1,194	648
R-squared	0.918	0.921	0.930	0.990	0.992	0.993
Ethnicity-country FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes: Data is at the level of ethnicity-country-year. Only minorities are part of the sample. Standard errors clustered at the country level are reported in parentheses. In columns (2) and (5), we keep every third year in which a group is present in the data, and in columns (3) and (6) we keep every fifth year in which a group is present in the data. *** p<0.01, ** p<0.05, * p<0.1.

4.H List of Countries

S.no.	Country	Years	Minorities	Baseline	IV Strategy
1.	Albania	6	2	✓	
2.	Australia	17	2	✓	✓
3.	Bangladesh	21	3	✓	✓
4.	Belarus	1	2	✓	✓
5.	Belgium	41	2	✓	
6.	Benin	23	3	✓	✓
7.	Bhutan	6	2	✓	
8.	Bolivia	15	3	✓	
9.	Botswana	48	9	✓	✓
10.	Brazil	36	2	✓	
11.	Bulgaria	18	3	✓	
12.	Cambodia	4	4	✓	✓
13.	Canada	65	2	✓	✓
14.	Central African Republic	10	3	✓	✓
15.	Chile	49	2	✓	
16.	Colombia	41	2	✓	
17.	Congo	5	4	✓	✓
18.	Costa Rica	66	2	✓	✓
19.	Cote d'Ivoire	3	4	✓	✓
20.	Croatia	14	5	✓	✓
21.	Czechoslovakia	3	3	✓	✓
22.	Ecuador	44	3	✓	
23.	Estonia	22	3	✓	
24.	Ethiopia	10	8	✓	
25.	France	61	3	✓	
26.	Gabon	5	3	✓	✓
27.	Ghana	15	4	✓	✓
28.	Greece	51	3	✓	
29.	Guatemala	18	3	✓	✓
30.	Guinea-Bissau	10	2	✓	✓
31.	Guyana	17	2	✓	
32.	Honduras	32	2	✓	
33.	India	63	19	✓	✓
34.	Indonesia	15	11	✓	
35.	Iran	4	10	✓	
36.	Iraq	4	2	✓	
37.	Israel	47	4	✓	
38.	Italy	49	5	✓	
39.	Japan	24	3	✓	

40.	Kenya	12	7	✓	✓
41.	Kosovo	4	5	✓	✓
42.	Kyrgyzstan	8	3	✓	✓
43.	Laos	2	5	✓	✓
44.	Latvia	21	3	✓	
45.	Lebanon	37	10	✓	
46.	Liberia	14	5	✓	
47.	Macedonia	16	4	✓	✓
48.	Malawi	20	2	✓	✓
49.	Malaysia	15	4	✓	✓
50.	Mali	21	2	✓	✓
51.	Mauritania	1	2	✓	✓
52.	Mauritius	38	6	✓	✓
53.	Moldova	20	3	✓	✓
54.	Montenegro	8	5	✓	✓
55.	Mozambique	15	2	✓	✓
56.	Myanmar	11	10	✓	✓
57.	Namibia	15	11	✓	✓
58.	Nepal	19	4	✓	
59.	New Zealand	6	2	✓	✓
60.	Nicaragua	24	3	✓	
61.	Nigeria	22	5	✓	✓
62.	Pakistan	17	7	✓	✓
63.	Panama	13	4	✓	
64.	Peru	44	3	✓	
65.	Philippines	36	3	✓	
66.	Poland	23	4	✓	✓
67.	Romania	18	3	✓	
68.	Russia	7	38	✓	
69.	Serbia	7	6	✓	✓
70.	Sierra Leone	20	3	✓	✓
71.	Singapore	17	3	✓	✓
72.	Slovenia	22	7	✓	✓
73.	South Africa	20	13	✓	✓
74.	Spain	36	4	✓	
75.	Sri Lanka	62	3	✓	✓
76.	Sudan	7	12	✓	✓
77.	Switzerland	67	2	✓	
78.	Tanzania	19	4	✓	✓
79.	Thailand	23	3	✓	
80.	Turkey	45	2	✓	
81.	Uganda	5	5	✓	✓
82.	Ukraine	11	4	✓	✓

83.	United Kingdom	68	6	✓	
84.	United States	68	5	✓	
85.	Yugoslavia	7	5	✓	
86.	Zambia	18	6	✓	✓
87.	Zimbabwe	5	2	✓	✓

4.I Proofs of Propositions

4.I.1 Proof of Proposition 1

Consider the case of party A. Vote share of party A among members of group j is given by:

$$\pi_{A,j} = Pr[U(f_j^A) > U(f_j^B) + \delta + \sigma_{i,j}]$$

Assuming that $\psi \geq \phi_j$ for all j, we get:

$$\pi_{A,j} = \frac{1}{2} + \phi_j[U(f_j^A) - U(f_j^B) - \delta]$$

Party A will win elections if more than half the population votes for it. Probability of winning for party A is given by:

$$p_A = Pr[\frac{\sum\limits_{j=1}^{3} n_j \pi_{A,j}}{\sum\limits_{j=1}^{3} n_j} > \frac{1}{2}]$$

This can simply be written as:

$$p_A = \frac{1}{2} + \frac{\psi \sum\limits_{j=1}^{3} \phi_j n_j (U(f_j^A) - U(f_j^B))}{\sum\limits_{j=1}^{3} \phi_j n_j}$$

Thus, party A solves:

$$\max_{f_j^A \geq 0} p_A = \frac{1}{2} + \frac{\psi \sum\limits_{j=1}^{3} \phi_j n_j (U(f_j^A) - U(f_j^B))}{\sum\limits_{j=1}^{3} \phi_j n_j}$$

$$s.t. \quad \sum_{j=1}^{3} n_j f_j^A \leq S$$

Solving the above optimization problem gives the equilibrium condition in 1.

4.I.2 Proof of Proposition 2

In a K district majoritarian election, probability of winning for party A in constituency k, as can be seen from the result under proportional electoral system, is given by:

$$p_A^k = \frac{1}{2} + \frac{\psi \sum_{j=1}^{3} \phi_j n_j^k (U(f_j^A) - U(f_j^B))}{\sum_{j=1}^{3} \phi_j n_j^k}$$

Party A will win the election if it wins more than half the votes in more than half the districts. If both parties win in equal number of districts, then the winner will be chosen randomly. Party A solves the following optimization problem under majoritarian elections:

$$\max_{f_j^A \geq 0} p_A \quad s.t. \quad \sum_{j=1}^{3} n_j f_j^A \leq S$$

Since the parties are symmetric, in equilibrium, $p_A^k = \frac{1}{2}$ for all districts. Thus, given a district k, we denote the probability of winning in any other given district, with a slight abuse of notation, as p_A^{-k}. When K=2, Probability of winning can be written as:

$$p_A = p_A^k p_A^{-k} + \frac{1}{2}[p_A^k(1 - p_A^{-k}) + p_A^{-k}(1 - p_A^k)]$$

This can be simplified to:

$$= \frac{1}{2} p_A^k + \frac{1}{4}$$

And when $K > 2$, probability of winning is:

$$\begin{aligned} p_A = &\sum_{i=\lfloor K/2 \rfloor}^{K-1} \binom{K-1}{i} p_A^k (p_A^{-k})^i (1 - p_A^{-k})^{K-1-i} \\ &+ \sum_{i=\lfloor K/2 \rfloor + 1}^{K-1} \binom{K-1}{i} (1 - p_A^k)(p_A^{-k})^i (1 - p_A^{-k})^{K-1-i} \\ &+ \frac{1}{2}[\frac{1 + (-1)^K}{2}][\binom{K-1}{\lfloor K/2 \rfloor - 1} p_A^k (p_A^{-k})^{(K/2)-1} (1 - p_A^{-k})^{K/2} \\ &+ \binom{K-1}{\lfloor K/2 \rfloor} (p_A^{-k})^{K/2} (1 - p_A^{-k})^{(K/2)-1} (1 - p_A^k)] \end{aligned}$$

This can be simplified to:

$$\begin{aligned}p_A =& \frac{1}{2^{K-1}}[\binom{K-1}{\lfloor K/2 \rfloor} p_A^k + \sum_{i=\lfloor K/2 \rfloor+1}^{K-1} \binom{K-1}{i}] \\ &+ \frac{1}{2^K}[\frac{1+(-1)^K}{2}][(\binom{K-1}{\lfloor K/2 \rfloor - 1} - \binom{K-1}{\lfloor K/2 \rfloor}) p_A^k + \binom{K-1}{\lfloor K/2 \rfloor}]\end{aligned}$$

Using this, we calculate:

$$\frac{dp_A}{dp_A^k} = C(K) = (\frac{1+(-1)^{K-1}}{2}) \binom{K-1}{\lfloor K/2 \rfloor} \frac{1}{2^{K-1}} + (\frac{1+(-1)^K}{2}) \binom{K}{\lfloor K/2 \rfloor} \frac{1}{2^K}$$

For the first order condition to the optimization problem, we need to calculate:

$$\frac{dp_A}{df_j^A} = \sum_{k=1}^{K} \frac{dp_A}{dp_A^k} \frac{dp_A^k}{df_j^A}$$

Substituting the expression for dp_A/dp_A^k, we can write this as:

$$\frac{dp_A}{df_j^A} = C(K) \sum_{k=1}^{K} \frac{dp_A^k}{df_j^A}$$

We can now easily solve the optimization problem to give the equilibrium condition given in 2. Consider the case where all groups are equally responsive to electoral promises i.e. $\phi_j = \phi$ for all j. Since $\sum_{j=1}^{3} n_j^k = 1$ for all k and $\sum_{k=1}^{K} n_j^k/n_j = K$ for all j, 2 can be simplified to:

$$U'(f_i^*) = U'(f_l^*) \quad \forall i, l$$

Now, consider the case where $n_j^k = n_j$ for all k. In this case, 2 can be simplified to:

$$\phi_i U'(f_i^*) = \phi_l U'(f_l^*) \quad \forall i, l$$

Both the above special cases indicate that when groups are evenly distributed across districts or when all groups are equally responsive to electoral promises, majoritarian elections give the same equilibrium political representation and per capita transfers as the proportional representation system.

4.I.3 Proof of Proposition 3

(a) When group 2 is concentrated, we have four types of constituencies based on the identity of groups residing in them: (1) Only group 1 and 3 reside (2) Only group 2 and 3 reside (3) Group 1, 2 and 3 all reside (4) Only group 3 resides. Densities D^m of constituency type m are:

$$D^1 = n_1^{1-\alpha} + n_3 \qquad D^2 = n_2^{1-\alpha} + n_3 \qquad D^3 = n_1^{1-\alpha} + n_2^{1-\alpha} + n_3 \qquad D^4 = n_3$$

Since constituencies have equal populations:

$$D^m a^m = \frac{1}{K} \quad \forall m$$

Where a^m is the area per consituency for each type m. Using this we get:

$$a^1 = \frac{1}{K(n_1^{1-\alpha} + n_3)} \quad a^2 = \frac{1}{K(n_2^{1-\alpha} + n_3)} \quad a^3 = \frac{1}{K(n_1^{1-\alpha} + n_2^{1-\alpha} + n_3)} \quad a^4 = \frac{1}{K(n_3)}$$

Number of consituencies K^m of each type can be calculated by dividing total area of occupied by all constituencies of a given type by a^m:

$$\begin{aligned}
K^1 &= K(n_1^\alpha - O \cdot min(n_1, n_2)^\alpha)(n_1^{1-\alpha} + n_3) \\
K^2 &= K(n_2^\alpha - O \cdot min(n_1, n_2)^\alpha)(n_2^{1-\alpha} + n_3) \\
K^3 &= K(O \cdot min(n_1, n_2)^\alpha)(n_1^{1-\alpha} + n_2^{1-\alpha} + n_3) \\
K^4 &= K(1 - n_1^\alpha - n_2^\alpha + O \cdot min(n_1, n_2)^\alpha)(n_3)
\end{aligned}$$

Proportion of group i in constituency of type m n_i^m:

$$\begin{aligned}
&n_1^1 = \frac{n_1^{1-\alpha}}{n_1^{1-\alpha} + n_3} \quad n_1^2 = 0 \quad n_1^3 = \frac{n_1^{1-\alpha}}{n_1^{1-\alpha} + n_2^{1-\alpha} + n_3} \quad n_1^4 = 0 \\
&n_2^1 = 0 \quad n_2^2 = \frac{n_2^{1-\alpha}}{n_2^{1-\alpha} + n_3} \quad n_2^3 = \frac{n_2^{1-\alpha}}{n_1^{1-\alpha} + n_2^{1-\alpha} + n_3} \quad n_2^4 = 0 \\
&n_3^1 = \frac{n_3}{n_1^{1-\alpha} + n_3} \quad n_3^2 = \frac{n_3}{n_2^{1-\alpha} + n_3} \quad n_3^3 = \frac{n_3}{n_1^{1-\alpha} + n_2^{1-\alpha} + n_3} \quad n_3^4 = 1
\end{aligned}$$

Since, $U(f_j) = \log(f_j)$. Therefore, $U'(f_j) = \frac{1}{f_j}$. Similar to the proof of proposition 2, we can obtain the first order conditions at equilibrium as:

$$\begin{aligned}
\gamma f_1 =& K\phi(n_1^\alpha - O \cdot min(n_1, n_2)^\alpha)(n_1^{1-\alpha} + n_3)(\frac{n_1^{-\alpha}}{\phi n_1^{1-\alpha} + \phi_3 n_3}) \\
&+ K\phi(O \cdot min(n_1, n_2)^\alpha)(n_1^{1-\alpha} + n_2^{1-\alpha} + n_3)(\frac{n_1^{-\alpha}}{\phi(n_1^{1-\alpha} + n_2^{1-\alpha}) + \phi_3 n_3})
\end{aligned}$$

$$\begin{aligned}\gamma f_2 =&K\phi(n_2^{\alpha} - O\cdot min(n_1,n_2)^{\alpha})(n_2^{1-\alpha}+n_3)(\frac{n_2^{-\alpha}}{\phi n_2^{1-\alpha}+\phi_3 n_3})\\&+K\phi(O\cdot min(n_1,n_2)^{\alpha})(n_1^{1-\alpha}+n_2^{1-\alpha}+n_3)(\frac{n_2^{-\alpha}}{\phi(n_1^{1-\alpha}+n_2^{1-\alpha})+\phi_3 n_3})\end{aligned}$$

$$\begin{aligned}\gamma f_3 =&K\phi_3(n_1^{\alpha} - O\cdot min(n_1,n_2)^{\alpha})(n_1^{1-\alpha}+n_3)(\frac{1}{\phi n_1^{1-\alpha}+\phi_3 n_3})\\&+K\phi_3(n_2^{\alpha} - O\cdot min(n_1,n_2)^{\alpha})(n_2^{1-\alpha}+n_3)(\frac{1}{\phi n_2^{1-\alpha}+\phi_3 n_3})\\&+K\phi_3(O\cdot min(n_1,n_2)^{\alpha})(n_1^{1-\alpha}+n_2^{1-\alpha}+n_3)(\frac{1}{\phi(n_1^{1-\alpha}+n_2^{1-\alpha})+\phi_3 n_3})\\&+K\phi_3(1-n_1^{\alpha}-n_2^{\alpha}+O\cdot min(n_1,n_2)^{\alpha})(\frac{1}{\phi_3})\end{aligned}$$

$$n_1f_1+n_2f_2+n_3f_3=S$$

The equilibrium value of per capita private transfers to group 1:

$$f_1=\frac{S\gamma f_1}{n_1\gamma f_1+n_2\gamma f_2+n_3\gamma f_3}$$

Calculating the denominator of the above expression using the first order conditions we get:

$$\begin{aligned}n_1\gamma f_1+n_2\gamma f_2+n_3\gamma f_3 =&K(n_1^{\alpha} - O\cdot min(n_1,n_2)^{\alpha})(n_1^{1-\alpha}+n_3)(\frac{\phi n_1^{1-\alpha}+\phi_3 n_3}{\phi n_1^{1-\alpha}+\phi_3 n_3})\\&+K(n_2^{\alpha} - O\cdot min(n_1,n_2)^{\alpha})(n_2^{1-\alpha}+n_3)(\frac{\phi n_2^{1-\alpha}+\phi_3 n_3}{\phi n_2^{1-\alpha}+\phi_3 n_3})\\&+K(O\cdot min(n_1,n_2)^{\alpha})(n_1^{1-\alpha}+n_2^{1-\alpha}+n_3)(\frac{\phi(n_1^{1-\alpha}+n_2^{1-\alpha})+\phi_3 n_3}{\phi(n_1^{1-\alpha}+n_2^{1-\alpha})+\phi_3 n_3})\\&+K(1-n_1^{\alpha}-n_2^{\alpha}+O\cdot min(n_1,n_2)^{\alpha})(n_3)(\frac{\phi_3 n_3}{\phi_3 n_3})\\=&K(n_1+n_2+n_3)=K\end{aligned}$$

When $n_1 < n_2$, we get from first order condition:

$$\frac{f_1}{S\phi}=\frac{\gamma f_1}{K\phi}=\frac{1-O}{w_1}+\frac{O}{w_3}$$

Where,

$$w_1=\phi+\frac{(\phi_3-\phi)(n_3)}{n_1^{1-\alpha}+n_3} \qquad w_3=\phi+\frac{(\phi_3-\phi)(n_3)}{n_1^{1-\alpha}+n_2^{1-\alpha}+n_3}$$

Derivative of w_1 and w_3 w.r.t. n_1:

$$w_1' = -\frac{(1-\alpha)(\phi_3-\phi)n_3 n_1^{-\alpha}}{(n_1^{1-\alpha}+n_3)^2} \qquad w_3' = -\frac{(1-\alpha)(\phi_3-\phi)n_3(n_1^{-\alpha}-n_2^{-\alpha})}{(n_1^{1-\alpha}+n_2^{1-\alpha}+n_3)^2}$$

As we can see $w_1' < 0$ and $w_3' < 0$ when $n_1 < n_2$. Therefore, $\frac{df_1}{dn_1} < 0$ in this case.
When $n_1 \geq n_2$, we can rewrite the first order condition as:

$$\frac{f_1}{S\phi} = \frac{\gamma f_1}{K\phi} = \frac{1-Or}{w_1} + \frac{Or}{w_3}$$

Where,

$$r = (n_2/n_1)^{\alpha}, \quad r' = -\alpha r(\frac{1}{n_1}+\frac{1}{n_2}), \quad r \in [0,1]$$

Differentiating:

$$\frac{1}{S\phi}\frac{df_1}{dn_1} = \frac{-(1-Or)w_1'}{w_1^2} + Or'(\frac{1}{w_3}-\frac{1}{w_1}) + \frac{-(Or)w_3'}{w_3^2}$$

The first additive term on the R.H.S. is positive and the second and third terms are negative. It can be seen that $\frac{df_1}{dn_1}$ is strictly decreasing in O and is positive as O tends to 0. Therefore, to prove that the expression $\frac{df_1}{dn_1} < 0$ when $O > O^*$ for some $O^* \in (0,1)$, it is sufficient to show tha $\frac{df_1}{dn_1} < 0$ when $O = 1$. Substituting O =1 and rearranging the above expression, we need to show:

$$-\frac{(1-r)w_1'}{w_1^2} < -r'(\frac{1}{w_3}-\frac{1}{w_1}) + \frac{rw_3'}{w_3^2}$$

Substituting the values of w_1, w_2, w_1', w_3', r, r' and simplifying, our expression is reduced to:

$$z - \frac{1}{z} < \frac{\alpha(n_2/n_1+1)}{(1-\alpha)(1-(n_2/n_1)^{\alpha})}$$

Where $z = 1 + \frac{\phi n_2^{1-\alpha}}{\phi n_1^{1-\alpha}+n_3}$

$$\implies \phi n_2^{1-\alpha}(2+\frac{\phi n_2^{1-\alpha}}{\phi n_1^{1-\alpha}+\phi_3 n_3}) < \frac{\alpha(n_2/n_1+1)(\phi(n_1^{1-\alpha}+n_2^{1-\alpha})+\phi_3 n_3)}{(1-\alpha)(1-(n_2/n_1)^{\alpha})}$$

As the ratio $\frac{\phi_3}{\phi}$ increases, the above inequality will be satisfied more easily. Therefore, it is sufficient to show that weak inequality holds in the above expression when $\phi_3 = \phi$. Using this and rearranging, we now need to show:

$$(n_1^{1-\alpha}n_2^{1-\alpha})(2+\frac{n_2^{1-\alpha}}{n_1^{1-\alpha}+n_3}) \leq \frac{\alpha(n_1+n_2)(n_1^{1-\alpha}+n_2^{1-\alpha}+n_3)}{(1-\alpha)(n_1^{\alpha}-n_2^{\alpha})}$$

This can be rearranged to give:

$$n_1^{3-2\alpha}X + n_1^{2-\alpha}n_3Y \leq 0$$

Where,

$$\begin{aligned}
X &= (2-3\alpha)q^{1-\alpha} - (2-\alpha)q - \alpha - \alpha q^{2-\alpha} \\
Y &= (2-3\alpha)q^{1-\alpha} - (2-\alpha)q - \alpha - \alpha q^{2-\alpha} - \alpha(1+q+\frac{n_3}{n_1^{1-\alpha}}(1+q)) \\
q &= \frac{n_2}{n_1}, \qquad q \in [0,1]
\end{aligned}$$

As we can see, $Y < X$ and n_3 can take any value in $(0,1)$, therefore it is both necessary and sufficient to show that $X \leq 0$. In fact, it is sufficient to show that:

$$x(q,\alpha) = (2-3\alpha)q^{1-\alpha} - (2-\alpha)q - \alpha \leq 0 \qquad \forall q \in [0,1], \quad \alpha \in (0,1)$$

Since x is continuous in q, the above condition will hold if it can be shown to hold at the boundaries and at each critical point in (0,1). At the boundaries:

$$\begin{aligned}
x(0,\alpha) &= -\alpha < 0 \\
x(1,\alpha) &= -3\alpha < 0
\end{aligned}$$

At critical point q^*:

$$\begin{aligned}
\frac{dx(q,\alpha)}{dq} &= (1-\alpha)(2-3\alpha)q^{-\alpha} - 2 + \alpha = 0 \\
\implies q^* &= (\frac{(1-\alpha)(2-3\alpha)}{2-\alpha})
\end{aligned}$$

$\therefore q^* \in (0,1)$ only when $\alpha \in (0, \frac{2}{3})$. Substituting the value of q^* and simplifying we need to show:

$$\begin{aligned}
x(q^*,\alpha) = \alpha((\frac{1-\alpha}{2-\alpha})^{\frac{1-\alpha}{\alpha}}(2-3\alpha)^{\frac{1}{\alpha}} - 1) &\leq 0 \\
\implies (\frac{2-\alpha}{1-\alpha})^{1-\alpha} &\geq 2-3\alpha
\end{aligned}$$

Let $t = 1-\alpha$. Now we need to show:

$$y(t) = (1+\frac{1}{t})^t - 3t + 1 \geq 0 \qquad \forall t \in (\frac{1}{3}, 1)$$

Again, since y(t) is continuous in t, we only need to show that the above condition is true at the boundary points and at each critical point in $(\frac{1}{3}, 1)$. At the boundaries:

$$y(\frac{1}{3}) = 4^{\frac{1}{3}} > 0$$
$$y(1) = 0$$

At the critical point:

$$\frac{dy(t)}{dt} = (1+\frac{1}{t})^t(ln(1+\frac{1}{t}) - \frac{1}{1+t}) - 3 = 0$$

Substituting the value of $(1+\frac{1}{t})^t$ in $y(t)$ and rearranging sides, we now need to show:

$$(3t-1)(ln(1+\frac{1}{t}) - \frac{1}{1+t}) \leq 3$$

Since $t \in (\frac{1}{3}, 1)$, therefore:

$$3t - 1 < 2 \qquad ln(1+\frac{1}{t}) < ln(4) \qquad \frac{1}{1+t} > \frac{1}{2}$$
$$\therefore (3t-1)(ln(1+\frac{1}{t}) - \frac{1}{1+t}) < 2(ln(4) - \frac{1}{2}) = 1.77 < 3$$

This implies that $x(q^*, \alpha) \leq 0$. Thus, $x(q, t) \leq 0$. Therefore, when $n_1 \geq n_2$, $\frac{df_1}{dn_1} < 0$ if and only if $O > O^*$ for some $O^* \in (0, 1)$.

(b) When group 2 is dispersed, settlement areas of each group are:

$$A_1 = n_1^\alpha \qquad A_2 = 1 \qquad A_3 = 1$$

In this case, there are two types of constituencies: (1) Group 1, 2 and 3 all reside and (2) Only group 2 and 3 reside. Densities of constituencies are:

$$D^1 = n_1^{1-\alpha} + n_2 + n_3 \qquad D^2 = n_2 + n_3$$

Since the populations across the K constituency are equal, we can calculate area per constituency:

$$a^1 = \frac{1}{K(n_1^{1-\alpha} + n_2 + n_3)} \qquad a^2 = \frac{1}{K(n_1 + n_2)}$$

Number of constituencies of each type:

$$K^1 = Kn_1^\alpha(n_1^{1-\alpha} + n_2 + n_3) \qquad K^2 = K(1 - n_1^\alpha)(n_2 + n_3)$$

Group proportions in each constituency type:

$$n_1^1 = \frac{n_1^{1-\alpha}}{n_1^{1-\alpha} + n_2 + n_3} \qquad n_1^2 = 0$$
$$n_2^1 = \frac{n_2}{n_1^{1-\alpha} + n_2 + n_3} \qquad n_2^2 = \frac{n_2}{n_2 + n_3}$$
$$n_3^1 = \frac{n_3}{n_1^{1-\alpha} + n_2 + n_3} \qquad n_3^2 = \frac{n_3}{n_2 + n_3}$$

We get first order conditions similar to (a). At equilibrium:

$$\gamma f_1 = K\phi(n_1^\alpha)(n_1^{1-\alpha} + n_2 + n_3)\frac{n_1^{-\alpha}}{\phi(n_1^{1-\alpha} + n_2) + \phi_3 n_3}$$

$$\begin{aligned}\gamma f_2 =& K\phi(n_1^\alpha)(n_1^{1-\alpha} + n_2 + n_3)\frac{1}{\phi(n_1^{1-\alpha} + n_2) + \phi_3 n_3} \\ &+ K\phi(1 - n_1^\alpha)(n_2 + n_3)\frac{1}{\phi n_2 + \phi_3 n_3}\end{aligned}$$

$$\begin{aligned}\gamma f_3 =& K\phi_3(n_1^\alpha)(n_1^{1-\alpha} + n_2 + n_3)\frac{1}{\phi(n_1^{1-\alpha} + n_2) + \phi_3 n_3} \\ &+ K\phi_3(1 - n_1^\alpha)(n_2 + n_3)\frac{1}{\phi n_2 + \phi_3 n_3}\end{aligned}$$

$$n_1 f_1 + n_2 f_2 + n_3 f_3 = S$$

Similar to the proof in (a), equilibrium per capita transfer to group 2 are:

$$f_1 = \frac{S\gamma f_1}{n_1\gamma f_1 + n_2\gamma f_2 + n_3\gamma f_3}$$

Calculating the denominator by substituting values from first order condition:

$$\begin{aligned}n_1\gamma f_1 + n_2\gamma f_2 + n_3\gamma f_3 =& K(n_1^\alpha)(n_1^{1-\alpha} + n_2 + n_3)\frac{\phi(n_1^{1-\alpha} + n_2) + \phi_3 n_3}{\phi(n_1^{1-\alpha} + n_2) + \phi_3 n_3} + \\ & K(1 - n_1^\alpha)(n_2 + n_3)\frac{\phi n_2 + \phi_3 n_3}{\phi n_2 + \phi_3 n_3} \\ =& K(n_1 + n_2 + n_3) = K\end{aligned}$$

Using this and the first order condition:

$$\frac{f_1}{S\phi} = \frac{\gamma f_1}{K\phi} = \frac{n_1^{1-\alpha} + n_2 + n_3}{\phi(n_1^{1-\alpha} + n_2) + \phi_3 n_3}$$

Differentiating and simplifying:

$$\frac{1}{S\phi}\frac{df_1}{dn_1} = \frac{(\phi_3 - \phi)n_3((1-\alpha)n_1^{-\alpha} - 1)}{(\phi(n_1^{1-\alpha} + n_2) + \phi_3 n_3)^2}$$

Since, $\phi_3 > \phi$, it follows:

$$\frac{df_1}{dn_1} > 0 \quad \text{if} \quad n_1 < (1-\alpha)^{\frac{1}{\alpha}}$$
$$\frac{df_1}{dn_1} < 0 \quad \text{if} \quad n_1 > (1-\alpha)^{\frac{1}{\alpha}}$$

$\therefore$ There is an inverted-U shaped relation between n_1 and f_1^* and hence between n_1 and G_1^* with peak at $n_1^* = (1-\alpha)^{\frac{1}{\alpha}}$.

Bibliography

Abele, A. E. and B. Wojciszke (2014): "Communal and agentic content in social cognition: A dual perspective model," in *Advances in experimental social psychology*, Elsevier, vol. 50, 195–255.

Abrahams, A., C. Oram, and N. Lozano-Gracia (2018): "Deblurring DMSP nighttime lights: A new method using Gaussian filters and frequencies of illumination," *Remote sensing of environment*, 210, 242–258.

Afridi, F., T. Dinkelman, and K. Mahajan (2018): "Why are fewer married women joining the work force in rural India? A decomposition analysis over two decades," *Journal of Population Economics*, 31, 783–818.

Afridi, F., V. Iversen, and M. Sharan (2017): "Women Political Leaders, Corruption, and Learning: Evidence from a Large Public Program in India," *Economic Development and Cultural Change*, 66, 000–000.

Akerlof, G. A. and R. E. Kranton (2000): "Economics and identity," *The quarterly journal of economics*, 115, 715–753.

Alesina, A. (1988): "Credibility and policy convergence in a two-party system with rational voters," *The American Economic Review*, 78, 796–805.

Alesina, A., S. Michalopoulos, and E. Papaioannou (2016): "Ethnic inequality," *Journal of Political Economy*, 124, 428–488.

Ancona, M., E. Ceolini, C. Öztireli, and M. Gross (2017): "Towards better understanding of gradient-based attribution methods for deep neural networks," *arXiv preprint arXiv:1711.06104*.

Arbabzadah, F., G. Montavon, K.-R. Müller, and W. Samek (2016): "Identifying individual facial expressions by deconstructing a neural network," in *German Conference on Pattern Recognition*, Springer, 344–354.

Arceo-Gómez, E. O., R. M. Campos-Vázquez, R. Y. B. Salas, and S. López-Araiza (2020): "Gender Stereotypes in Job Advertisements: What Do They Imply for the Gender Salary Gap?" Mexico. Retrieved from http://conference. iza. org/conference_files . . . .

Arras, L., F. Horn, G. Montavon, K.-R. Müller, and W. Samek (2017a): "What is relevant in a text document?: An interpretable machine learning approach," *PloS one*, 12, e0181142.

Arras, L., G. Montavon, K.-R. Müller, and W. Samek (2017b): "Explaining Recurrent Neural Network Predictions in Sentiment Analysis," in *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 159–168.

Bach, S., A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek (2015): "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PloS one*, 10, e0130140.

Ban, R. and V. Rao (2008): "Tokenism or agency? The impact of women's reservations on village democracies in South India," *Economic Development and Cultural Change*, 56, 501–530.

Banducci, S. A., T. Donovan, and J. A. Karp (1999): "Proportional representation and attitudes about politics: results from New Zealand," *Electoral Studies*, 18, 533–555.

——— (2004): "Minority representation, empowerment, and participation," *The Journal of Politics*, 66, 534–556.

Banerjee, A. and R. Somanathan (2001): "A simple model of voice," *The Quarterly Journal of Economics*, 116, 189–227.

Banfi, S., S. Choi, and B. Villena-Roldan (2019): "Deconstructing Job Search Behavior," Unpublished manuscript.

Banfi, S. and B. Villena-Roldan (2019): "Do high-wage jobs attract more applicants? Directed search evidence from the online labor market," *Journal of Labor Economics*, 37, 715–746.

Bardhan, P. K., D. Mookherjee, and M. P. Torrado (2005): "Impact of Reservations of Panchayat Pradhans on Targeting in West Bengal," *BREAD Working Paper*.

——— (2010): "Impact of political reservations in West Bengal local governments on anti-poverty targeting," *Journal of Globalization and development*, 1.

Barnes, T. D. and S. M. Burchard (2013): ""Engendering" politics: The impact of descriptive representation on women's political engagement in Sub-Saharan Africa," *Comparative Political Studies*, 46, 767–790.

Baskaran, T., B. Min, and Y. Uppal (2015): "Election cycles and electricity provision: Evidence from a quasi-experiment with Indian special elections," *Journal of Public Economics*, 126, 64–73.

Beaman, L., R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova (2009): "Powerful women: does exposure reduce bias?" *The Quarterly journal of economics*, 124, 1497–1540.

Beaman, L., E. Duflo, R. Pande, and P. Topalova (2007): "Women politicians, gender bias, and policy-making in rural India." Tech. rep.

——— (2012): "Female leadership raises aspirations and educational attainment for girls: A policy experiment in India," *science*, 335, 582–586.

Becker, S., M. Ackermann, S. Lapuschkin, K.-R. Müller, and W. Samek (2018): "Interpreting and explaining deep neural networks for classification of audio signals," *arXiv preprint arXiv:1807.03418*.

Belot, M., P. Kircher, and P. Muller (2017): "How Wage Announcements Affect Job Search Behaviour - A Field Experimental Investigation," Unpublished manuscript.

Besley, T., R. Pande, L. Rahman, and V. Rao (2004): "The politics of public good provision: Evidence from Indian local governments," *Journal of the European Economic Association*, 2, 416–426.

Besley, T., R. Pande, and V. Rao (2007): "Political economy of panchayats in South India," *Economic and Political Weekly*, 661–666.

Bettencourt, L. M. (2013): "The origins of scaling in cities," *science*, 340, 1438–1441.

Bhalotra, S. and I. Clots-Figueras (2014): "Health and the political agency of women," *American Economic Journal: Economic Policy*, 6, 164–97.

Bhalotra, S., I. Clots-Figueras, G. Cassan, and L. Iyer (2014): "Religion, politician identity and development outcomes: Evidence from India," *Journal of Economic Behavior & Organization*, 104, 4–17.

Bhalotra, S. R., D. Clarke, and J. Gomes (2020): "Maternal mortality and women's political participation," *CEPR Discussion Paper No. DP14339*.

Binder, A., M. Bockmayr, M. Hägele, S. Wienert, D. Heim, K. Hellweg, A. Stenzinger, L. Parlow, J. Budczies, B. Goeppert, et al. (2018): "Towards computational fluorescence microscopy: Machine learning-based integrated prediction of morphological and molecular tumor profiles," *arXiv preprint arXiv:1805.11178*.

Blais, A. and I. Lago (2009): "A general measure of district competitiveness," *Electoral Studies*, 28, 94–100.

Bobo, L. and F. D. Gilliam (1990): "Race, sociopolitical participation, and black empowerment," *American Political Science Review*, 84, 377–393.

Boix, C. (1999): "Setting the rules of the game: the choice of electoral systems in advanced democracies," *American political science review*, 93, 609–624.

Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2019): "Beliefs about gender," *American Economic Review*, 109, 739–73.

Bormann, N.-C., L.-E. Cederman, and M. Vogt (2017): "Language, religion, and ethnic civil war," *Journal of Conflict Resolution*, 61, 744–771.

Bormann, N.-C. and M. Golder (2013): "Democratic electoral systems around the world, 1946–2011," *Electoral Studies*, 32, 360–369.

Born, M. P. and T. W. Taris (2010): "The impact of the wording of employment advertisements on students' inclination to apply for a job," *The Journal of social psychology*, 150, 485–502.

Braga, M. and F. Scervini (2017): "The performance of politicians: The effect of gender quotas," *European Journal of Political Economy*, 46, 1–14.

Burgess, R., R. Jedwab, E. Miguel, A. Morjaria, and G. Padró i Miquel (2015): "The value of democracy: evidence from road building in Kenya," *American Economic Review*, 105, 1817–51.

Burn, I., P. Button, L. F. M. Corella, and D. Neumark (2019): "Older Workers Need Not Apply? Ageist Language in Job Ads and Age Discrimination in Hiring," Tech. rep., National Bureau of Economic Research.

Bustelo, M., A. M. Díaz Escobar, J. Lafortune, C. Piras, L. M. Salas Bahamón, J. Tessada, et al. (2020): "What is The Price of Freedom?: Estimating Women's Willingness to Pay for Job Schedule Flexibility," Tech. rep., Inter-American Development Bank.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014): "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 82, 2295–2326.

Campa, P. (2011): "Gender quotas, female politicians and public expenditures: quasi-experimental evidence," .

Campos-Vazquez, R., G. Esquivel, and R. Badillo (2020): "How has labor demand been affected by the COVID-19 pandemic? Evidence from job ads in Mexico," CEPR Press.

Cascio, E. U. and E. Washington (2013): "Valuing the vote: The redistribution of voting rights and state funds following the voting rights act of 1965," *The Quarterly Journal of Economics*, 129, 379–433.

Cederman, L.-E., K. S. Gleditsch, I. Salehyan, and J. Wucherpfennig (2013): "Transborder ethnic kin and civil war," *International Organization*, 67, 389–410.

Cederman, L.-E., A. Wimmer, and B. Min (2010): "Why do ethnic groups rebel? New data and analysis," *World Politics*, 62, 87–119.

Cesaretti, R., J. Lobo, L. M. Bettencourt, S. G. Ortman, and M. E. Smith (2016): "Population-area relationship for Medieval European cities," *PloS one*, 11, e0162678.

Chattopadhyay, R. and E. Duflo (2004): "Women as policy makers: Evidence from a randomized policy experiment in India," *Econometrica*, 72, 1409–1443.

Chaturvedi, R. and S. Chaturvedi (2020): "It's All in the Name: A Character Based Approach To Infer Religion," *arXiv preprint arXiv:2010.14479*.

Chiplunkar, G., E. Kelley, and G. Lane (2020): "Which jobs are lost during a lockdown? Evidence from vacancy posting in India," Unpublished Manuscript.

Chollet, F. et al. (2015): "Keras," .

Chowdhury, A. R., A. C. Areias, S. Imaizumi, S. Nomura, and F. Yamauchi (2018): *Reflections of employers' gender preferences in job ads in India: an analysis of online job portal data*, The World Bank.

Clots-Figueras, I. (2011): "Women in politics: Evidence from the Indian States," *Journal of Public Economics*, 95, 664–690.

——— (2012): "Are female leaders good for education? Evidence from India," *American Economic Journal: Applied Economics*, 4, 212–44.

Coffey, D., A. Gupta, P. Hathi, N. Khurana, D. Spears, N. Srivastav, and S. Vyas (2014): "Revealed Preference for Open Defecation: Evidence from a New Survey in Rural North India," *Economic & Political Weekly*, 49, 43–55.

Coffey, D., A. Gupta, P. Hathi, D. Spears, N. Srivastav, and S. Vyas (2017): "Understanding Open Defecation in Rural India: Untouchability, Pollution, and Latrine Pits," *Economic & Political Weekly*, 52, 59–66.

Colomer, J. M. (2004): "The strategy and history of electoral system choice," in *The Handbook of electoral system choice*, Springer, 3–78.

Dal Bó, E., F. Finan, and M. Rossi (2013): "Strengthening state capabilities: The role of financial incentives in the call to public service," *Quarterly Journal of Economics*, 128(3), 1169–218.

Das, S., A. Mukhopadhyay, and R. Saroy (2017): "Efficiency consequences of affirmative action in politics: Evidence from india," *IZA Discussion Paper*.

De Luca, G., R. Hodler, P. A. Raschky, and M. Valsecchi (2018): "Ethnic favoritism: An axiom of politics?" *Journal of Development Economics*, 132, 115–129.

Deininger, K., S. Jin, H. K. Nagarajan, and F. Xia (2015): "Does Female Reservation Affect Long-Term Political Outcomes? Evidence from Rural India," *The Journal of Development Studies*, 51, 32–49.

Deming, D. and L. B. Kahn (2018): "Skill requirements across firms and labor markets: Evidence from job postings for professionals," *Journal of Labor Economics*, 36, S337–S369.

Deshpande, A., D. Goel, and S. Khanna (2018): "Bad karma or discrimination? Male–female wage gaps among salaried workers in India," *World Development*, 102, 331–344.

Devlin, C. and R. Elgie (2008): "The effect of increased women's representation in parliament: The case of Rwanda," *Parliamentary Affairs*, 61, 237–254.

Dhar, D., T. Jain, and S. Jayachandran (2018): "Reshaping adolescents' gender attitudes: Evidence from a school-based experiment in India," Tech. rep., National Bureau of Economic Research.

DICKENS, A. (2018): "Ethnolinguistic favoritism in african politics," *American Economic Journal: Applied Economics*, 10, 370–402.

DIMICO, A. (2017): "Size matters: the effect of the size of ethnic groups on development," *Oxford Bulletin of Economics and Statistics*, 79, 291–318.

DING, Y., Y. LIU, H. LUAN, AND M. SUN (2017): "Visualizing and understanding neural machine translation," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1150–1159.

DOLL, C. N. (2008): "CIESIN thematic guide to night-time light remote sensing and its applications," *Center for International Earth Science Information Network of Columbia University, Palisades, NY.*

DOLLAR, D., R. FISMAN, AND R. GATTI (2001): "Are women really the "fairer" sex? Corruption and women in government," *Journal of Economic Behavior & Organization*, 46, 423–429.

DOW, J. K. (2011): "Party-system extremism in majoritarian and proportional electoral systems," *British Journal of Political Science*, 41, 341–361.

DOWNS, A. (1957): "An economic theory of political action in a democracy," *Journal of political economy*, 65, 135–150.

DOZAT, T. (2016): "Incorporating nesterov momentum into adam," .

DREZE, J. AND P. SRINIVASAN (1997): "Widowhood and poverty in rural India: Some inferences from household survey data," *Journal of Development Economics*, 54, 217–234.

EGGERS, A. C., A. FOWLER, J. HAINMUELLER, A. B. HALL, AND J. M. SNYDER JR (2015): "On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races," *American Journal of Political Science*, 59, 259–274.

EJDEMYR, S., E. KRAMON, AND A. L. ROBINSON (2018): "Segregation, ethnic favoritism, and the strategic targeting of local public goods," *Comparative Political Studies*, 51, 1111–1143.

FARRELL, D. M. (2011): *Electoral systems: a comparative introduction*, Macmillan International Higher Education.

FEARON, J. D. (2003): "Ethnic and cultural diversity by country," *Journal of economic growth*, 8, 195–222.

FERREIRA, F. AND J. GYOURKO (2014): "Does gender matter for political leadership? The case of US mayors," *Journal of Public Economics*, 112, 24–39.

FJELDE, H. AND K. HÖGLUND (2016): "Electoral institutions and electoral violence in Sub-Saharan Africa," *British Journal of Political Science*, 46, 297–320.

FLETCHER, E., C. MOORE, AND R. PANDE (2018): "Women and Work in India: Descriptive Evidence and a Review of Potential Policies," Unpublished Manuscript.

Flory, J., A. Leibbrandt, and J. List (2015): "Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions," *Review of Economic Studies*, 82(1).

Forsythe, E., L. Kahn, F. Lange, and D. Wiczer (2020): "Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims," *Journal of Public Economics*, 189, 104238.

Franceschet, S. and J. M. Piscopo (2008): "Gender quotas and women's substantive representation: Lessons from Argentina," *Politics & Gender*, 4, 393–425.

Franck, R. and I. Rainer (2012): "Does the leader's ethnicity matter? Ethnic favoritism, education, and health in sub-Saharan Africa," *American Political Science Review*, 294–325.

Francois, P., I. Rainer, and F. Trebbi (2015): "How is power shared in Africa?" *Econometrica*, 83, 465–503.

Gajwani, K. and X. Zhang (2014): "Gender and public goods provision in Tamil Nadu's village governments," *The World Bank Economic Review*, 29, 234–261.

Gallagher, M. (1991): "Proportionality, disproportionality and electoral systems," *Electoral studies*, 10, 33–51.

——— (1992): "Comparing proportional representation electoral systems: Quotas, thresholds, paradoxes and majorities," *British Journal of Political Science*, 22, 469–496.

Gangadharan, L., T. Jain, P. Maitra, and J. Vecci (2016): "Social identity and governance: The behavioral response to female leaders," *European Economic Review*, 90, 302–325.

Gangopadhyay, S. and W. Wadhwa (2004): *Are Indian female-headed households more vulnerable to poverty*, Bazaar Chintan New Delhi, India.

Gaucher, D., J. Friesen, and A. C. Kay (2011): "Evidence that gendered wording in job advertisements exists and sustains gender inequality." *Journal of personality and social psychology*, 101, 109.

Geruso, M. and D. Spears (2018): "Neighborhood sanitation and infant mortality," *American Economic Journal: Applied Economics*, 10, 125–62.

Ghosh, S. and A. Mitra (2019): "Ethnic identities, public spending and political regimes," Tech. rep., School of Economics Discussion Papers.

Goldin, C. (2014): "A grand gender convergence: Its last chapter," *American Economic Review*, 104, 1091–1119.

Goldin, C. and L. F. Katz (2011): "The cost of workplace flexibility for high-powered professionals," *The Annals of the American Academy of Political and Social Science*, 638, 45–67.

Grembi, V., T. Nannicini, and U. Troiano (2016): "Do fiscal rules matter?" *American Economic Journal: Applied Economics*, 1–30.

Gulzar, S., N. Haas, and B. Pasquale (2020): "Does Political Affirmative Action Work, and for Whom? Theory and Evidence on India's Scheduled Areas," *American Political Science Review*, 11, 1230–1246.

Gupta, A., N. Khalid, D. Desphande, P. Hathi, A. Kapur, N. Srivastav, S. Vyas, D. Spears, and D. Coffey (2019): "Changes in Open Defecation in Rural North India: 2014-2018," *IZA Discussion Paper*.

Hägele, M., P. Seegerer, S. Lapuschkin, M. Bockmayr, W. Samek, F. Klauschen, K.-R. Müller, and A. Binder (2020): "Resolving challenges in deep learning-based analyses of histopathological images using explanation methods," *Scientific reports*, 10, 1–12.

Hayashi, R. and N. Matsuda (2020): "COVID-19 impact on job postings: Real time assessment using Bangladesh and Sri Lanka online job portals," Asian Development Bank, ADB Briefs.

He, H., D. Neumark, and Q. Weng (2019): "Do Workers Value Flexible Jobs? A Field Experiment," Tech. rep., National Bureau of Economic Research.

He, K., X. Zhang, S. Ren, and J. Sun (2015): "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 1026–1034.

Heath, O., G. Verniers, and S. Kumar (2015): "Do Muslim voters prefer Muslim candidates? Co-religiosity and voting behaviour in India," *Electoral Studies*, 38, 10–18.

Helleseter, M. D., P. Kuhn, and K. Shen (2020): "The Age Twist in Employers' Gender Requests Evidence from Four Job Boards," *Journal of Human Resources*, 55, 428–469.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012): "Measuring economic growth from outer space," *American economic review*, 102, 994–1028.

Herrera, H., M. Morelli, and T. Palfrey (2014): "Turnout and power sharing," *The Economic Journal*, 124, F131–F162.

Hershbein, B. and L. Kahn (2018): "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," *American Economic Review*, 108, 1737–1772.

Hicks, D. L., J. H. Hicks, and B. Maldonado (2016): "Women as policy makers and donors: Female legislators and foreign aid," *European Journal of Political Economy*, 41, 46–60.

Hinton, G. E., N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov (2012): "Improving neural networks by preventing co-adaptation of feature detectors," *arXiv preprint arXiv:1207.0580*.

Hodler, R. and P. A. Raschky (2014): "Regional favoritism," *The Quarterly Journal of Economics*, 129, 995–1033.

Horst, F., S. Lapuschkin, W. Samek, K.-R. Müller, and W. I. Schöllhorn (2019): "Explaining the unique nature of individual gait patterns with deep learning," *Scientific reports*, 9, 1–13.

Huber, J. D. (2012): "Measuring ethnic voting: Do proportional electoral laws politicize ethnicity?" *American Journal of Political Science*, 56, 986–1001.

Hyytinen, A., J. Meriläinen, T. Saarimaa, O. Toivanen, and J. Tukiainen (2018): "When does regression discontinuity design work? Evidence from random election outcomes," *Quantitative Economics*, 9, 1019–1051.

Ibanez, M. and G. Reiner (2018): "Sorting through affirmative action: three field experiments in Colombia," *Journal of Labor Economics*, 36(2).

Iyer, L. and A. Mani (2019): "The road not taken: Gender gaps along paths to political power," *World Development*, 119, 68–80.

Jenkins, M. W. and V. Curtis (2005): "Achieving the 'good life': Why some people want latrines in rural Benin," *Social science & medicine*, 61, 2446–2459.

Jensen, R. (2012): "Do labor market opportunities affect young women's work and family decisions? Experimental evidence from India," *The Quarterly Journal of Economics*, 127, 753–792.

Kartal, M. (2015): "Laboratory elections with endogenous turnout: proportional representation versus majoritarian rule," *Experimental Economics*, 18, 366–384.

Khanna, T. and M. Das (2016): "Why gender matters in the solution towards safe sanitation? Reflections from rural India," *Global public health*, 11, 1185–1201.

Klasen, S. and J. Pieters (2015): "What Explains the Stagnation of Female Labor Force Participation in Urban India?" *The World Bank Economic Review*, 29, 449–478.

Kuhn, P. and K. Shen (2013): "Gender discrimination in job ads: Evidence from china," *The Quarterly Journal of Economics*, 128, 287–336.

Kuhn, P., K. Shen, and S. Zhang (2020): "Gender-targeted job ads in the recruitment process: Facts from a Chinese job board," *Journal of Development Economics*, 102531.

Kunicova, J. and S. Rose-Ackerman (2005): "Electoral rules and constitutional structures as constraints on corruption," *British journal of political science*, 35, 573–606.

Kunze, A. (2018): "The gender wage gap in developed countries," *The Oxford Handbook of Women and the Economy*, 369.

Lapuschkin, S., A. Binder, K.-R. Muller, and W. Samek (2017): "Understanding and comparing deep neural networks for age and gender classification," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 1629–1638.

Lapuschkin, S., S. Wäldchen, A. Binder, G. Montavon, W. Samek, and K.-R. Müller (2019): "Unmasking clever hans predictors and assessing what machines really learn," *Nature communications*, 10, 1–8.

LeCun, Y., B. E. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. E. Hubbard, and L. D. Jackel (1990): "Handwritten digit recognition with a back-propagation network," in *Advances in neural information processing systems*, 396–404.

Lijphart, A. (1992): "Democratization and constitutional choices in Czecho-Slovakia, Hungary and Poland 1989-91," *Journal of theoretical politics*, 4, 207–223.

Lizzeri, A. and N. Persico (2001): "The provision of public goods under alternative electoral incentives," *American Economic Review*, 91, 225–239.

Mahajan, K. and S. Sekhri (2020): "Access to toilets and violence against women," .

Marinescu, I. and R. Wolthoff (2020): "Opening the black box of the matching function: The power of words," *Journal of Labor Economics*, 38, 535–568.

Marshall, M. G., T. R. Gurr, and K. Jaggers (2016): "Polity IV project: Political regime characteristics and transitions, 1800-2015," *Center for Systemic Peace*.

Mas, A. and A. Pallais (2017): "Valuing alternative work arrangements," *American Economic Review*, 107(12).

McCallum, A., K. Nigam, et al. (1998): "A comparison of event models for naive bayes text classification," in *AAAI-98 workshop on learning for text categorization*, Citeseer, vol. 752, 41–48.

McCrary, J. (2008): "Testing for manipulation of the running variable in the regression discontinuity design, forthcoming," *Journal of Econometrics*.

Meijer, J. R., M. A. Huijbregts, K. C. Schotten, and A. M. Schipper (2018): "Global patterns of current and future road infrastructure," *Environmental Research Letters*, 13, 064006.

Menzio, G., I. Telyukova, and L. Visschers (2016): "Directed search over the life cycle. Special issue in honor of Dale Mortensen," *Review of Economic Dynamics*, 19, 38–62.

Meyersson, E. (2014): "Islamic Rule and the Empowerment of the Poor and Pious," *Econometrica*, 82, 229–269.

Michalopoulos, S. and E. Papaioannou (2013a): "National institutions and subnational development in Africa," *The Quarterly journal of economics*, 129, 151–213.

——— (2013b): "Pre-colonial ethnic institutions and contemporary African development," *Econometrica*, 81, 113–152.

Milesi-Ferretti, G. M., R. Perotti, and M. Rostagno (2002): "Electoral systems and public spending," *The Quarterly Journal of Economics*, 117, 609–657.

Min, B. and K. Gaba (2014): "Tracking electrification in Vietnam using nighttime lights," *Remote Sensing*, 6, 9511–9529.

Min, B., K. M. Gaba, O. F. Sarr, and A. Agalassou (2013): "Detection of rural electrification in Africa using DMSP-OLS night lights imagery," *International journal of remote sensing*, 34, 8118–8141.

Moen, E. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2).

Moser, R. G. (2008): "Electoral systems and the representation of ethnic minorities: Evidence from Russia," *Comparative Politics*, 40, 273–292.

Myerson, R. B. (1999): "Theoretical comparisons of electoral systems," *European Economic Review*, 43, 671–697.

Ng, A. Y. and M. I. Jordan (2002): "On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes," in *Advances in neural information processing systems*, 841–848.

Ningrum, P., T. Pansombut, and A. Ueranantasun (2020): "Text mining of online job advertisements to identify direct discrimination during job hunting process: A case study in Indonesia," *Plos One*, 15(6), e0233746.

Ortman, S. G., A. H. Cabaniss, J. O. Sturm, and L. M. Bettencourt (2014): "The pre-history of urban scaling," *PloS one*, 9, e87902.

——— (2015): "Settlement scaling and increasing returns in an ancient society," *Science Advances*, 1, e1400066.

Ortman, S. G., K. E. Davis, J. Lobo, M. E. Smith, L. M. Bettencourt, and A. Trumbo (2016): "Settlement scaling and economic change in the Central Andes," *Journal of Archaeological Science*, 73, 94–106.

Paul, H. (2014): "ICOW Colonial History Data Set, version 1.0," .

Pennington, J., R. Socher, and C. D. Manning (2014): "GloVe: Global Vectors for Word Representation," in *Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.

Persson, T. and G. Tabellini (2004): "Constitutional rules and fiscal policy outcomes," *American Economic Review*, 94, 25–45.

Persson, T. and G. E. Tabellini (2002): *Political economics: explaining economic policy*, MIT press.

——— (2005): *The economic effects of constitutions*, MIT press.

Prakash, N., M. Rockmore, and Y. Uppal (2019): "Do criminally accused politicians affect economic outcomes? Evidence from India," *Journal of Development Economics*, 102370.

Qiang, J., Z. Qian, Y. Li, Y. Yuan, and X. Wu (2020): "Short text topic modeling techniques, applications, and performance: a survey," *IEEE Transactions on Knowledge and Data Engineering.*

Reynolds, A., B. Reilly, and A. Ellis (2008): *Electoral system design: The new international IDEA handbook*, International Institute for Democracy and Electoral Assistance.

Rickard, S. J. (2012): "A non-tariff protectionist bias in majoritarian politics: Government subsidies and electoral institutions," *International Studies Quarterly*, 56, 777–785.

Rogerson, R., R. Shimer, and R. Wright (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 43(4), 959–988.

Samek, W., A. Binder, G. Montavon, S. Lapuschkin, and K.-R. Müller (2016): "Evaluating the visualization of what a deep neural network has learned," *IEEE transactions on neural networks and learning systems*, 28, 2660–2673.

Santos, A. C., J. A. Roberts, M. L. Barreto, and S. Cairncross (2011): "Demand for sanitation in Salvador, Brazil: A hybrid choice approach," *Social science & medicine*, 72, 1325–1332.

Shi, S. (2002): "A directed search model of inequality with heterogeneous skills and skill-biased technology," *Review of Economic Studies*, 69(2), 467–91.

Silva, T. (2016): "Who Gets What? Formateur Advantage and Salience in Portfolio Allocation in Presidential Multiparty Systems," Tech. rep., Working Paper.

Srinivasan, V., S. Lapuschkin, C. Hellge, K.-R. Müller, and W. Samek (2017): "Interpretable human action recognition in compressed domain," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 1692–1696.

Stopnitzky, Y. (2017): "No toilet no bride? Intrahousehold bargaining in male-skewed marriage markets in India," *Journal of Development Economics*, 127, 269–282.

Sturm, I., S. Lapuschkin, W. Samek, and K.-R. Müller (2016): "Interpretable deep neural networks for single-trial EEG classification," *Journal of neuroscience methods*, 274, 141–145.

Swamy, A., S. Knack, Y. Lee, and O. Azfar (2001): "Gender and corruption," *Journal of Development Economics*, 64, 25–55.

Taris, T. W. and I. A. Bok (1998): "On gender specificity of person characteristics in personnel advertisements: A study among future applicants," *The Journal of psychology*, 132, 593–610.

Thomas, A. W., H. R. Heekeren, K.-R. Müller, and W. Samek (2019): "Analyzing neuroimaging data through recurrent deep learning models," *Frontiers in neuroscience*, 13, 1321.

Trebbi, F., P. Aghion, and A. Alesina (2008): "Electoral rules and minority representation in US cities," *The Quarterly Journal of Economics*, 123, 325–357.

Van der Windt, P., M. Humphreys, and R. S. de la Sierra (2018): "Gender quotas in development programming: Null results from a field experiment in Congo," *Journal of Development Economics*, 133, 326–345.

Vogt, M., N.-C. Bormann, S. Rüegger, L.-E. Cederman, P. Hunziker, and L. Girardin (2015): "Integrating data on ethnicity, geography, and conflict: The ethnic power relations data set family," *Journal of Conflict Resolution*, 59, 1327–1342.

Wagner, W. (2014): "The Overstated Merits of Proportional Representation: The Republic of Macedonia as a Natural Experiment for Assessing the Impact of Electoral Systems on Descriptive Representation," *Ethnopolitics*, 13, 483–500.

Wright, R., P. Kircher, B. Julien, and V. Guerrieri (2021): "Directed Search and Competitive Search Equilibrium: A Guided Tour," *Journal of Economic Literature*, 59(1), 90–148.

Wucherpfennig, J., N. B. Weidmann, L. Girardin, L.-E. Cederman, and A. Wimmer (2011): "Politically relevant ethnic groups across space and time: Introducing the GeoEPR dataset," *Conflict Management and Peace Science*, 28, 423–437.

Yin, J. and J. Wang (2014): "A dirichlet multinomial mixture model-based approach for short text clustering," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 233–242.