For Whom Does the Phone (Not) Ring? Discrimination in the Rental Housing Market in Delhi, India*

Saugato Datta† Vikram Pathania‡

Abstract

Using an audit experiment carried out using one of India's largest real estate and rental websites, we find strong evidence of discrimination against Muslim applicants, both in terms of probability of being contacted and the number of contacts, relative to upper-caste Hindu (UC) applicants, in the rental housing market in Delhi and its largest suburbs. While the probability that a landlord responds to an upper-caste applicant is 0.35, this is only 0.22 for a Muslim applicant. We also find suggestive evidence that when landlords respond to both UC and Muslim applicants, they call back the UC applicant sooner. Muslim applicants are especially disadvantaged when applying to rent one-bedroom houses; there is an additional 20% points reduction in the probability of a callback. In contrast, we find no clear evidence that landlords are less likely to respond to Scheduled Castes and Other Backward Classes. However, a bounding exercise suggests that for both these groups, our estimates may understate the true differentials in callback ratios as a result of our failure to perfectly link all callbacks to a listing.

JEL classification: I3, J7, R21

*We gratefully acknowledge UNU-WIDER for its generous grant that enabled this research project.
†Managing Director, ideas42, 80 Broad Street, New York, NY, USA. Email: saugato@ideas42.org
‡Department of Economics, University of Sussex. Email: v.pathania@sussex.ac.uk

1 Introduction

The persistence of disparate educational, economic, financial and health outcomes between social groups defined along racial, ethnic, gender or religious lines has been a long standing concern for public policy in many countries across the world. Scholars in fields ranging from anthropology to sociology to criminal justice have consequently employed a variety of disciplinary perspectives to attempt to understand the forces - whether legal, social, institutional, behavioural or political - that generate and sustain such gaps in outcomes between members of different social groups (Dohan 2003; Hirsch 1983; Lamb 2005; Pager and Shepherd 2008)

While acknowledging the possible role of innate or learned differences in ability or preferences between groups, discrimination against marginalized or historically disadvantaged groups by members of more powerful groups is widely held to be a plausible source of at least some (and sometimes a major) part of observed inter-group differences. Indeed, at least since Becker (1957), economists have sought to understand, establish, and quantify the possible role of discrimination to explain relatively worse outcomes among members of marginalized/disadvantaged groups, racial or ethnic minorities, and women. The outcomes studied include employment, wages, access to credit, job performance, housing patterns, occupational choice etc (Aigner and Cain 1977; Altonji and Blank 1999; Arrow et al. 1973; Bertrand and Mullainathan 2003; Blau and Kahn 2006; Blau and Ferber 1987).

Within this broad area of research, a growing literature focuses on the existence, drivers and magnitude of discrimination in housing markets (Choi et al. 2005; Galster 1991; Turner et al. 2002). Motivated at least in part by a large literature that links housing discrimination to a variety of adverse social and economic consequences for marginalized groups, including worsening residential segregation (Denton 1999; South and Crowder 1998), poorer educational and employment outcomes (Yinger 1995), and lower rates of saving (Kain and Quigley 1972), this body of research seeks to quantify the extent of such discrimination and understand what drives it. These thus emerge as key questions for further research.

As with the study of discrimination in other domains (see for instance, Bertrand and Mullainathan (2003), and Banerjee et al. (2009) for labor markets), the difficulty of clearly attributing some or all of the observed differences in housing outcomes between groups to discrimination against individuals belonging to those groups using observational data alone has led economists to use experimental methods, including audit studies, to more clearly identify the role of discrimination.[1] While the bulk of such audit studies have been carried out in developed countries (Ahmed and Hammarstedt 2008; Andersson et al. 2012; Choi et al. 2005; Yinger 1995), a small but growing number of audit studies have by now been carried out in developing countries.

In India, discrimination is a salient issue for policy. There are persistent differences in key outcomes between various social groups differentiated by religion as well as caste ? historically the primary axis of social differentiation in South Asia (Beteille 1992). A growing body of literature has

[1]See Fix and Struyk (1993) for an overview of this research for the United States.

begun to quantify caste based discrimination in labour markets (Banerjee et al. 2009; Deshpande 2011; Siddique 2011; Thorat and Neuman 2012). There are several historical and anthropological accounts of enforced housing segregation along caste lines (Beteille et al. 1969; Dumont 1980; Ghurye 1961; Srinivas 1957). There is considerable anecdotal evidence, some of which has received widespread media coverage, about the existence of discrimination against Muslims, the country's largest religious minority.[2]

However, until recently there has been little experimental evidence on the extent to which religious and caste minorities in India experience housing discrimination. To the best of our knowlegde, there is only one published experimental study of housing discrimination in India. Thorat et al. (2015) use a variety of audit techniques to document the differential treatment of Muslims and Scheduled Castes in the housing market of India?s national capital region. We conduct a web-based audit of the market for rental properties offered directly by owners/landlords using a sample of 170 rental properties in the Delhi region. Our findings complement and extend the results in Thorat et al. (2015) by adopting a different experimental strategy to generate what we argue are likely cleaner measures of differential treatment attributable only to perceived differences in caste and religion.

First, rather than employ a face-to-face or telephonic audit as in Thorat et al. (2015), we use a strictly impersonal, web-based approach that involves no interaction whatsoever between the fictitious tenants and prospective landlords. Following Bertrand and Mullainathan (2003) and Banerjee et al. (2009) we argue that this generates greater confidence that any measured differences are not driven by investigator or enumerator effort, affect, presentation, etc. Our study should thus provide more accurate estimates of the effect of caste/religion than a telephonic or face-to-face audit.

This choice is not costless: we can only use measures based on landlords' call patterns, unlike Thorat et al. (2015) who can measure, for example, the kinds of properties prospective tenants are steered towards. Given the trend globally to generate data on discrimination using the kind of impersonal audit we use here, however, we argue below that this tradeoff is worth making. At the very least, our results should be seen as extending those in Thorat et al. (2015) by providing a cleaner, albeit more limited-scope, measure of discrimination in housing markets that can be compared to their results from in-person and telephonic audits.

Second, while the "resume audit" literature typically restricts itself to measuring whether or not landlords contact potential tenants from different social groups at different rates using a binary variable, we can also measure how many times, in what order, and at what intervals landlords contact tenants. Our experimental set-up thus permits us to derive plausible continuous measures of landlord interest, effort, and persistence in calling back tenancy applicants of different castes

[2] See "They Said We Don't Give Flats to Muslims, Alleges 25-Year-Old Woman in Mumbai", http://www.ndtv.com/india-news/they-said-we-dont-give-flats-to-muslims-alleges-25-year-old-woman-in-mumbai-766353, accessed on May 27, 2015; and "Denied Flat Because She's Muslim, Delhi Academic Asks Kejriwal for Help", http://thewire.in/2015/07/24/video-denied-flat-because-shes-muslim-delhi-academic-asks-kejriwal-for-help-7165/, accessed on July 24, 2015.

and religions.

Third, we are able to estimate whether the extent of measured discrimination varies by neighbourhood quality, apartment size, and landlord background, subject to the usual caveat of diminishing statistical power.

We find strong evidence of discrimination against Muslim applicants, both in terms of probability of being contacted and the number of contacts. Where the probability that a landlord contacts an upper-caste applicant is 0.35, this is only 0.22 for a Muslim applicant. A Muslim applicant must respond to 45.5 listings to receive 10 landlord callbacks, while an UC applicant must respond to only 28.6 listings to receive the same number. A similar pattern obtains in the case of the number of callbacks. A Muslim applicant would need to send about 21 expressions of interest to get 10 callbacks, whereas an UC candidate would only need to send just over 12. Both these results point to significant disadvantages faced by Muslim applicants relative to upper-caste Hindus, who must expend significantly more effort to find housing. However, we fail to find statistically significant evidence of bias against Scheduled Castes (SC) or Other Backward Classes (OBC).

However, a bounding exercise suggests that, our estimates may understate the true differentials in callback ratios as a result of our failure to perfectly link all callbacks to a listing. Finally, we find some suggestive evidence that landlords wait longer to call Muslim (and to a lesser extent Scheduled Caste) applicants back after receiving a query from them than they do to call back upper-caste Hindus.

We also find some heterogeneity in landlord responses to applicant types. First, female landlords are about 13.4% points less likely to respond to an SC applicant. Second, landlords offering 1-bedroom properties are 20% points less likely to respond to Muslims. As a rule, applicants to 1-bedroom properties tend to be single men or women. Since all our applicants are male, this implies that the housing rental market is especially hostile to single Muslim men. Third, landlords offering more expensive properties are 11% points less likely to respond to OBC applicants.

The plan of the rest of the paper is as follows. Section 2 draws upon a large literature in economics, social anthropology and law to motivate studying housing discrimination in general. Section 3 considers the case of India, arguing that housing discrimination is an important policy question in general but particularly in India, where there is both a history of enforced housing segregation by caste and contemporary anecdotal evidence of discrimination against religious minorities and some caste groups. Section 4 situates this study within the literature on this issue in India. Section 5 describes the audit exercise. Section 6 presents key descriptive statistics. Section 7 describes the key results. Section 8 discusses their significance, the implications for future research, and concludes.

2 Literature Review: Housing Discrimination Globally

Besides labor and credit markets, a large and growing sub-set of the literature on the causes of divergent socio-economic outcomes between social groups focuses on the existence, drivers and magnitude of discrimination in housing markets (Choi et al. 2005; Turner et al. 2002). The theoretical justification for this literature arises in the first instance from considerations of fairness and the right of individuals and households to live where they choose (Danziger and Lin 2000; Massey and Denton 1993) but also from policy imperatives arising from the recognition of the effects of discriminatory housing practices on a variety of outcomes. Given existing spatial inequalities in the quality of public services in most countries, housing discrimination limits the ability of individuals and households from disadvantaged groups to access quality schooling and healthcare, which in turn affects their schooling, labor and credit market outcomes, and could therefore lead to the persistence of measured inter-group inequality (Galster 1991; Yinger 1995).

As with other forms, of discrimination (e.g., discrimination in labor and credit markets), it is difficult for researchers to infer unequal treatment in the housing market from aggregate data alone, since some relevant characteristics of applicants may be visible to potential employers, lenders or landlords but unknown to the researcher. As a result, studies that decompose observed outcomes by observable applicant characteristics risk finding biased results, where a portion of the variation that is in fact due to some characteristic of the applicant unobservable to the researcher but observable to the landlord or interviewer is erroneously attributed to discrimination (Altonji and Blank 1999). Econometrically, such studies are the canonical case of omitted variable bias.

Efforts to mitigate this bias have to a variety of econometric techniques, such as the use of instrumental variables, etc. Beyond this, researchers studying discrimination have used quasi-experimental approaches. An influential quasi-experiment in the field of workplace discrimination is Goldin and Rouse (2000), which uses the effect of the introduction of blind auditions into orchestra hiring on women musicians' hiring outcomes to measure gender discrimination.

However, the dominant response of researchers - whether in labor, credit or housing markets - to the problems with observational studies is to turn to experimentation. Experimental approaches in the field of discrimination have centered on audit studies. In the area of labor discrimination, researchers have used hiring audits, where comparable minority and non-minority candidates are sent to actual interviews to measure the existence and extent of differential treatment (Altonji and Blank 1999). These studies are closely related to "mystery shopping" studies of various kinds (e.g. Mullainathan et al. (2012) on the market for financial advice) where trained actors are enlisted to model various characteristics or needs for a service provider.

The counterparts of these studies in the case of housing markets are housing audits, where researchers send actors from different groups of interest who are trained to otherwise present as identical and follow pre-set scripts, in order to measure differences in the apartments or houses they are shown, rental prices quoted, and of course ability to rent. Such systematic audit or correspondence studies have played an important role in enabling researchers to measure the extent of discrimination against minorities in housing markets in the United States (Choi et al. 2005; Yinger 1995) more credibly than would be possible using aggregate data. In the United States in particular, these studies have been institutionalized by the Department of Housing and Urban Development which sponsors regular housing audits in many large metropolitan areas in order to be able to track temporal and spatial changes in minorities? outcomes in the housing market (e.g., Turner et al. (2002)).

Audit studies clearly offer many advantages over cross-sectional analyses. However, as Bertrand and Mullainathan (2003) point out in the context of labor market discrimination, they have at least one major shortcoming. The fact that actors are aware of the purpose of the experiment (i.e., that the experiment is not and cannot be double-blind) increases the inherent difficulty of ensuring the absence of observational differences between mock job candidates in a pair and could therefore introduce error into the estimates of discrimination. For instance, actors of different races may (perhaps subconsciously) behave differently, triggering different behaviours on the part of the landlord than if they had truly been "identical".

These shortcomings have led researchers to devise what we will call "impersonal" or "automated" audit studies (with resume audits a la Bertrand and Mullainathan (2003) being a prime example), where researchers vary the perceived group membership of a fictitious applicant (for a job or an apartment) without enlisting actors to interact with the potential employer, landlord, or other service provider. These studies rely on the ability to apply remotely for jobs or apartments using mail or internet-based applications. As Bertrand and Mullainathan (2003) and Banerjee et al. (2009) note for resume audits, these studies are able to provide the cleanest possible evidence of differential treatment but only for a much earlier stage of the hiring process. In the labor-market case this means that they can only measure whether a potential employer called an applicant back, and not whether the potential employee was in fact offered a job, or what salary was offered. Nonetheless, as Bertrand and Mullainathan (2003) note, "to the extent that the search process has even moderate frictions, one would expect that reduced interview rates would translate into reduced job offers".

The growth of classifieds and internet-based housing portals, which make it possible for a potential tenant to express interest in an apartment or house offered for rent without interacting with the putative landlord either in person or over the telephone, has made it possible to implement such "remote" audits in the housing market. Ahmed and Hammarstedt (2008) and Andersson et al. (2012) are prominent examples of such studies in a developed country context.

As in the case of labor-market discrimination studies, such audits in the case of housing markets must restrict themselves to measuring an early-stage outcome (i.e., whether the landlord applied to responded to a potential tenant's expression of interest) without being able to measure differences in downstream outcomes such as what apartments were offered or the terms on which apartments were made available. Nonetheless, as in the case of resume audits, impersonal housing audits do provide the cleanest possible evidence of differential treatment of different categories of people seeking housing of all study methods. And, building upon the argument in Bertrand and Mullainathan (2003), we argue that the existence of search frictions means that those from groups that receive fewer landlord callbacks are also likely to be offered fewer houses, or to take longer to find a house that meets their requirements.

In what follows, we argue for the relevance of these considerations in the study of housing markets in urban India, and put the contributions of the present study in the context of the existing literature on this topic.

3 Housing Discrimination in India: Theoretical Considerations, Historical Experience

There are several reasons that suggest, a priori, that discrimination might impinge upon the results of certain categories of individuals' or families' quests to rent or buy housing in India. Further, it appears ex ante likely that such discrimination may be faced both by individuals who occupy a historically subordinate position within the hierarchy of caste, the complex system of hierarchical social relations that long governed, and in important ways still governs some aspects of Indian society (notably marriage, see Banerjee et al. (2013)), and/or by Muslims.

First, the strong correlations between caste and religion and socio-economic status, occupational choice, housing outcomes, etc. seen in the data suggest a possible role for discrimination against members of these groups along some or all of these dimensions.[3] In general, upper-caste Hindus have better economic outcomes than both non-upper-caste Hindus (including but not limited to Scheduled Castes and the Other Backward Classes or OBCs) and Muslims.[4]

Using data from India's National Sample Survey, a nationally representative repeat cross-sectional household survey, Deshpande (2005) provides a detailed account of levels and patterns of consumption expenditure (used as a proxy for income in the absence of reliable income data) by social group over a twenty-year period. She finds that Scheduled Castes/Tribes (SC/ST) have lower per capita consumption expenditure than other groups, and while each social group's consumption expenditure has risen over the twenty-year period studied, the increase is lower for SC/ST than the rest of the population. Deshpande (2001) uses five indicators of standard of living (land holding, occupation, education, ownership of consumer durables, and of livestock to construct a "Caste Development Index" (CDI), and finds that in the early 1990s, there was no state where the CDI of SC/ST populations was higher than that of non-SC/ST populations. Desai and Dubey (2012) analyze data from a nationally representative survey of 41, 554 households conducted in 2005 to argue that there continue to be persistent disparities in education, income and social networks by caste. There are also persistent disparities in housing quality. According to the 2011 Census, only 34% of SC households have a latrine on the premises, compared with 46.7% of all households in the country (Office of the Registrar General and Census Commissioner 2012).

The existence of affirmative action for members of the Scheduled Castes, Tribes and Other Backward Classes necessitates official data-collection about these groups. Since there is no such national policy for Muslims, data on Hindu-Muslim differentials are more sparse. Nevertheless,

[3] The argument here draws heavily on Banerjee et al. (2009).

[4] Apart from a small (albeit culturally and economically powerful) elite descended from the Arab- or Persian-origin dominant ruling and administrative classes of medieval and pre-colonial India, most Indian Muslims are descended from members of the lower Hindu castes who converted to Islam to escape discrimination or to better their economic prospects (but who, in many cases, continue to follow the caste-ordained occupations of their Hindu forebears).

what data there are is suggestive. For example, while the widespread prevalence of disguised unemployment or underemployment in developing countries limit the utility of official measures of unemployment, it is striking that both Scheduled Castes and Muslims are over-represented among those whom the Indian state classifies as "marginal workers" - those employed for less than 6 months of the prior year. This figure is 10.9% for Scheduled Castes and 6.5% for Muslims, while it is 3.5% for the country as a whole (Banerjee et al. 2009). According to the Sachar Committee, a government committee set up to probe the socio-economic status of Muslims in India, Muslims had the highest unemployment rate of any socio-religious group in India (Sachar 2006). Shariff (1995), one of the few analyses of socio-economic differentials between Hindus and Muslims in India, finds that 22% of Muslims had a monthly household per capita expenditure of less than Rs.110 in 1987-88, compared with 13.1% of Hindus. It is worth noting that given well-documented caste differentials within the Hindu population, the difference between Muslims and the non-Scheduled Caste/Tribe Hindu population would likely be even higher.

Secondly, the nature of caste and religion in India lend strong a priori plausibility to the idea that members of lower-caste and minority groups are likely to face discrimination in the housing market in particular. Along with restrictions on commensality, inter-marriage, education and occupational choice, housing segregation is central to the logic of caste (Beteille et al. 1969; Dumont 1980; Ghurye 1961; Srinivas 1957). Just as there were restrictions on inter-dining and inter-marriage between members of "higher" and "lower" castes (particularly those belonging to the groups formerly referred to as Untouchable and, since Independence as Scheduled Castes), housing in most Indian villages and towns was organized along caste and community lines to conform with strongly held notions of purity and pollution where the presence of certain groups within a certain distance of higher-status groups was held to be polluting. For example, most Indian villages had hamlets occupied by members of specific caste groups (such as the "Agraharam", or the "Brahmin" quarter of traditional Tamil villages and towns) with those lower in the social hierarchy often being prohibited from entering the parts of the village occupied by those ostensibly of "higher" status, and from using the same places of worship, sources of drinking water, or other public facilities as the latter.

Thirdly, while modernization and urbanization have loosened the ties of caste and clan so that large Indian cities contain many neighborhoods that are, at least in theory, available to anyone who can afford to live there, there is considerable anecdotal evidence that members of India's Muslim community (and, to some extent, other religious minorities), some linguistic groups, and those belonging to the Scheduled Caste or Other Backward Classes (historically disadvantaged groups whose members are eligible for some forms of positive discrimination in education and public-sector employment) continue to face difficulties in accessing the housing of their choice. For example, there have been a series of media reports that present cases where even upwardly-mobile, middle-class, professional or elite Muslims face discrimination when they look for housing, even in India's biggest cities.[5]

[5] See "Shutting Out the Other", http://www.thehindu.com/features/magazine/task-of-house-hunting-in-india-becomes-complicated-if-you-dont-belong-to-the-right-religion/

And while relatively little media attention focuses on SCs and OBCs, it is plausible that Scheduled Castes and Other Backward Castes - both groups whose members have historically been considered "unclean" or "inferior" by dominant caste groups, and who were traditionally not allowed to live in the same areas as those from more powerful group - would face similar hurdles. For example, preliminary results from a recent nationally representative household survey of over 40,000 households show that 52% of Brahmins and 24% of non-Brahmin "upper-caste" households practiced untouchability either directly, or in that they were hesitant to admit a member of the Scheduled Castes into the kitchen (NCAER 2014). Indeed, Vithayathil and Singh (2012) find that segregation by caste is greater than that by class in all seven Indian mega-cities that they study.

4 Existing Literature on Housing Discrimination in India

As in other countries, spatial inequalities in the provision of public services (e.g., schools, hospitals, roads, etc.) and the signalling role of an individual's address mean that being unable to access the housing and area of one's choice is (quite apart from being unjust) a cause of other persistent gaps (such as in educational attainment, health status, and employment status. Housing discrimination is thus a policy concern. Reliable evidence on housing discrimination in India emerges as a matter of importance if these problems are to be addressed.

However, although Vithayathil and Singh (2012) argue for the relevance of audit studies to the research agenda on urban housing and patterns of residential segregation in India, there is a surprising paucity of rigorous empirical evidence on this issue, with Thorat et al. (2015), which experimentally measures the extent of discrimination against non-upper-caste Hindus and Muslims in the city of Delhi and its suburbs using face-to-face and telephonic audits, being a notable exception.

Our study builds on this study but differs along some important dimensions, which we discuss in detail later, but which have to do both with the nature of the experiment we implement and the measures of landlord interest and effort we are able to capture and analyse. Briefly, our audit methodology, manipulates perceived applicant identity through the names used on web-based applications for rental housing to ensure that inadvertent and unconscious experimenter effects do not bias the results of the audit, thus side-stepping some lingering concerns with audit studies that rely on live subjects and/or telephonic conversations with landlords.

Secondly, we exploit the fact that landlords can and do call back potential tenants multiple times, and that we can track the times of their calls, to develop measures of landlord interest and effort. For instance, we can measure not just whether landlords were more likely to call back upper-caste tenants than Muslims or Scheduled Castes, but also whether they were more likely to call back the former more frequently, or sooner, or more persistently. The nature of our experiment thus leads to a rich set of observable and quantifiable landlord behaviours which we exploit to deepen our analysis beyond a simple analysis of callback rates. Since we have information on features of

---

`article7286237.ece` accessed on September 18, 2015.

advertised houses, and can infer some limited information about landlords from their names, we can also test for interactions between landlord and apartment characteristics and callbacks.

5 The Experiment

This section describes our experimental strategy in greater detail.

5.1 Location and Sample

Our experiment, which was carried out entirely remotely, exploited one of India?s most popular online housing search platforms. Over the course of a roughly two-month period in the summer of 2015, we regularly scanned the most recently posted rental listings for Delhi and its two largest contiguous suburbs on this website and identified a convenience sample of 171 listings posted directly by landlords (i.e., not ones posted by an agent from a rental or property agency), taking care to avoid sending more than one set of applications to a given landlord in order to avoid arousing any suspicions that would result if a landlord received an application in two rounds of sending and noticed that the same number was now attached to a different name.

The landlords in our study were seeking tenants for apartments or houses in India's capital and second-largest city, Delhi, and its two largest contiguous suburbs, Gurgaon in the neighboring state of Haryana and the New Okhla Industrial Development Area, usually referred to as NOIDA, in the neighboring state of Uttar Pradesh.

These three administrative units form the core of what is known as the National Capital Region (or NCR), which is envisaged by urban planners as eventually constituting a single commuter zone centered on the national capital, Delhi. While the entire NCR extends in a wide arc around the city and is not yet a fully realized vision, the three areas in our study are in many ways a single economic entity since they are connected by the same rail-based mass transit system, the Delhi Metro. The development of the Metro and the growth of many service- and manufacturing-sector industries in NOIDA and Gurgaon have led to a substantial population, particularly of middle-class white-collar workers, who live in one city of the three and work in another. For the purposes of our study, it is thus reasonable to think of these three administrative units as constituting a single housing market with several sub-regions (South, East, West and North Delhi, Gurgaon and NOIDA to a first approximation).

Although the behavior of agents is interesting in its own right, we chose to focus on landlords for this initial study since doing so avoids any potential issues with principal-agent problems arising, for example, from agents having imperfect information about landlords' preferences.[6] While no explicit attempt was made at representativeness, the final counts of apartments applied to in each part of the city should be broadly indicative of rental housing flows in those areas, since we sampled the most recent landlord-posted advertisements in each region.[7]

[6] We intend to further explore both landlord and agent behavior and possible differences between them in a companion study, currently in the pilot phase.

[7] We did attempt to over sample Muslim landlords. Since the flow of such advertisements is very sparse, in practice

5.2 Names and Contact Strategy

As discussed above, past research suggests that both Scheduled Castes and Muslims, as well as individuals belonging to the type loosely known as Other Backward Classes (OBCs), may face discrimination in many aspects of life in India (see Deshpande, 2006). We therefore chose to send applications from fictitious tenants belonging to four social categories - Upper-Caste Hindu (UC); Muslim, Scheduled Caste (SC) and OBC.

We sent four queries to each landlord we contacted. Queries were sent using the online web form (see Appendix 1 for a sample of the form in which the landlord received the web query) with one fictitious candidate from each of these four categories applying to each landlord. Care was taken not to send a flurry of applications, but rather to apply sequentially with time gaps between applications. We randomised the order in which we sent the applications so that each social type was equally likely to be first, second, third, or fourth to apply to a given landlord.

We used two names for each type, with both having common male first names beginning with the letter "A" and last names that denoted caste and religion (with Muslim applicants having both first and last names that denoted their religion, since Muslims have distinctive first and last names).[8] All applicants were male. For our Scheduled Caste and Other Backward Caste names, we relied on government lists of the castes included in these categories in Delhi and its surrounding Hindi-speaking states for maximum signalling value. We also drew on previous qualitative work carried out by Banerjee et al. (2009); they field-tested possible SC, Muslim and OBC names to identify those most widely recognized by middle-class residents of Delhi.

Queries were identical except for the name and email of the applicant, which reflected the assigned social type, the associated cellphone number (which was mated to type to aid clear assignment to social type) and the time and intra-ad order at which the query was sent which, as we discuss above, was randomized to avoid order effects.

In all, we report results from landlord responses to 681 unique applicants to 171 apartments. While our design was fully blocked, there was one instance where the listing was deleted in the midst of sending the four applications. As a result, we sent OBC queries to 171 listings but UC, SC, and Muslim queries to 170. Because of this almost fully blocked design, we do not need to worry about balance across social categories when it comes to listing and landlord characteristics.[9]

5.3 Data and Analysis

Apart from the details on the listings (e.g., square footage, monthly rent, location, etc.), our key data come from whether, how often, and at what intervals the landlords to whom we sent queries called our fictitious tenants.

that meant we sampled most of those landlords during the study period.

[8]The Hindu first names were Anil, Arun, and Amit; the Muslim first names were Abbas and Arif. The OBC last names were Yadav and Ahir; the UC last names were Gupta and Sharma; the SC last names were Paswan and Manjhi; and the Muslim last names were Khan and Ahmed.

[9]Later, we present a table showing the balance across applicant categories.

Calls were received on cellphones carrying Indian SIM cards procured for this experiment. Each SIM (i.e., each number) was mated permanently to a particular type (e.g., SC or Muslim). Call log data was downloaded into Excel and coded into STATA to enable the analysis. Keeping count of the number of callbacks and number of unique callers to each type of applicant was thus trivial.

We did not answer phone calls, so we do not have any direct way of gauging what landlords were calling to say (except in cases where landlords also sent either a text or email, which was only in a small minority of cases[10]). In many cases landlords called back multiple times, allowing us to measure not just whether they called but how many times and how soon after the receipt of an online application expressing in an interest in their apartment landlords called our fictitious ?tenants?.

However, it is not enough for our purposes to track simply the total number of unique numbers that called back our different categories of tenants. There are two reasons for this. First, we need to link callbacks to listings in order to analyze the effects of landlord and listing characteristics. Secondly, there is a distinct possibility (borne out by our subsequent findings) that not all calls received were in response to the queries we sent out, or at least not directly.

To elaborate, there was a strong possibility that some calls were spam, telemarketing calls, or calls from brokers or agents whom we had not contacted but who had somehow got hold of our fictitious tenants' numbers and were calling to offer them their services. While including such "spurious" calls in raw counts is not entirely uninformative (after all, it would be striking if spammers, too, were more interested in some social categories than others), we need to focus our estimates on genuine calls from landlords we applied to in order to make reliable inferences about landlord behavior. In addition, calls that we cannot link to a listing cannot be used in our regression estimates.

5.4 Tracing Callers

The difficulty of linking callbacks to listings was exacerbated by the fact that while the majority of listings do not list the landlord's actual cellphone or land-line number. Rather, the housing portal assigns them a masking number, presumably to protect themselves from spam and to protect their privacy.[11] The portal?s practice is similar to that employed by online classified services in the US, such as Craigslist, which provide users with a masking email so as to protect their privacy. Potential tenants (and researchers) thus mostly see only a specially assigned number, dialling which connects them to the advertiser's actual number, which they (and we) do not see. In practice, this meant that except for a small subset of those we applied to, we could not immediately link a callback to a specific listing, since the callback came from the landlord's actual number, and not the one on the online listing, which is essentially a call-forwarding service.

[10] It is worth noting that all texts or emails were simply requests for the prospective tenant to get in touch to discuss the listing, so were "affirmative callbacks" in the sense that the landlord was contacting the prospective tenant to move the discussion further. No landlord emailed or texted to deny a prospective tenant a shot at his/her apartment.

[11] A reasonable concern, given that we found that many people we had not contacted had nevertheless got hold of the phone numbers we employed in our experiment.

However, the use of web-based call-tracking resources such as Truecaller[12], publicly-available information, and some supplemental calling back allowed us to solve this problem to a great extent.

We received callbacks from 118 unique numbers, whether landlines or cellphones. Of these, we were conclusively able to identify 22 as either telemarketing calls, misdials or wrong numbers, calls from property agents who appear to have gained access to our number, or (in one case) to a homeowner other than one in our sample. We drop such spam calls (and the associated calling numbers) from further analysis since none of these categories are informative about the preferences and effort of the landlords in our sample.

Of the remaining 96 unique (potentially legitimate) numbers, we are able to conclusively establish that 89 (or 92.7%) correspond to a listing we applied to. We do this in four ways. First, some landlords do list their actual numbers on the listing, making it trivial to link their callbacks to a listing. Second, some landlords also emailed and identified themselves (and provided a number). Third, Truecaller and other web-based search engines, together with the call logs that showed us when a call from a given number was first receive, enabled us to trace a large number of the remainder. Finally, we carried out an intensive period of calling hitherto unidentified numbers back about 10 days after the conclusion of the experiment (and about a month since most numbers had first called) where we attempted to determine whether these numbers were in fact landlords we had applied to, and if this was indeed the case, which listing each number corresponded to.[13]

The success of our tracing attempts makes us confident in asserting that any results we see based only on traced calls (which are the only ones we can use in our regressions) are not driven by differential success in tracing callers to different categories of applicants. It is worth noting that of the 7 numbers we were unable to either trace to a landlord or tag as spam, 4 called only once. Only 3 numbers that called our experimental numbers more than once (a mere 2.5% of the 118 numbers that called us at any point) are thus ones that remain untraceable and unclassifiable. Nonetheless, to ensure that our results, which are based on traced calls only, are not driven by the call patterns of these 7 untraceable numbers, we run a bounding exercise on our callback ratios to see how they would change under assumptions chosen to go against our hypothesis.

6 Descriptive Statistics: Listings, Landlords, and Applicants

6.1 Listing and Landlord Characteristics

Table 1 provides an overview of the features of the properties in the sample. A large majority of listings (71%) were for two- or three-bedroom apartments, with 20% were for one-bedroom properties, and 9% had four bedrooms. The distribution of properties differed somewhat between the city and suburbs, with fewer one-bedroom flats in the suburbs. As we would expect, city flats were about one-and-a-half times more expensive per square foot (Rs. 28.5 psf compared with Rs.

[12] `https://www.truecaller.com`

[13] For this, we used Skype numbers set up to show our experimental numbers as Caller IDs to dispel any concerns call recipients might have about a non-India number showing up as calling about a flat in Delhi.

18.1 psf in Gurgaon or NOIDA), and were smaller on average (at a little over 1100 sf, compared with an ample 1600+ sf in the suburbs). We also report those landlord characteristics that we were able to discern from names with a high degree of certainty, which were religion and gender. About 12% of the landlords we applied to were Muslim, the lion's share of these in the city. This is slightly higher than Muslims' share in the city's population, which is estimated at around 10%.[14] About 13% of our landlord sample was female.

6.2 Applicant and Application Characteristics

Table 2 provides an overview of application characteristics. Since our design was fully blocked, with each listing receiving one application of each type (UC, SC, OBC and M) there can be no differences in the proportion of each type applying to any kind of apartment or landlord. As discussed earlier, the position of each type within the applicant pool for each listing was randomized to avoid order effects, since it is possible that an applicant who applies before others may have an advantage. As we see in Column (1) there are negligible differences in mean position of types within a listing.

As Columns (2) and (3) show, the staggered nature of our application procedure does introduce some variation in when within the day or week a particular kind of applicant applied to rent a property. OBC and SC applications were more likely to have been sent over the weekend and outside of daytime hours in India. The likely effects of these differences are unclear. Perhaps the most likely effect is on time to first response: it seems reasonable that a query sent within office hours or on a weekday would be acted on immediately - although the opposite is also possible, since this is personal work for landlords and may in fact be harder to attend to during work hours or on a workday. Indeed, queries received during work hours to be more likely to be missed due to inattention. Thus, while it is not clear whether and how these differences in sending time matter, we control for them in some specifications. Our findings will be informative for future audit studies, which could seek to maximize response rates by optimizing sending times. Finally, we expect that applications to older listings would, ceteris paribus, elicit less response. Again, due to the fully blocked design, the average gap between the application date and the listing date is about the same across the four categories.

7 Results

it is instructive to first discuss the patterns in the raw counts of calls, callers, and all responses (inclusive of emails and texts.)

Table 3 displays raw counts of responders and responses by applicant type. Column 1 of Panel A presents the count of all unique calling numbers by type, and appears to suggest that there is a large callback differential between UC applicants and all other types. However, our tracing exercise tempers these conclusions somewhat. As we can see from the move between Column 1 and Column

[14] As noted before, the actual share of Muslim landlords in the flow of listings was much smaller; we oversampled the Muslim landlords.

2, part of the UC-vs-Others differential in Column 1 is due to the fact that our UC applicants were spammed more.

This is, of course, interesting in its own right, particularly because some of the spam callers were in fact property agents whom we did not contact, but who contacted our fictitious tenants independently, while others presumably wanted to offer them good or services. We do not explore this issue further here, although we flag it for future research. For our purposes, we must exclude those numbers we identify as belonging to unsolicited callers (uncontacted brokers or landlords, telemarketers, etc.) from the analysis. This substantially narrows what initially seems to be a large callback differential between UC and SC, for example.

After we exclude the spam callers, we are to match most of the remaining calling numbers to landlords in our sample, as discussed above. Column (3) contains the counts of the landlords we could trace among the callers. Untraceable numbers are not the only source of the differences between Columns 2 and 3 in Panel A of Table 3. It is worth noting that 13 landlords called from two numbers each, explaining why the numbers in Column 3 are so much smaller than those in Column 2. Once Column 3 is supplemented by information about landlords who texted or emailed, we get the total number of landlords who responded in Column (6). These numbers are our fundamental raw data for the results presented in the remainder of this section. As Column (6) shows, roughly similar numbers of landlords respond to upper-castes and SC applicants. Somewhat fewer respond to OBCs, and substantially fewer to Muslims.

Panel B of Table 3 displays analogous counts for the number of calls, emails, and texts. Since we did not answer calls, landlords (or others) could persist in calling applicants. In many cases they did, leading to a much larger number of calls than callers. The relevant numbers are, again, Column 6, which measures the total number of non-spam contacts received by each type. Here, we continue to see a big difference between UC and Muslim, with UC applicants receiving 60% more calls than Muslim applicants. We also see a substantial difference between UC and SCs here, despite there being no difference at all in the number of callers: UCs receive 18% more calls than SCs, despite being called by almost exactly the same number of landlords. Finally, we should note that while we see that OBCs receive more calls than even UCs, this is driven entirely by one landlord who called only the OBC applicant 33 times.

Result 1: Landlords are significantly less likely to respond to Muslims applicants.

Our central result can be seen from Table 4, which presents differences in mean response rates (defined as the percentage of landlords who either called, texted, or emailed an applicant) between our UC applicant and others. While the probability of a landlord responding to the UC applicant is 0.35, the corresponding probability is 0.22 for Muslim applicants. The difference of 0.13 is statistically significant at conventional levels of significance. A simple way to scale this difference is by calculating the number of landlords each type must contact in order to have a pool of 10 apartments to consider. Whereas an UC applicant needs to send out just under 29 queries to hear from 10 landlords, a Muslim applicant needs to send out nearly 45 queries to achieve the same degree of interest. Muslims must therefore expend considerably greater time and effort, including search time, to have access to a similar-sized pool of potential rental properties as upper castes.

The regression counterpart of these results can be found in Table 5, which presents OLS regressions for the probability of being called back at the applicant level. The coefficients of interest are those on the dummy variables for each applicant type. Once again we see that a Muslim candidate is about 12.4% points less likely to be contacted by a landlord than an UC candidate, and that this coefficient is highly statistically significant. This result survives the addition of controls for sending patterns and is essentially unchanged by a specification that uses landlord fixed effects.

Result 2: Relative to UC applicants, Muslim applicants receive fewer responses.

As discussed earlier, the fact that we do not answer calls means that landlords make multiple attempts to contact those applicants they are interesting in pursuing. We can thus use the number of calls and other forms of contact (and not just the number of callers, which has been our key measure so far) as an additional measure of landlord interest.

In Table 4, we present the mean count of landlord responses per application by each applicant type. This number is 0.82 for UCs and 0.48 for Muslims. The difference of 0.34 per listing is strongly statistically significant. Put differently, a Muslim applicant would need to send about 21 expressions of interest to get 10 callbacks, whereas an UC candidate would only need to send just over 12. Muslims must expend significantly greater time and effort to elicit a comparable number of calls.

Table 6 implements regressions using the count of responses rather than merely the probability of being called back. The coefficient on the Muslim dummy is negative and statistically significant, indicating that Muslims get 0.58 fewer responses. Muslim applicants are at a significant disadvantage compared with upper caste applicants.

Result 3: Differences in the probability of response and count of responses between UC and SC /OBC types are not statistically significant.

As the rows for OBCs and SCs in Tables 4. 5, and 6 show, we do find statistically significant evidence of discrimination against these two types as compared to UCs. The probability that a landlord contacts an OBC applicant is 0.30 (see Table 4), which is lower than the 0.35 for an UC applicant. The difference of 0.05 is not, however, statistically significant at conventional levels. The corresponding difference between UC and SC is a trivial 0.01. Almost exactly as many landlords call our SC applicants back as call back our upper-caste applicants.

This pattern is replicated in the corresponding columns for the mean number of callbacks. OBCs have a higher number of callbacks in aggregate (and therefore per listing), but the difference is not statistically significant.[15] The point estimate of the difference between mean number of callbacks for SCs and UCs points to a disadvantage for SCs but is once again statistically indistinguishable from zero.

The regressions in Table 5 confirm this finding. The coefficient on the OBC dummy for the probability of being called back is negative but not significant. The one for the SC dummy is

[15] Recall that one of the landlords called the OBC applicant 33 times.

positive but statistically and economically insignificant. There are also no significant results for OBCs and SCs when it comes to the response count (Table 6).

A caveat to these three results is in order, however. Recall that our regressions and callback ratios are constrained to use only data from the 89 numbers (and emails and texts) which called our experimental numbers that were neither spam/unrelated to our experiment, nor untraceable. We address this concern explicitly through a bounding exercise that we present later.

Result 4: There is suggestive evidence that landlords who respond to both UC and Muslim (or SC) applicants are more likely to call UCs sooner as compared to Muslims (SC).

Table 7 looks at the length of time that elapsed between our fictitious applicant sending an online query to a landlord and the first time the landlord contacted the applicant. Our results, while not statistically significant (in part due to the pairwise regressions only being able to utilize those listings where a landlord responded to both types in the pair), are suggestive. The point estimates suggest that landlords wait about 6.5 hours longer before calling a Muslim applicant than they do for an upper-caste candidate. The results for SC candidates are smaller and magnitude but of the same sign. Both groups, therefore, likely would need to search longer for housing before being able to find a place to rent.

Table 8 displays the findings from a related analysis that investigates whether landlords who respond to two types are more likely to respond first to one type over the other. One hypothesis is that to whom the landlord first responds is independent of the order in which the applications were received. In this case, the null (assuming no bias) is that the proportion of landlords who first respond to a given type in a pair is 0.5 (in half the cases, the landlords should respond to one type, and in the other half to the other type.) A more plausible hypothesis (again assuming no bias) is that landlords simply respond in the same order in which they receive applications. We test the pairwise response patterns against the two nulls in Table 8. As before, since we are restricted to landlords who responded to both types in a pair, we have diminished statistical power. Nevertheless, there is some suggestion that landlords who respond to both UCs and SCs, first call UCs (almost significant at 10%).

Result 5: There is heterogeneity by gender and religion of landlord, and by size and rental price of the listed property in the likelihood of response to the various applicant types

Table 9 displays the results of regressing interactions of landlord and property features with applicant type, on the likelihood of response. Although the small proportion of female or muslim landlords points to limited statistical power, we still find some interesting patterns. First, female landlords are about 13.4% points less likely to respond to an SC applicant. Second, muslim landlords are 22% points less likely to respond at all (the corresponding point estimate for females is about -16% but borderline insignificant at the 10% level.) Third, landlords offering 1-bedroom properties are 20% points less likely to respond to Muslims. As a rule, applicants to 1-bedroom properties tend to be single men or women. Since all our applicants are male, this implies that the housing rental market is especially hostile to single Muslim men. Fourth, landlords offering more expensive properties are 11% points less likely to respond to OBC applicants.

7.1 Call Ratio Bounds

While there are only 7 untraceable numbers, these numbers do not call back the applicant types evenly. It is therefore important to understand the extent to which our result (or no result in the case of SCs and OBCs) is driven by our differential ability to trace callers who called back our different types. Intuitively, if we were relatively less successful in tracing landlords who called UCs relative to tracing those who called SCs, our results would understate the extent of discrimination faced by the latter because our callback ratios would be lower than the true ones - an artefact of our relative success in tracing various kinds of callers.

To address this concern, we carry out a bounding exercise. Consider the simplest case of two types: H and M. In a balanced experiment such as ours, each type applies to the same set of N landlords. Define the following:

- $N_{i,T}$: number of callers for type i that are confirmed as landlords in the sample.
- $N_{i,U}$: number of untraced callers (could be landlords or spammers) that *only* contacted type i.
- C: number of untraced callers common to (i.e. called) both types. Then the number of untraced callers who contacted a type i is $N_{i,U} + C$, and the number of all callers to a type is $N_{i,U} + N_{i,T} + C$

We are interested in the ratio (r) of the number of landlords who called H vs M. If there is no bias, $E(r) = 1$. Due to untraceable callers, we will not observe the true ratio in the experiment. However, we can put bounds on the true ratio by making assumptions on the number of landlords in the three groups of untraced callers - those who only called H, those who called M, and those who called both. Let:

- h be the (unobserved) number of landlords among the untraced callers who only contacted type H. Note $0 \leq h \leq N_{H,U}$
- m be the number of landlords among the untraced callers who only contacted type M. Note $0 \leq m \leq N_{M,U}$
- c be the number of landlords among the untraced callers who contacted both types. Note $0 \leq c \leq C$.

It is trivial to show that for any c, r is highest if we assume $h = N_{H,U}$ and $m = 0$ i.e. all untraced callers who only contacted H are landlords, and all those who only contacted M are spammers. The converse is true for a lower bound on r (conditional on any c). How do these bounds change with c? Consider the upper bound:

$$r^u = \frac{N_{H,T} + N_{H,U} + c}{N_{M,T} + c}$$

This is maximised for $c = 0$. In plain English, count as landlords all the untraced callers who only contacted H, and exclude as spammers all untraced callers who contacted M. The lower bound is minimised by assuming the converse i.e. exclude as spammers any untraced callers who only contacted H and include as landlords any untraced callers who contacted M . Hence the upper and lower bounds on r

$$r^u = \frac{N_{H,T} + N_{H,U}}{N_{M,T}}$$
$$r^l = \frac{N_{H,T} + C}{N_{M,T} + N_{M,U} + C} \tag{1}$$

The results of this exercise, presented in Table 10, are instructive. The first column replicates the counts of responding landlords that we could trace. In column (2), we calculate callback ratios - the ratio of landlords who called back UC applicants to those who called back each of the other types. A ratio higher than 1 indicates that non-UC applicants faced a disadvantage. This ratio is higher than 1 for all non-UC categories. In the case of Muslims, this ratio is 1.55, i.e. the number of landlords who contact UC applicants is 55% more than those who contact Muslim candidates. Columns (3) and (4) of Table 10 present the upper and lower bounds using the expressions in equation 1.

Even if we stack the decks against finding any discrimination against the non-UC types, we see that the minimum plausible callback ratio (UC: SC and UC:OBC) never falls below 1, which would imply equal treatment. In other words, our bounding exercise suggests that we are actually capturing the lower bound of the estimate of discrimination against SCs, which could in fact be higher depending on the true provenance of the 7 untraced calls. The upper bounds for SCs and OBCs are substantially higher, at 1.1 and 1.22 respectively. Given this, we call below for interpreting our "null" result for SCs and OBCs, but particularly SCs, with caution.

A final piece of evidence supports this note of caution. Recall that 22 numbers that called one or more of our applicant categories were dropped from the analysis because they were not from landlords whom we contacted. While some of these were pure spam or telemarketing, we identified 12 of these numbers as belonging either to property agents or a potential landlord offering a different apartment or house than the ones applied for during our audit. There is of course no way to map these unsolicited callers to listings, or indeed to know whether these individuals received information about all 4 of our applicants or only a sub-set. But it is suggestive that while 6 of such agents/landlords called an UC applicant, only one of them called an SC applicant.

We summarise the results of the bounding exercise.

Result 6: A simple bounding exercise to account for untraced callers confirms a substantial bias against Muslims, and is suggestive of some bias against OBCs and

SCs.

8 Discussion, Interpretation and Conclusion

Our results indicate that Muslims in particular face serious disadvantages in the search for rental housing. To get an expression of interest from 10 landlords, a UC applicant has to apply to about 29 listings while Muslim applicant must apply to almost 45 listings - about 60% more. Although we lack statistical power, point estimates suggest discrimination on other dimensions as well. Landlords who do respond to Muslim applicants as well as UC applicants, tend to respond sooner to UCs, and call more frequently. We do not believe that our findings are biased by the inability to trace the identity of some callers. A simple bounding exercise confirms that even with conservative assumptions about the untraced callers, about 50% more landlords respond to UCs as compared to Muslims. We also find that landlords offering 1-bedroom properties are particularly reluctant to respond to Muslim applicants. Since male applicants for 1-bedroom properties are commonly perceived to be single men, this suggests that single Muslim men may be finding it especially challenging to find suitable housing in Delhi and its suburbs.

Our paper complements and extends the findings of Thorat et al. (2015). They too find significant discrimination against Muslims. But they also find evidence of discrimination against SCs while we don't find any statistically significant difference in the ratio of landlords who respond to UCs v. SCs. However, the two studies use very different audit techniques are the findings are not easily comparable. Thorat et al. (2015) employ telephonic or face-to-face audits, and landlords may be loathe to overtly discriminate in such settings. Note for instance, that 99/7% of the UC applicants in their studies received a positive response while only 33% of the UC applicants in our study received a response. Why then do their findings match ours for Muslim applicants? Interviewer effects may also come in play (see Bertrand and Mullainathan (2003)). For instance, to signal "Muslim" identity, did the auditors consciously or unconsciously dress and behave in a "stereotypical" fashion, which in turn could have evoked negative responses from landlords? Such stereotypes may be harder to convey for SC applicants. Another reason could be that in our study landlords were clearly able to identify Muslims but were not as certain about the caste identity of the SC names. Finally, the bounding exercise in our study did point towards discrimination against SCs and OBCs. It is plausible that with a larger sample and 100% call tracing, we would have also found significant effects.

Online housing markets offer anonymity and flexibility, making them convenient platforms to conduct "clean" discrimination audits. Understanding discrimination in such settings is also increasingly policy relevant, as more markets and transactions move online. There are important questions about how discrimination manifests itself in online settings that facilitate anonymity. In our context, landlords may be more comfortable in discriminating online. In turn, disadvantaged groups may evolve different coping strategies. Specialised markets or agents may emerge who assist the disadvantaged in finding housing. Alternatively, disadvantaged applicants may seek to "disguise" their identity in order to at least get the proverbial "foot through the door". Yet another question is how much of the observed discrimination is taste based versus statistical. If the latter plays a major role, signalling strategies must adapt to the online setting. In addition, there are quirks that may be idiosyncratic to the Indian setting. For instance, dietary preferences are often cited a major reason to discriminate across tenants - many upper caste landlords are vegetarian, and prefer vegetarian tenants. Many of these questions will be explore in the ongoing companion study.

References

Ahmed, Ali M and Mats Hammarstedt, "Discrimination in the rental housing market: A field experiment on the Internet," *Journal of Urban Economics*, 2008, *64* (2), 362–372.

Aigner, Dennis J and Glen G Cain, "Statistical theories of discrimination in labor markets," *Industrial and Labor relations review*, 1977, pp. 175–187.

Altonji, Joseph G and Rebecca M Blank, "Race and gender in the labor market," *Handbook of labor economics*, 1999, *3*, 3143–3259.

Andersson, Lisa, Niklas Jakobsson, and Andreas Kotsadam, "A field experiment of discrimination in the Norwegian housing market: Gender, class, and ethnicity," *Land Economics*, 2012, *88* (2), 233–240.

Arrow, Kenneth et al., "The theory of discrimination," *Discrimination in labor markets*, 1973, *3* (10), 3–33.

Banerjee, Abhijit, Esther Duflo, Maitreesh Ghatak, and Jeanne Lafortune, "Marry for What? Caste and Mate Selection in Modern India," *American Economic Journal: Microeconomics*, 2013, *5* (2), 33–72.

_ , Marianne Bertrand, Saugato Datta, and Sendhil Mullainathan, "Labor market discrimination in Delhi: Evidence from a field experiment," *Journal of Comparative Economics*, 2009, *37* (1), 14–27.

Becker, Gary S, *The economics of discrimination*, University of Chicago press, 1957.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination," Technical Report, National Bureau of Economic Research 2003.

Beteille, Andre, *The backward classes in contemporary India*, Oxford University Press, USA, 1992.

_ et al., "Castes old and new. Essays in social structure and social stratification.," *Castes old and new. Essays in social structure and social stratification.*, 1969.

Blau, Francine D and Lawrence M Kahn, "The US gender pay gap in the 1990s: Slowing convergence," *Industrial & Labor Relations Review*, 2006, *60* (1), 45–66.

_ and Marianne A Ferber, "Discrimination: empirical evidence from the United States," *The American Economic Review*, 1987, pp. 316–320.

Choi, Seok Joon, Jan Ondrich, and John Yinger, "Do rental agents discriminate against minority customers? Evidence from the 2000 Housing Discrimination Study," *Journal of Housing Economics*, 2005, *14* (1), 1–26.

Danziger, Sheldon and Ann Chih Lin, *Coping with poverty: The social contexts of neighborhood, work, and family in the African-American community*, University of Michigan Press, 2000.

Denton, Nancy A, "Half empty or half full: segregation and segregated neighborhoods 30 years after the Fair Housing Act," *Cityscape*, 1999, pp. 107–122.

Desai, Sonalde and Amaresh Dubey, "Caste in 21st century India: competing narratives," *Economic and political weekly*, 2012, *46* (11), 40.

Deshpande, Ashwini, "Caste at birth? Redefining disparity in India," *Review of Development Economics*, 2001, *5* (1), 130–144.

_ , "Levels and Patterns of Consumption Expenditure Among Scheduled Castes and Tribes," Technical Report, Indian Institute for Dalit Studies, New Delhi 2005.

_ , "The grammar of caste: economic discrimination in contemporary India," *OUP Catalogue*, 2011.

Dohan, Daniel, *The price of poverty: Money, work, and culture in the Mexican American barrio*, Univ of California Press, 2003.

Dumont, Louis, *Homo hierarchicus: The caste system and its implications*, University of Chicago Press, 1980.

Fix, Michael and Raymond Struyk, "Clear and convincing evidence: Measurement of discrimination in America," Technical Report, The Field Experiments Website 1993.

Galster, George C, "Housing discrimination and urban poverty of African-Americans," *Journal of Housing Research*, 1991, *2* (2), 87–122.

Ghurye, Govind Sadashiv, *Caste, class and occupation*, Popular Book Depot Bombay, 1961.

Goldin, Claudia and Cecilia Rouse, "Orchestrating Impartiality: The Impact of" Blind" Auditions on Female Musicians," *The American Economic Review*, 2000, *90* (4), 715–741.

Hirsch, Arnold R, *Making the second ghetto: Race and housing in Chicago 1940-1960*, University of Chicago Press, 1983.

Kain, John F and John M Quigley, "Housing market discrimination, home-ownership, and savings behavior," *The American Economic Review*, 1972, pp. 263–277.

Lamb, Charles M, *Housing segregation in suburban America since 1960: Presidential and judicial politics*, Cambridge University Press, 2005.

Massey, Douglas S and Nancy A Denton, *American apartheid: Segregation and the making of the underclass*, Harvard University Press, 1993.

Mullainathan, Sendhil, Markus Noeth, and Antoinette Schoar, "The market for financial advice: An audit study," Technical Report, National Bureau of Economic Research 2012.

NCAER, "Biggest caste survey: One in four Indians admit to practicing untouchability," November 2014.

Office of the Registrar General and Census Commissioner, "Census of India," Technical Report, Ministry of Home Affairs 2012.

Pager, Devah and Hana Shepherd, "The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets," *Annual review of sociology*, 2008, *34*, 181.

Sachar, Rajindar, "Sachar Committee Report," Technical Report, Government of India 2006.

Shariff, Abusaleh, "Socio-economic and demographic differentials between Hindus and Muslims in India," *Economic and Political Weekly*, 1995, pp. 2947–2953.

Siddique, Zahra, "Evidence on caste based discrimination," *Labour Economics*, 2011, *18*, S146–S159.

South, Scott J and Kyle D Crowder, "Housing discrimination and residential mobility: Impacts for blacks and whites," *Population Research and Policy Review*, 1998, *17* (4), 369–387.

Srinivas, Mysore Narasimhachar, "Caste in modern India," *The Journal of Asian Studies*, 1957, pp. 529–548.

Thorat, Sukhadeo and Katherine S Neuman, *Blocked by caste: economic discrimination in modern India*, Oxford University Press, 2012.

Thorat, Sukhdeo, Anuradha Banerjee, K Mishra Vinod, and Firdaus Rizvi, "Urban Rental Housing Market Caste and Religion Matters in Access," *Economic and Political Weekly*, 2015, *50* (26-27).

Turner, Margery A, Stephen Ross, George C Galster, and John Yinger, "Discrimination in metropolitan housing markets: national results from phase 1 of the housing discrimination study (HDS)," Technical Report 2002.

Vithayathil, Trina and Gayatri Singh, "Spaces of Discrimination," *Economic & Political Weekly*, 2012, *47* (37), 60–66.

Yinger, John, *Closed doors, opportunities lost: The continuing costs of housing discrimination*, Russell Sage Foundation, 1995.

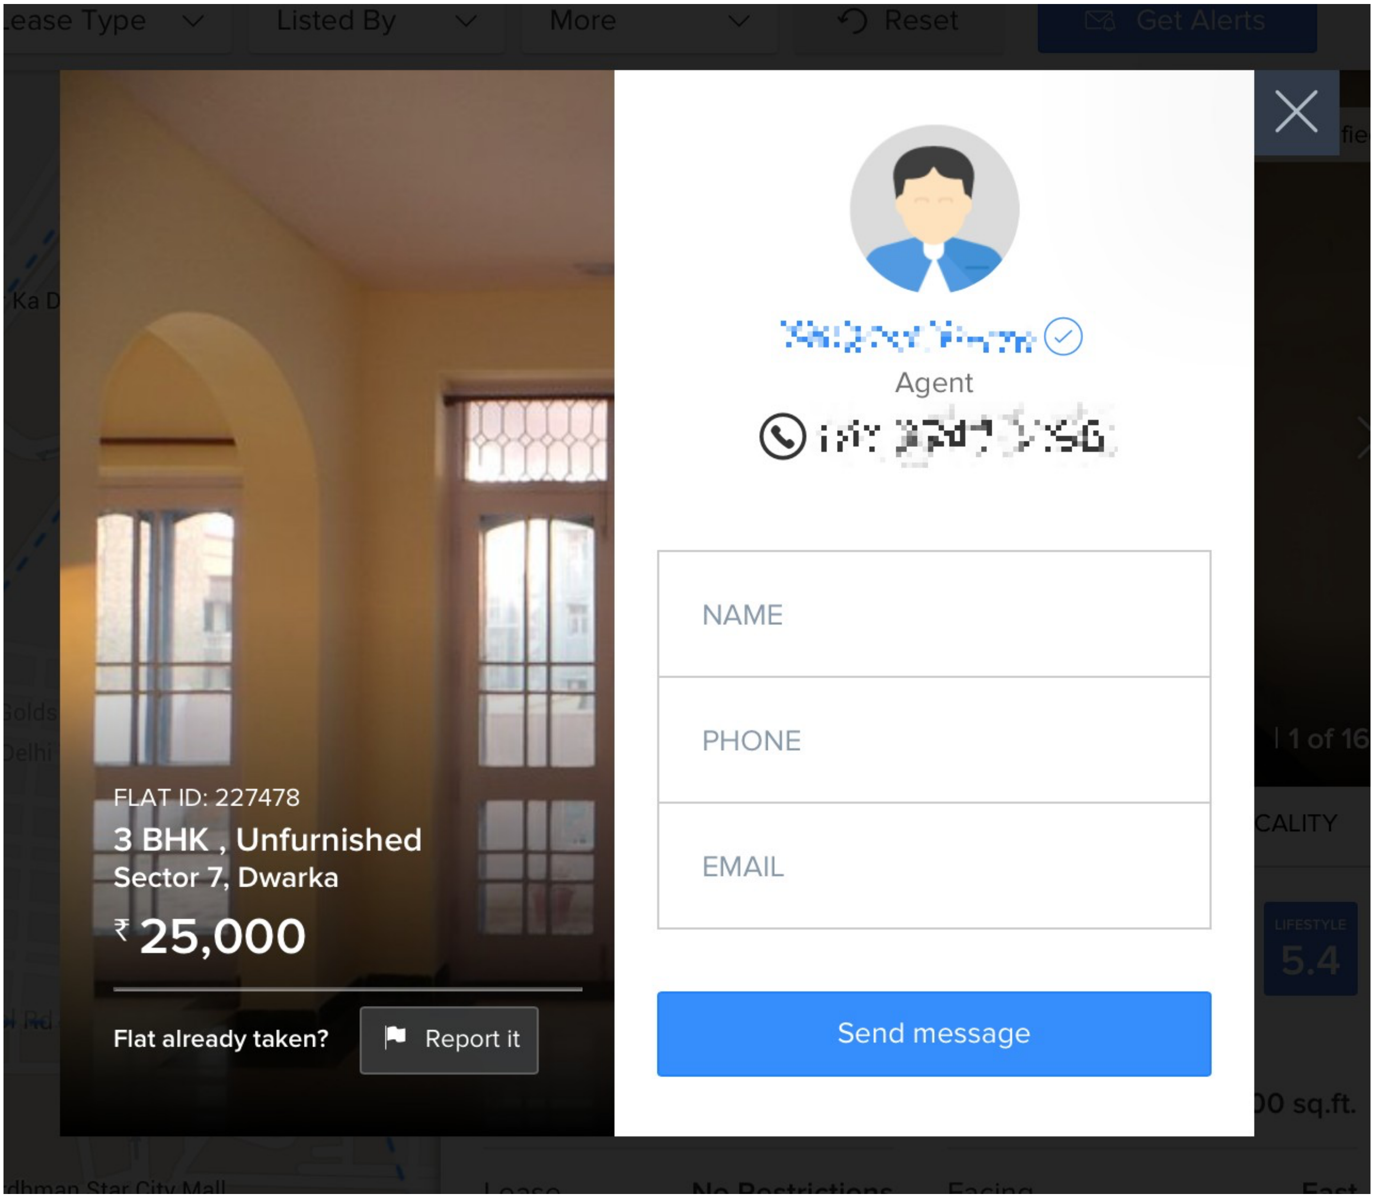

Figure 1: Online Rental Application Form

Table 1: **Listing & Landlord Summary Statistics**

	Delhi City	**Suburbs**[1]	**All**
		House characteristics	
Num bedrooms:[2]			
1	0.26	0.10	0.20
2 or 3	0.67	0.78	0.71
4+	0.07	0.12	0.09
Rent (Rs)	36100.90	29333.32	33726.31
	(53348.86)	(24292.45)	(45353.01)
Floor area (Sq ft)	1138.92	1616.93	1306.64
	(654.86)	(702.55)	(707.91)
Rent/sq. ft	28.47	18.12	24.84
	(20.74)	(10.37)	(18.45)
		Landlord characteristics	
% Female[3]	0.14	0.10	0.13
% Muslim[4]	0.16	0.03	0.12
N	111	60	171

Standard errors in parenthesis.

[1] Suburbs - Gurgaon and NOIDA.

[2] 1-1.5 bedrooms coded as 1 bedroom, 2-3.5 as 2-3 bedrooms.

[3] We were unable to code gender for13 of the 171 landlords (due to missing first name e.g. only initial). The reported female % is computed over all 171 landlords.

[4] We were unable to code religion for 1 of the 171 landlords. The reported muslim % is computed over all 171 landlords.

Table 2: **Application Summary Statistics**

	Order	Daytime	Weekday	Gap (days)	N
UC	2.49	0.84	0.59	4.05	170
	(1.11)	(0.37)	(0.49)	(7.59)	
OBC	2.50	0.74	0.51	4.09	171
	(1.12)	(0.44)	(0.50)	(7.58)	
SC	2.47	0.72	0.51	4.08	170
	(1.12)	(0.45)	(0.50)	(7.61)	
M	2.51	0.84	0.59	4.03	170
	(1.13)	(0.37)	(0.49)	(7.60)	

Standard errors in parenthesis.
Order is the chronological position of the applicant type within the set of 4 applications sent to each landlord.
Daytime: 6:01AM-18:59PM; Weekday: Mon-Fri.
Gap is number of days between date of applying and date the ad was posted.

Table 3: **Counts of Responders & Responses**

A. Counts of Responders

	Unique Callers		Landlords who			
	Total	Excl. spam	Called[1]	Texted	Emailed	Responded[2]
	(1)	(2)	(3)	(4)	(5)	(6)
UC	80	66	56	5	2	59
OBC	67	59	50	6	2	51
SC	63	58	55	4	3	58
M	52	40	35	4	3	38

B. Counts of Responses

			Traced to landlords:			
	All calls	Excl. spam	Calls	Texts	Emails	Total
	(1)	(2)	(3)	(4)	(5)	(6)
UC	192	157	132	5	2	139
OBC	165	149	142	6	2	150
SC	126	112	111	4	3	118
M	101	80	75	4	3	82

[1] In Panel A, Column (3) differs from (2) because some numbers cannot be traced, and because some landlords called from more than one number.

[2] In Panel A, Column (6) is not the sum of (3)-(5) since some landlords both, called and texted or emailed.

Table 4: **Fraction Landlords Responding & Mean Responses per Landlord**

	Fraction responding	**Diff. vs. UC**	**Mean responses (All landlords)[1]**	**Diff. vs. UC**	**Mean responses (Responding landlords)[2]**	**Diff. vs. UC**	**Range responses (min-max)**
UC	0.35	-	0.82	-	2.36		0-9
	(0.48)		(1.52)		(1.74)		
OBC	0.3	0.05	0.88	-0.05	2.94	-0.58	0-33
	(0.46)	(0.05)	(2.90)	(0.25)	(4.74)	(0.70)	
SC	0.34	0.01	0.69	0.13	2.03	0.33	0-10
	(0.48)	(0.05)	(1.50)	(0.16)	(1.96)	(0.34)	
M	0.22	0.13**	0.48	0.34**	2.16	0.2	0-6
	(0.42)	(0.05)	(1.12)	(0.14)	(1.42)	(0.32)	

Responses include calls, emails and texts.

[1] Total traced responses divided by number of landlords contacted.

[2] Total traced responses divided by number of landlords who responded to that type.

Table 5: **Probability of Response by Landlord**

	(1)	(2)	(3)
Muslim	-0.124***	-0.124***	-0.124***
	[0.034]	[0.034]	[0.039]
OBC	-0.049	-0.04	-0.039
	[0.032]	[0.032]	[0.037]
SC	-0.006	0.005	0.003
	[0.034]	[0.034]	[0.040]
Order=2[1]		0.018	0.018
		[0.030]	[0.034]
Order=3		-0.021	-0.007
		[0.033]	[0.047]
Order=4		0.017	0.028
		[0.034]	[0.049]
Gap (days)[2]		-0.009***	-0.021
		[0.003]	[0.034]
Weekday (Mon-Fri)		0.081	0.129*
		[0.056]	[0.065]
Daytime (6AM-6:59PM)		0.032	-0.031
		[0.061]	[0.065]
Suburbs		0.014	
		[0.066]	
2-3 beds		0.101	
		[0.065]	
4+ beds		0.381***	
		[0.115]	
Rent (Rs/sqft)		-0.003**	
		[0.001]	
Landlord FE			x
Constant	0.347***	0.269**	0.369***
	[0.037]	[0.106]	[0.120]
Observations	681	681	681
R-squared	0.012	0.079	0.671

OLS regression coefficients (linear probability models); the dependant variable is a dummy for any response from the landlord.
Robust standard errors in brackets, clustered on landlord.
*** p<0.01, ** p<0.05, * p<0.1

[1] Order is the chronological position of the applicant within the set of 4 applications to a landlord.

[2] Gap in days between date application sent, and date ad posted.

Table 6: **Count of Responses by Landlords**

	(1)	(2)	(3)
Muslim	-0.528***	-0.531***	-0.531***
	[0.146]	[0.146]	[0.147]
OBC	0.07	0.175	0.092
	[0.263]	[0.273]	[0.220]
SC	-0.164	-0.064	-0.136
	[0.142]	[0.161]	[0.134]
Order=2[1]		-0.268	-0.237
		[0.220]	[0.202]
Order=3		-0.513*	-0.352
		[0.295]	[0.277]
Order=4		-0.503*	-0.412
		[0.296]	[0.298]
Gap (days)[2]		-0.058***	-0.236
		[0.017]	[0.182]
Weekday (Mon-Fri)		0.28	0.409
		[0.309]	[0.299]
daytime (6AM-6:59PM)		0.453	0.212
		[0.280]	[0.281]
Suburbs		-0.269	
		[0.265]	
2-3 beds		0.036	
		[0.472]	
4+ beds		0.694	
		[0.522]	
Rent (Rs/sqft)		-0.017**	
		[0.008]	
Ad fixed effect			x
Constant	-0.201	0.106	-19.081***
	[0.142]	[0.618]	[1.290]
Observations	681	681	681

Poisson regression coefficients (dep. variable is count of responses to an applicant)
Robust standard errors in brackets, clustered on landlord.
*** p<0.01, ** p<0.05, * p<0.1

[1] Order is the chronological position of the applicant within the set of 4 applications to a landlord.

[2] Gap in days between date application sent, and date ad posted.

Table 7: **Pairwise Comparison: Time to First Response (hours)**

	Hours between time application sent & first response received					
	UC v. OBC	UC v. SC	UC v. M	OBC v. SC	OBC v. M	SC v. M
OBC	-0.386					
	[5.939]					
SC		2.767		-0.331		
		[4.735]		[2.400]		
M			6.75		-0.244	0.977
			[7.036]		[4.776]	[6.719]
Constant	49.021	29.782**	26.800***	50.924***	21.524***	31.483**
	[33.069]	[11.327]	[1.539]	[8.374]	[6.875]	[11.944]
Landlord FE	X	X	X	X	X	X
Observations	78	84	62	88	46	58
R-squared	0.947	0.966	0.987	0.993	0.688	0.907

Robust standard errors in brackets, clustered on landlord.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$.
OLS regressions at the applicant level.
Each regression only includes landlords who replied to both types in the relevant pair.
Controls include the rank order of the type with the application set, weekday dummy, daytime dummy, gap in days between date ad posted and date application sent, and landlord FE.

Table 8: **Pairwise Comparison: Who Receives the First Response?**

Type 1 v 2	N[1]	Type 1 1st applied[2]	Type 1 1st response[3]	Null Hypothesis 1[4]	p-value	Null Hypothesis 2[5]	p-value
UC v OBC	39	0.46	0.54	0.5	0.63	0.46	0.33
UC v SC	42	0.52	0.64	0.5	**0.07**	0.52	0.12
UC v M	31	0.61	0.65	0.5	0.11	0.61	0.71
OBC v SC	44	0.45	0.61	0.5	0.13	0.45	**0.03**
OBC v M	23	0.43	0.57	0.5	0.53	0.43	0.21
SC v M	29	0.41	0.52	0.5	0.85	0.41	0.26

[1] Number of landlords who responded to both types (includes those who responded to others as well.)
[2] The fraction of landlords to whom Type 1 applied to before Type 2.
[3] The fraction of landlords that responded to Type 1 before Type 2.
[4] The null is that landlords first respond to either type independent of the order in which the types apply.
[5] The null is that landlords first respond to the two types in the same order in which the types apply.

Table 9: **Probability of Response: Interactions of Applicant type with Landlord/Property Features**

	Interacting characteristic of landlord/property (Z):			
	Female	Muslim	One bed	High price
Z	-0.157	-0.224*	-0.139	-0.146**
	[0.095]	[0.130]	[0.085]	[0.073]
OBC	-0.043	-0.043	-0.06	-0.110**
	[0.034]	[0.035]	[0.040]	[0.043]
OBC*Z	-0.048	-0.015	0.039	0.044
	[0.088]	[0.081]	[0.031]	[0.048]
SC	-0.018	-0.007	0.002	-0.049
	[0.038]	[0.037]	[0.041]	[0.046]
SC*Z	0.134*	0.091	0.015	0.069
	[0.078]	[0.105]	[0.060]	[0.052]
Muslim	-0.140***	-0.140***	-0.105***	-0.184***
	[0.037]	[0.038]	[0.039]	[0.046]
Muslim*Z	0.036	-0.005	-0.200***	-0.052
	[0.077]	[0.008]	[0.070]	[0.049]
Constant	0.289***	0.276***	0.305***	0.338***
	[0.083]	[0.080]	[0.083]	[0.090]
Observations	630	677	681	681
R-squared	0.041	0.049	0.055	0.044

Robust standard errors in brackets, clustered on landlord.
*** p<0.01, ** p<0.05, * p<0.1.
OLS regressions at the applicant level (linear probability model)
Controls include the rank order of the applicant with the application set, weekday dummy, daytime dummy, gap in days between date ad posted and date application sent.

Table 10: **Caller Ratio Bounds**

	Traced responders	Ratio (UC: .)	Upper bound	Lower bound
UC	59	1		
OBC	51	1.16	1.22	1.13
SC	58	1.02	1.1	1.02
M	38	1.55	1.68	1.5

See text for the derivation of the upper and lower bounds.