Discrimination, Preferences, and Residential Sorting in Delhi's Rental Housing Market

Hamza Syed

September 1, 2026

Abstract

Household location decisions are central to their access to critical public services and economic opportunities. Residential segregation is routinely read as revealed preference, but in a rental market location reflects both what households want and what landlords, brokers and resident associations will permit. I separate the two using a survey of 3,253 residents and 615 brokers in Delhi, combined with randomized neighbourhood-choice and vignette experiments, gatekeeper elicitation, and list experiments. Own-group residential preferences are substantial across religious and caste groups. Supply-side exclusion of Muslims is severe: brokers report that a Muslim applicant can secure housing in 24 of 100 properties against 86 for an upper-caste Hindu, resident associations bar Muslim tenants in 26 percent of areas against under 1 percent for Hindus, and switching an otherwise identical vignette profile to a Muslim name lowers acceptance by 19 to 23 percentage points, rising to 41 to 48 among brokers. The two forces run in opposite directions: most Muslim households express willingness to live in mixed or majority-Hindu neighbourhoods, so observed sorting is a constrained equilibrium rather than free choice. Because the market is in chronic excess demand, exclusion operates almost entirely on quantity rather than price: roughly 97 percent of the identity-induced shift in landlord response is refusal rather than markup. That concentrates the binding constraint in countable, regulable institutions, brokers and resident welfare associations, and makes it a tractable target for enforcement.

1 Introduction

Where a household lives determines what it has access to: jobs, schools, clinics, transport, and the social networks through which information about all of these travels. Neighbourhood environments shape educational attainment, earnings and intergenerational mobility (Chetty et al., 2016), which makes the housing market one of the principal mechanisms through which economic opportunity is allocated within a city. If that market allocates location by identity

[1]I would like to thank IGC and STICERD for funding the full and pil0t surveys respectively. I would also like to acknowledge the mentorship of my supervisors Prof. Tim Besley, Prof. Maitreesh Ghatak, and Dr. Michael Callen. I would also like to thank Dr. Sam Asher, Dr. Gharad Bryan, Dr. Catherine Szermann, Dr. Gilat Levy, Dr. Michael Greenstone for their feedback and comments throughout the project. Additionally, I would like to thank Ms. Sanjana Nahar, Mr. Ayush Dwivedi, and Ms. Zoheen Shehzad for research assistance.

rather than by willingness to pay and match-quality, the resulting spatial sorting is not merely unequal but also inefficient.

Whether it does so has been studied most intensively in the United States, where residential segregation remains a central explanation for socioeconomic gaps across race and ethnicity (Logan and Parman, 2025; Berry and Wiener, 2020; Christensen et al., 2021; Rosen et al., 2021; Zhao et al., 2006; Ondrich et al., 1999). Audit and correspondence methods have established that discrimination in housing access is real and persistent (Bertrand and Mullainathan, 2004; Asplund et al., 2020; Turner and Ross, 2003a,b; Datta and Pathania, 2016), and recent work shows that it operates not only through outright refusal but through steering: minority households are directed toward neighbourhoods with less economic opportunity and greater exposure to crime and pollution, with measurable welfare costs (Christensen and Timmins, 2022, 2023). Very few studies focus on developing countries ((Datta and Pathania, 2016; Thorat et al., 2015; Gbadegesin and Ojo, 2013)), which are an important context to study this in, as explained later.

That literature is overwhelmingly about race and ethnicity. Religion has received far less attention as a determinant of economic exclusion, and where it has been studied the empirical difficulty has been that of separating it from everything it comes with. In European settings, a Muslim applicant is typically also an immigrant, of distinct national origin, and often racialised, so the religious effect is confounded by construction. Isolating it has required elaborate design. Adida et al. (2010) hold national origin fixed and vary only religious signalling in French résumés, finding that a Muslim candidate is roughly 2.5 times less likely to receive an interview callback than an otherwise identical Christian one. Di Stasio et al. (2021) extend the logic across five European labour markets, distinguishing a "disclosed Muslim" penalty from the stigma of merely originating in a Muslim-majority country. In housing, the pattern recurs and is, if anything, sharper: meta-analysing 25 correspondence studies conducted in OECD countries between 2006 and 2017, comprising more than 300 effect estimates and over 110,000 inquiries to landlords and agents, Flage (2018) finds that majority applicants are almost twice as likely as minority applicants to be selected, and *more* than twice as likely when the comparison group is Arab or Muslim. Religion-related hostility is meanwhile a global rather than a Western phenomenon: Muslims were harassed by state or non-state actors in 143 of the 198 countries Pew Research Center monitors in 2023, and India records among the highest levels of religion-related social hostility of any large country (Pew Research Center, 2026).

India offers something the European literature cannot: a setting in which religion can be

observed without its usual companions. Hindus and Muslims in Delhi are, for the most part, co-national, co-ethnic, co-linguistic and long-settled in the same city. An applicant's religion is legible from a name, but that name carries no signal of immigration status, foreign origin, or race. What the correspondence literature in Europe achieves through careful design, the Indian setting supplies as a feature of the population. This paper exploits that: the identity variation it studies is religious variation, and the caste comparison it reports alongside serves as a benchmark establishing that what is measured is not generic minority disadvantage.

The housing-discrimination literature has two further gaps, and they matter most for interpreting segregation. Audit studies measure access but are silent on preferences. Equilibrium sorting models recover preferences but assume households choose freely across neighbourhoods (Bayer et al., 2007). Neither measures both on the same population. The field therefore lacks an empirical handle on the counterfactual it most often invokes, whether integration would occur absent discrimination, because observed sorting confounds what households want with what the supply side will permit. The point is not hypothetical: documenting Muslim and Scheduled Caste segregation across 1.5 million Indian neighbourhoods, Asher et al. (ming) state explicitly that their analysis cannot disentangle discrimination from homophily, and identify that separation as the natural next step.

There is a second gap, less often noticed, which is that the U.S. case is institutionally unusual. It is a market that broadly clears on price, under an enforced statutory prohibition on housing discrimination, with credit scores, tenant screening services and public rental histories. Remove those features and there is no reason to expect discrimination to take the same form, or to respond to the same policy. Reviewing the housing-discrimination literature, Ifeanacho et al. (2023) find it concentrated almost entirely in the United States, Germany and Canada, the three settings where none of the conditions that characterise most of the developing world's urban rental markets obtain.

This paper addresses both gaps. I measure own-group residential preferences and identity-specific access on the same households and neighbourhoods in Delhi, in a market that does not clear on price, in which contract enforcement is weak, formal tenant information is largely absent, and no anti-discrimination statute reaches private lettings. The central object is the *wedge between desired and feasible sorting*: the gap between where minority households would choose to live and where the supply side will let them.

Contributions. *First*, joint measurement. I field randomized neighbourhood-choice experiments and randomized supply-side vignettes on the same respondents, and elicit access constraints directly from gatekeepers. This identifies preferences and access from separate, purpose-built sources of variation rather than from the residual of one another.

Second, gatekeeper elicitation. Rather than inferring exclusion from applicant-side callback rates, I ask the 615 brokers who control the local rental inventory how many of the properties in their service area would accept an applicant of each identity. This yields spatially resolved, market-share-weighted access measures that a correspondence design cannot produce, and it *locates* discrimination in the intermediary rather than the owner. It also captures a margin that correspondence designs cannot reach: properties an applicant never learns exist, because the broker withholds them before any listing is seen.

Third, and conceptually, I characterise discrimination in a market that rations rather than prices. Delhi's rental market is in chronic excess demand (Section 2.3). This has two consequences the existing literature, built on clearing markets, does not confront. It switches off Becker's competitive-erosion mechanism: a landlord facing a queue forgoes nothing by refusing a Muslim applicant. And it makes the mapping from taste intensity to exclusion severely convex: under rationing, mild and widely shared distaste is sufficient to produce near-total exclusion, because acceptance requires not merely being acceptable but winning a queue. One does not need to posit extreme animus to generate the magnitudes I find, and the evidence suggests one should not.

I combine four sources of data collected in Delhi: a survey of 3,253 resident households, a survey of 615 rental brokers, neighbourhood-level spatial data on amenities and religious institutions, and randomized choice and vignette experiments embedded in both surveys. Neighbourhoods are sampled through a stratified design that, within each stratum, deliberately oversamples areas with a history of communal conflict.

The paper establishes three reduced-form facts. First, own-group residential preferences are real and substantial. Facing identical hypothetical menus, upper-caste Hindu respondents choose the own-group neighbourhood 62.6% of the time against 47.1% for others; the corresponding own-versus-other gaps are even larger for Muslim (42.9 vs 9.3%) and Scheduled Caste (40.4 vs 16.8%) respondents, whose own-group areas are rarely chosen by outsiders. Second, supply-side gatekeeping against Muslim applicants is severe and operates through several channels: brokers report a Muslim applicant would be accepted in only 24 of 100 properties against 86 for an

upper-caste Hindu; resident associations bar Muslims in 26% of areas against under 1% for Hindus; and in a vignette experiment, switching an otherwise identical profile from a Hindu to a Muslim name lowers acceptance by 19–23 percentage points pooled, rising to 41–48 points among brokers. Third, the two forces do not coincide: most Muslim households reveal a willingness to live in mixed or majority neighbourhoods, yet brokers report they can access only about a quarter of the market, sofeasible sorting is strictly more segregated than desired sorting. This wedge is the paper's central object.

Two results follow from taking rationing seriously. Discrimination here operates overwhelmingly on the *quantity* margin: identity changes produce large and precise refusal effects, but the price margin barely moves. Of the probability mass displaced from acceptance at the offered rent, roughly 97 percent lands on refusal and under 3 percent on acceptance at a higher rent. In a market that cleared on price the ordering would be reversed. That contrast is itself evidence that identity functions as a rationing device.

And the exclusion appears to be enacted rather than owned. Direct measures show severe discrimination; list experiments that shield respondents from social-desirability pressure show almost no *admitted* discrimination, even as the same indirect method uncovers widespread but suppressed majoritarian sentiment. Brokers, who screen applicants before landlords ever see them and who bear none of the tenant's default risk, are nonetheless the most exclusionary link in the chain, which suggests their conduct reflects the norm they expect to face rather than a taste they hold.

Policy. The binding constraint sits in identifiable institutions. This matters more here than the standard framing suggests, because India has no Fair Housing Act analogue and the leading authority affirms the autonomy of housing societies to restrict membership by religion. The policy question is not how to enforce an existing right but how to create a regulatory hook where none exists, and the state already operates a registration-and-conduct regime for real estate agents in the primary sale market that does not currently reach rental brokerage.

The remainder proceeds as follows. Section 2 describes the setting and the regulatory status of brokers and RWAs. Section 3 presents the conceptual framework. Section 4 describes the data and design, and Section 5 the measurement of each object. Section 6 presents the reduced-form results. Section 7 develops the structural model and welfare analysis, Section 8 the enforcement implications, and Section 9 concludes.

2 Context

2.1 Segregation, Sorting, and Opportunity in Urban India

Frictions in housing markets impose welfare and efficiency costs beyond their distributional ones, by limiting households' ability to reach productive locations. A large literature in urban economics emphasises the role of housing markets in facilitating labour mobility: housing-supply constraints can limit migration to high-productivity cities and generate aggregate productivity losses by preventing workers from locating where they are most productive (Moretti, 2011; Hsieh and Moretti, 2019). Location also governs access to local public goods, and households sort across neighbourhoods partly in response to differences in amenities and service provision (Tiebout, 1956; Bayer et al., 2007), so barriers to residential mobility may reinforce existing socioeconomic disparities rather than merely reflect them. These mechanisms are especially consequential in developing-country cities, where access to jobs, schools, and public services is often highly spatially concentrated, and where informal institutions and social networks play a large role in determining opportunity.

Residential segregation along religious and caste lines remains a prominent feature of urban India. A large literature documents the spatial concentration of Muslims, Scheduled Castes (SCs), and Scheduled Tribes (STs) in particular neighbourhoods and localities, often accompanied by unequal access to public services, infrastructure, and housing markets (Asher et al., ming; Singh, 2019).

Among religious groups, Muslims have attracted particular attention. Using ward- and neighbourhood-level data, Susewind (2017) shows that Muslim residential segregation is widespread across Indian cities, although its extent and form vary substantially across locations. He argues that commonly used narratives of "ghettoisation" often overlook considerable heterogeneity in the degree and form of segregation, while emphasising the role of insecurity, historical trajectories, and perceptions of exclusion in shaping residential patterns. More recent work demonstrates that Muslim segregation can be measured at a much finer spatial scale and finds substantial intra-city variation across neighbourhoods (Shah, 2024).

Qualitative and historical studies similarly document the emergence of Muslim-majority neighbourhoods in many Indian cities. Gayer and Jaffrelot (2012) describe how processes of marginalisation, communal violence, and restricted residential opportunities have contributed to the concentration of Muslims in particular urban spaces. Galonnier (2015) argues that Muslim

residential clustering can simultaneously reflect three distinct processes: involuntary segregation resulting from discrimination and insecurity (the "ghetto"), voluntary clustering driven by cultural or religious preferences (the "enclave"), and exclusionary practices by dominant groups (the "citadel"). This perspective highlights the difficulty of interpreting observed residential concentration as evidence of either preferences or exclusion alone.

Recent evidence suggests that segregation may have intensified in some contexts. Reporting documents increasing movement of Muslim households into Muslim-majority neighbourhoods in Delhi following episodes of communal violence and rising perceptions of insecurity, with many households citing safety concerns and anticipated discrimination as important considerations in their housing decisions (Reuters, 2024).

Residential segregation is not limited to religion. Using Census data from 2001 and 2011, Singh (2019) finds that caste-based residential segregation persisted or increased in a majority of Indian cities over the decade. The persistence of caste segregation suggests that urbanisation alone does not necessarily erode social boundaries, and may instead reproduce existing forms of social stratification in spatial form. These patterns may have consequences beyond residential composition itself: Muslims and Dalits are disproportionately concentrated in neighbourhoods with weaker access to public services and infrastructure, potentially contributing to persistent disparities in education, labour-market access, and welfare (Asher et al., ming; Alesina et al., 1999). Exclusion from particular neighbourhoods may therefore shape not only where households live but also their access to labour markets, schooling, and urban services, with the concentration of disadvantaged groups in lower-opportunity areas a plausible channel for persistent inequality.

Taken together, the literature suggests that observed segregation reflects a combination of preferences, discrimination, safety concerns, institutional restrictions, and historical processes. Existing work, however, has largely been unable to quantify the relative contribution of these mechanisms. This paper contributes by separately measuring residential preferences, discriminatory access constraints, broker-mediated restrictions, and neighbourhood amenities within a common empirical framework.

2.2 Religion as an Axis of Economic Exclusion

Evidence from outside India

The economics of discrimination has developed largely around race in the United States and ethnicity in Europe. Religion has been studied less, and in Western settings it is difficult to

study cleanly, because religious minority status is bundled with immigration, national origin and race. The studies that have made progress did so by engineering the separation.

Adida et al. (2010) conduct a correspondence test in the French labour market in which applicants share Senegalese origin and differ only in religious signalling: first name, and the religious affiliation of a former employer and a volunteering organisation. The Muslim applicant is approximately 2.5 times less likely to be called for interview. In a follow-up survey they find matched second-generation Muslim households have lower income than comparable Christian households. Di Stasio et al. (2021) generalise the design across Germany, the Netherlands, Norway, Spain and the United Kingdom, sampling applicants from countries with substantial Muslim populations and varying whether the application discloses religious affiliation. This permits them to separate an anti-Muslim penalty from a "Muslim by default" origin stigma, and they find the disclosed-Muslim penalty operates independently of origin.

In housing markets the evidence is more direct. Ahmed and Hammarstedt (2008) document that Arab/Muslim-named applicants to Swedish rental listings receive substantially fewer callbacks, enquiries and viewings than Swedish-named applicants with identical characteristics. Flage (2018) aggregates the field: pooling 25 OECD correspondence studies from 2006 to 2017, comprising more than 300 estimates and over 110,000 inquiries, he finds majority applicants almost twice as likely to be selected as minorities generally, and more than twice as likely relative to Arab or Muslim applicants specifically. Religion-related hostility is meanwhile a global rather than a Western phenomenon: Muslims were harassed by state or non-state actors in 143 of the 198 countries Pew Research Center monitors in 2023 (Pew Research Center, 2026).

Two results from this literature bear directly on the design of the present paper, and I return to both in Section 6. First, the *information* margin. Ahmed et al. (2010) ask whether supplying landlords with additional information about applicants reduces discrimination in the Swedish rental market, and Auspurg et al. (2017) find in a German field experiment that signalling high occupational status significantly reduces discrimination rates, evidence consistent with a statistical mechanism operating in those markets. My vignette design includes a directly comparable reliability-signal arm. That the signal moves acceptance in Germany but not in Delhi is informative about the mechanism, and is what the framework of Section 3 predicts for a market that rations rather than prices.

Second, the *tightness* margin. Christensen and Timmins (2023) find discriminatory constraints in U.S. rental markets substantially stronger for listings on the market fewer than three

days than for those listed more than a week, and stronger still once a property manager has already received inquiries from other applicants. Discrimination intensifies where the landlord has a queue, the comparative static that Result 2 formalises and that motivates treating excess demand as a primitive rather than a background condition.

Why the Indian setting identifies religion

The confound that European studies design around is largely absent here. Hindus and Muslims in Delhi share nationality, and broadly share language, region of origin, physical appearance and length of settlement. A name signals religion without simultaneously signalling foreign birth, migration status or race. The religious effect is therefore identified from population structure rather than from an experimental workaround, and the estimates are not contaminated by the origin stigma that Di Stasio et al. (2021) must model explicitly.

The setting also permits a benchmark unavailable elsewhere. Scheduled Caste applicants are a stigmatised minority who are neither religiously distinct from the majority nor distinguishable by any of the markers that separate Muslims from Hindus. Comparing the two penalties therefore separates religion-specific exclusion from generic minority disadvantage. The broker data make this comparison sharp: the Muslim access gap relative to upper-caste Hindus is 62.0 percentage points, while the gap for lower-caste Hindus is 3.7 percentage points. Whatever is operating in this market is religion-specific, and is not a proxy for poverty, low status, or the general tendency of gatekeepers to prefer applicants like themselves. For this reason I retain the caste comparison in the main text while reporting caste-specific analysis in Appendix A.

The Indian trajectory

Concern about anti-Muslim hostility in India has grown over the past decade. Monitoring organisations have documented substantial increases in anti-Muslim and anti-minority hate speech, while international human rights reporting has highlighted discrimination and violence directed toward religious minorities (Reuters, 2025; Human Rights Watch, 2024; U.S. Department of State, 2024). Pew Research Center places India among the large countries with the highest levels of religion-related social hostilities (Pew Research Center, 2026). These developments raise a basic economic question: to what extent does residential segregation reflect household preferences, and to what extent does it emerge as a response to discrimination, insecurity, and restricted access to housing markets?

I record this trajectory because it is the context in which the data were collected, but the paper's contribution does not rest on it. The mechanisms it identifies, rationing under excess demand, ex-ante screening under weak contract enforcement, and gatekeeping by intermediaries, are structural features of the market rather than functions of the political moment. They would operate at lower intensity in a calmer period, and they operate on other identity margins in other settings. What a period of heightened salience provides is statistical power and, through the conflict-exposure design of Section 4.1, variation in salience within the city that the analysis can exploit.

2.3 Delhi's Rental Market and Why It Does Not Clear

Delhi provides a useful setting for studying identity-based housing exclusion. As the national capital and one of India's principal urban centres, it draws people from across the country in search of livelihoods. Rental housing accounts for more than 30 percent of urban residence in India, particularly among migrants, young workers, students, and lower-income households (Knight Frank, 2019), and Delhi alone accounts for roughly 5 percent of all rental housing in the country. Yet access to rental housing frequently depends not only on price and availability but also on informal social rules governing who counts as an acceptable tenant. Media reports and anecdotal evidence point to Muslims and unmarried tenants facing particular difficulty in finding housing in several parts of India (Adukia, 2024; The Caravan, 2019; Chandran, 2017; Mander, 2025; The Indian Express, 2018).

The framework in Section 3 takes chronic excess demand in this market as a primitive. The remainder of this subsection establishes it, because the assumption does substantial work and is not innocuous.

The deficit. The Delhi government's Economic Survey for 2020–21 records 6.75 million people living in low-income settlements in the capital, 695 slum settlements, 1,797 unauthorised colonies and 362 urban villages, and an additional requirement of 2.4 million housing units. Nationally, the Technical Group on Urban Housing Shortage put the 2012 shortfall at 18.78 million units, ninety-five percent of it among low-income households.

Withdrawal of rental supply. The more striking fact is that the rental sector has contracted even as the deficit has grown. Rental housing fell from 53 percent of the occupied urban residential stock in 1971 to 28 percent in 2011 (Gandhi et al., 2021), while vacancy rose: the 2011

Census records 11.1 million vacant urban units, 12.4 percent of the stock, against 1.8 million in 1971. A shortage and a large, growing vacant stock coexist. Gandhi et al. (2021) attribute this to the interaction of pro-tenant rent control with weak contract enforcement, and estimate that permitting free rent revision reduces district vacancy by 2.8 to 3.1 percentage points. Gross rental yields run at 2 to 4 percent, below the risk-free rate, so the return to letting is thin and the option value of holding a unit vacant is correspondingly high.

Consequence. Excess demand at prevailing rents means the marginal unit has a queue. A landlord who declines an applicant on identity grounds lets to the next in line at the same rent. **The private cost of discriminating is approximately zero.** This is the single most important institutional fact for what follows, and it is why Becker's prediction that competition erodes discrimination does not apply here.

2.4 Tenancy Law, Contract Enforcement, and the Anatomy of a Transaction

Rent and tenancy are state subjects. Delhi is governed by the Delhi Rent Control Act 1958, which fixes standard rents, restricts revision, and constrains the grounds on which a landlord may recover possession. Landlords have adapted rather than complied. The dominant instrument is the eleven-month *leave-and-licence* agreement, executed under the Indian Easements Act 1882, which grants permission to occupy without transferring an interest and which, being under twelve months and unregistered, falls outside rent control and permits revocation on notice. Landlords are typically reluctant to register even these, for fear that registration invites future statutory capture. Roughly 19 percent of households in formal urban rental housing hold a written contract, and roughly 14 percent in slums (Gandhi et al., 2021).

Enforcement is the second constraint. Contract enforcement in India took an average of 1,445 days as of 2019, and Gandhi et al. (2021) find judicial density negatively related to vacancy. The Model Tenancy Act 2021 would require written agreements, cap deposits, and impose a sixty-day deadline on a dedicated dispute-resolution tier; adoption has been uneven and Delhi has not adopted it yet. This also means that there is no formal data on the tenant–landlord relationship, a shortcoming that the sampling design tries to address.

Why this matters for discrimination. The screening problem is the mirror image of the enforcement problem. A landlord who cannot rely on ex post remedies, who cannot swiftly evict, cannot easily recover arrears, and holds no enforceable written instrument, must resolve

his exposure *ex ante, at the door*. But the information available at the door is thin: incomes are largely informal, and there are no credit scores, tenant databases, or rental histories. Screening falls back on whatever is cheap and observable, and identity is the cheapest observable there is.

This generates a testable prediction: because identity is being used as a *substitute* for absent verification, offering verification directly should reduce the identity penalty. Section 6 shows that it does not.

How a rental transaction works

Brokers, community networks, and online portals are the main avenues for rental housing search in Delhi and most major cities in India, with over 90 percent of the tenants in the survey reporting having used one or more of these avenues. The rest rely on door-to-door searches for 'To-let' signs.[1] Brokers often post and manage advertisements on these websites at the behest of the landlords. Some housing websites like NoBroker explicitly market themselves on direct contact with the landlord, saving the tenant and landlord the brokerage.

Brokers handle the administrative documents and pre-screen potential tenants before taking the shortlisted applicants to the landlords. This was confirmed in the focus group discussions and informal conversations with brokers in Delhi. The landlords specify the type of tenants they want and brokers screen the applicants accordingly. Identity, occupation, family composition, dietary preferences and similar attributes are established pre-viewing in the first meeting with the broker.

After the pre-screening, the broker arranges the viewing. Most take the potential tenant to the house in their own vehicle or accompany the tenant on their own vehicle. They normally charge half a month's rent from the tenant and the same from the landlord after the deal is formalised, although this could be higher or lower based on negotiations. There is normally no fee for showing a house, and brokers service an area of around 1.5 to 2 kilometres around their shop.

Post-viewing, there can be negotiations over rent and deposits. The brokers generally do the first round of negotiations. Landlords specify an ideal rent and the amount they would be willing to negotiate. Applications can be rejected by brokers if they are below the landlords'

[1] 99acres captures roughly 30 percent of online property-search traffic in Delhi, MagicBricks roughly 25 percent and Housing.com roughly 20 percent, with the remainder fragmented across direct-from-owner sites, classifieds and social-media groups, many of the listings on which are posted by brokers. Source: Semrush Traffic Analytics, accessed 15 May 2026.

expectations. In some cases the landlords can also be involved in these negotiations, at the broker's request.

Once the rent and deposit are finalised, the applicant pays a token amount to hold the property until further checks like police verification and identity documents are processed. Once the formalities are done, the parties sign an eleven-month leave-and-licence agreement, which can be renewed before its expiry, generally with some increase in rent. Although, as mentioned earlier, a majority of the agreements end up being informal.

The critical feature is that pre-screening precedes the viewing. Exclusion at the pre-screen is invisible to the applicant, who never learns which properties existed. A correspondence study observes pre-screening outcomes for properties the applicant already knows about; it cannot observe the inventory withheld. Gatekeeper elicitation can.

2.5 The Gatekeepers and the Legal Status of Exclusion

Three actors exercise discretion over who is an acceptable tenant, and the discretion is exercised along identity lines.

Landlords hold the residual right to select and bear the consequences of selection: the default and possession risk described above. Datta and Pathania (2016) find in a Delhi correspondence study that landlords are markedly less responsive to Muslim than to comparable upper-caste Hindu applicants.

Brokers. What they do: hold and generate the listing inventory for a service area of roughly 1.5 to 2 kilometres; pre-screen applicants before any landlord contact; vouch for the tenant to the landlord, substituting reputational intermediation for the formal screening infrastructure the market lacks; draft the agreement; assist with verification formalities.

What they do *not* do: own the property; set the rent; sign the contract; **bear the tenant's default risk**.

That last point carries analytical weight. The broker is an agent whose income is a commission on a completed match and whose continuing business depends on his standing with the landlords in his area. He internalises none of the cost of a bad tenant and all of the cost of a referral the landlord finds objectionable. His screening rule is therefore not "will this tenant perform?" but "will this tenant be accepted?", which means his revealed penalty is not his own risk assessment, and need not be his own taste.

Resident Welfare Associations administer eligibility norms that are often unwritten but enforced in practice, and that operate on the landlord rather than the tenant: a landlord who lets to a prohibited category faces sanction from his neighbours. Reporting based on the membership records of all 1,960 registered housing societies in Delhi found 68 percent with no Muslim member; Muslim membership across societies barely exceeds 3 percent (TwoCircles.net and Newslaundry, 2024).

Why this concentration matters

The enforcement framing of this paper rests on a feature of the Delhi, and several developing country cities', market that the existing literature has largely treated as background: the discretion that shapes tenant access is concentrated in two identifiable types of institution, brokers and Resident Welfare Associations, whose conduct is at least partly subject to regulation. Brokers mediate a large share of rental transactions and typically serve a compact area, so the population of relevant gatekeepers in any neighbourhood is small and observable. RWAs administer tenant-eligibility norms that are frequently unwritten but enforceable in practice. Both are countable institutions rather than a diffuse population of private owners, which is what makes the discrimination this paper measures a tractable target for policy rather than an immutable feature of preferences.

The legal position

India has no general statutory prohibition on discrimination in private housing. Article 15(2) of the Constitution reaches shops, public restaurants, hotels and places of public entertainment; it does not extend to private lettings. And in *Zoroastrian Cooperative Housing Society Ltd. v. District Registrar, Cooperative Societies (Urban)*, (2005) 5 SCC 632, the Supreme Court *upheld* a housing society's bye-law restricting membership to members of a single religious community, holding that nothing in the governing Cooperative Societies Act prohibited such a restriction, that it was not contrary to public policy, and that the associational autonomy protected by Article 19(1)(c) was engaged.

Media reports document incidents where neighbours protested against a Muslim couple buying a house in a Hindu locality in Moradabad.[2] They were forced to re-sell their house to a Hindu family, and the Hindu family which sold the house to Muslims also faced backlash. Several such

[2]BBC News, accessed 15 May 2026: https://www.bbc.co.uk/news/articles/cp837p125ywo

incidents have been documented in the media in different parts of the country.

Identity-based eligibility rules in housing societies are therefore not, in general, unlawful. This reframes rather than weakens the policy contribution: the object is not to enforce an existing right but to identify where a regulatory hook can be created at low institutional cost. Section 8 argues that the broker, not the society, is that point.

3 Conceptual framework

3.1 What exclusion is

Say that property j is *excluded* for household i if i cannot secure a tenancy at j at the prevailing rent r_j, notwithstanding that i is willing and able to pay it. Exclusion is a restriction on the *choice set*, not a shift in the ranking over a fixed choice set.

Exclusion takes three observationally distinct forms, each measured with a different instrument:

Table 1: Forms of exclusion and their measurement

Form	Mechanism	Measured by
Refusal	Landlord declines at any rent	Vignette penalty β_1 (Table 6)
Non-showing	Broker withholds; applicant never learns	Broker access (Figure 4)
Rule-based	RWA norm makes the let infeasible	Reported RWA rules (Figure 6)

Two adjacent objects should not be conflated with exclusion. *Price discrimination* is a markup conditional on acceptance (Table 7): a movement along the choice set, not a restriction of it. Distinguishing the two is not pedantry; Result 1 shows their relative size into a test. *Anticipated exclusion* is self-restriction of search in expectation of refusal, which produces sorting in excess of what realised exclusion alone would generate (Pager and Shepherd, 2008).

3.2 A simple model

Environment. A city has K neighbourhoods, $n = 1, \ldots, K$. Neighbourhood n contains N_n rental properties indexed j, with $N = \sum_n N_n$. There are M households seeking a tenancy. Each household i has type $g \in \{HUC, MUS, SC\}$, Hindu upper caste, Muslim, Scheduled Caste, with population shares μ_t.

Assumption 1 (Excess demand)**.** $M > N$. $q \equiv M/N > 1$ is the queue ratio.

Preferences. The utility household i derives from property j in neighbourhood $n(j)$ is

$$U_{ij} \;=\; \phi A_n \;-\; r_j \;-\; \tau\, c_{i,n(j)} \;+\; \eta_{g(i)}\, s_{g(i),n(j)} \;+\; \xi_{n(j)} \;+\; \epsilon_{ij} \tag{1}$$

where A_n is an amenity index, r_j rent, c_{in} commuting cost, $s_{g(i)n}$ the share of n's residents of type g, ξ_n unobserved neighbourhood quality, and ε_{ij} an i.i.d. type-I extreme value shock. The parameter η_t is own-group residential preference. Equation (1) is what the neighbourhood-choice experiment identifies; Table 5 reports η_t.

Supply. The landlord of property j receives, from letting to i,

$$\Pi_{ij} \;=\; r_j \;-\; \kappa_{ij} \;-\; d_{g(i),j} \tag{2}$$

where κ_{ij} is the expected cost of default, damage and failure to recover possession, large in this market and largely unverifiable ex ante, and $d_{g,j}$ is the non-pecuniary cost of a type-g tenant, decomposed as

$$d_{g,j} \;=\; \underbrace{d^L_{g,j}}_{\text{landlord's own taste}} \;+\; \underbrace{d^N_{g,n(j)}}_{\text{expected RWA/neighbour sanction}} \tag{3}$$

The decomposition in (3) is not trivial. d^L is a preference; d^N is a norm. They have identical implications for behaviour and opposite implications for policy.

Rationing. Under assumption 1, property j receives a queue Q_j of applicants, $|Q_j| = l \geq 1$, i.e. the queue length. The landlord's problem is *not* to accept or reject a single applicant. It is to select

$$i^*_j \;=\; \underset{i \in Q_j}{\operatorname{argmax}}\; \Pi_{ij} \tag{4}$$

Acceptance requires *winning the queue*, not merely being acceptable. Everything below follows from that substitution.

3.3 Three implications

Result 1 (Rationing shows up as refusal)**.** *If the landlord could set a type-specific rent he would post $r_j(g) = r_j + d_{g,j}$, every type would be housed, and discrimination would appear entirely as a price premium. Under assumption 1, with posted rents that adjust sluggishly and a queue at the prevailing rent, the landlord has no reason to price the disamenity: he selects a different tenant.*

Discrimination appears as refusal.

The market therefore reveals its own regime, and the relative magnitude of the extensive and intensive margins is a *test*. Section 6 finds identity penalties of 19 to 23 percentage points pooled, rising to 41 to 48 percentage points among brokers, against a price margin that is small, thinly populated, and, because the most discriminating respondents refuse rather than quote, selected toward zero. That is what a rationed market looks like.

Result 2 (Exclusion is convex in taste)**.** *Take $d_{H,j} = 0$, $d_{M,j} = d > 0$, and κ type-independent. From (4), a Muslim applicant is selected only if the queue contains no Hindu applicant:*

$$\Pr(\textit{accept} \mid M) \;=\; \mu_M^{k-1} \tag{5}$$

Two things follow. First, acceptance in (5) does not depend on the *magnitude* of d at all, only on its sign. Once there is a queue, any strictly positive distaste produces exclusion falling geometrically in queue length. One does not need extreme animus to explain broker-reported access of 24 properties in 100.

This resolves what would otherwise be a contradiction in the evidence. The direct measures show severe discriminatory behaviour; the list experiments show almost nobody will admit to discriminating, even anonymously. Under (5) both are true simultaneously, because from any individual gatekeeper's point of view nothing that felt like discrimination occurred, only a mild preference exercised over a queue. The behaviour is severe; the taste need not be. **Discrimination here is an aggregation phenomenon.**

Second, $\partial \Pr(\text{accept} \mid M)/\partial k < 0$: exclusion rises with market tightness. This is directly testable and is confirmed in the U.S. setting by Christensen and Timmins (2023), who find discriminatory constraints substantially stronger for listings on the market fewer than three days (relative response 45.1 percent) than for those listed more than seven days (77.5 percent), and stronger still when a property manager has already received inquiries from other applicants.

Result 3 (Information does not move the penalty)**.** *Let ω be a signal that reduces the posterior on κ. Under (4), acceptance depends on the* ranking *within the queue, not on the level of expected surplus. A signal available to all applicants shifts every κ down and leaves the ranking intact. So long as $d > 0$, the Hindu applicant still wins.*

A null on the reliability parameter β_2 in Section 6 therefore does *not* cleanly establish

that discrimination is taste-based. It establishes that reliability beliefs are not the only thing operating, that some component survives the provision of information. That component may be taste d^L, norm d^N, or statistical discrimination on a dimension the signal does not address.

3.4 The intermediary

The broker b earns a commission on a completed match and forwards i to j only if he expects acceptance. Because he bears no κ, his revealed penalty decomposes as

$$d^B = \underbrace{d_b^{\text{own}}}_{\text{his own taste}} + \underbrace{\mathbb{E}_b\left[d_j^L\right]}_{\text{expected landlord taste}} + \underbrace{\mathbb{E}_b\left[d_n^N\right]}_{\text{expected RWA norm}} \tag{6}$$

Only the first term is a preference the broker holds. The remainder is *derived demand for discrimination*.

Equation (6) makes sense of a result that is otherwise puzzling: the broker penalty substantially exceeds the landlord penalty despite the broker bearing none of the risk. A plausible reading is that the broker screens on behalf of the *strictest* landlord he might approach. A referral the landlord finds objectionable is costly to a relationship he must maintain across many future transactions, while a rejected applicant costs him one commission from a queue of substitutes. Over-exclusion is his optimal response to asymmetric costs. This is the analogue of the mechanism Christensen and Timmins (2022) document, in which agents cater to the prejudiced beliefs of their clients.

3.5 Relation to frictional models of discrimination

The second-generation literature asks how discrimination survives competition. Becker's answer, that it should not, depends on a market that clears. Black (1995) shows that in an equilibrium search model, discrimination by a minority of prejudiced employers depresses minority outcomes and persists, because search frictions prevent workers from costlessly relocating. Bowlus and Eckstein (2002) estimate such a model structurally. Surveying the field, Lang and Lehmann (2012) conclude that taste-based models generate differentials under search *even when strong prejudice is not widespread*, precisely the convexity of Result 2, while noting that no model of the time explained the extensive-margin (unemployment-rate) differential.

This paper's relation to that literature is threefold. The friction here is not primarily search cost but *rationing*: the constraint is not that the minority household must search longer, but

that a subset of the stock is unavailable at any price. The extensive margin is the dominant one, the housing analogue of the unexplained unemployment-rate differential, and Result 1 supplies a mechanism. And empirically, these models are typically estimated by inverting observed outcomes; here the choice-set restriction is measured directly, from the gatekeepers who impose it.

3.6 The labour-market analogue

Table 2: Structural correspondence between the two markets

Housing	Labour
Landlord	Employer
Broker	Recruiter, temp agency, referral network
RWA	Closed shop, occupational licensing board
Queue under excess demand	Job rationing under a binding wage floor
Refusal	Non-hiring; unemployment-rate differential
Rent markup	Wage differential
Neighbourhood as input to opportunity	Firm or occupation as input to human capital

The parallel is not merely expositional. The labour literature has concentrated on the wage margin because wages are observed and hiring rejections are not. Housing markets with gatekeeper elicitation reverse this: the refusal is observable and the price is the harder object. The two literatures are complementary exactly where each is weak.

3.7 Access, sorting, and the wedge

Let $a_{ij} = \Pr(i_j^* = i)$. Define neighbourhood-level access for type g:

$$p_{gn} \;=\; \frac{1}{N_n} \sum_{j \in n} a_{gj} \tag{7}$$

Equation (7) is precisely what the broker access measure estimates: the share of properties in a service area at which a given type would be accepted. Figure 4 reports $\hat{p}_{gn}$ averaged over n.

Combining preferences and access, the probability that i resides in n is the feasibility-weighted logit

$$P_{in} \;=\; \frac{p_{g(i)n} \exp(V_{in})}{\sum_m p_{g(i)m} \exp(V_{im})} \tag{8}$$

which collapses to the standard conditional logit when $p \equiv 1$.

Let D^* be the sorting distribution under $p \equiv 1$ (desired) and D the distribution under $\hat{p}$ (feasible). The wedge $W = D - D^*$ is the paper's central object; Section 7 uses it for welfare calculations.

3.8 Rents are equilibrium objects

Equations (1) and (8) treat r_j as given. It is not. Exclusion redistributes demand: if Muslim households are barred from a large share of the stock, their demand concentrates on the accessible remainder, bidding rents there *up* relative to quality. Against this, if Muslim-composition areas are disfavoured by all groups, which the estimated area main effect near 0.08 suggests, stigma pushes rents *down*. The net sign is ambiguous, which makes it worth estimating. I report a hedonic in Section 6:

$$\ln r_j \;=\; \alpha + \boldsymbol{\beta}' A_j + \gamma\, \text{MuslimArea}_j + \varepsilon_j \tag{9}$$

Two consequences. The counterfactuals in Section 7 hold rents fixed and are therefore *partial equilibrium*: equalising access shifts demand across neighbourhoods and rents must adjust, partly offsetting the access gain (Bayer et al., 2007). And the scraped listing rents are asking prices for the stock accessible to *someone*; they are not transaction prices and need not be the price a Muslim applicant faces.

4 Data

The analysis combines four sources of data collected in Delhi: a survey of resident households, a survey of rental brokers, neighbourhood-level spatial data on amenities and religious institutions, and a set of randomized housing-choice experiments embedded in both surveys. The design is intended to observe, for the same neighbourhoods, both the preferences households reveal over locations and the access constraints they face on the supply side, the two objects the framework in Section 3 shows must be measured separately.

4.1 Sampling Design

Respondents were selected through a multi-stage design stratified by local religious composition. In the first stage, I mapped neighbourhood amenities and religious institutions across the city and used them to classify neighbourhoods into strata reflecting their religious composition, in particular areas with mosques versus those with temples, alongside mixed areas.

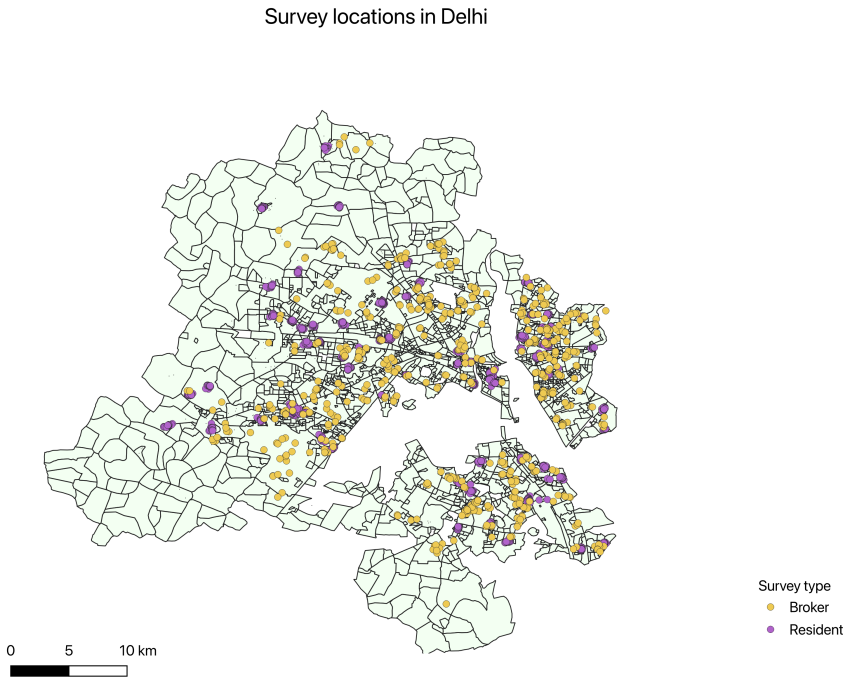


Figure 1: Survey areas

In the second stage, neighbourhoods were drawn at random within strata. Within each selected neighbourhood, enumerators began from a random starting point and surveyed every fifth household along a fixed route, a random-walk procedure that limits enumerator discretion in selecting respondents. Stratifying on religious composition ensures adequate representation of both minority-concentrated and majority neighbourhoods. All resident and broker locations were geocoded, allowing the survey data to be linked to the spatial amenity measures and to neighbourhood area type.

Figure 1 shows the rough locations of broker and resident surveys in Delhi. The sampling design has good coverage of Delhi with significant neighbourhood overlap between the two survey types.

4.2 Resident Survey

The resident survey covers 3,253 households. For each respondent it records demographic characteristics, religion and caste, the religious composition of the current neighbourhood, housing and neighbourhood satisfaction, and recent housing-search experiences, together with the experimental modules described below. Religion and caste are collected directly and used to construct the three identity groups that recur throughout the analysis: Hindu upper caste, Muslim,

Table 3: Religious composition of survey respondents

Religion	Frequency	Percent	Delhi (2011 Census)
Atheist	1	0.03	–
Buddhist	4	0.12	0.12
Christian	8	0.25	0.96
Hindu	2,708	83.25	80.21
Islam	496	15.25	12.78
Jain	9	0.28	1.39
Sikh	26	0.80	4.43
Did not state	1	0.03	0.08
Others	0	0.00	0.02
Total	3,253	100.00	100.00

Notes: Sample frequencies and shares are from the resident survey. Census shares are the religious distribution of Delhi in the 2011 Census of India (Office of the Registrar General and Census Commissioner, India, 2011). Columns may not sum to 100 due to rounding.

and Scheduled Caste.

Table 3 compares the religious composition of the sample with the religious composition of Delhi in the 2011 Census (Office of the Registrar General and Census Commissioner, India, 2011).

4.3 Broker Survey

The broker survey covers 615 rental brokers. Brokers are central informants about the supply side of the market: they mediate a large share of rental transactions and frequently screen prospective tenants before any landlord interaction. As discussed in Section 2.4, a broker typically serves an area of roughly 1.5 to 2 kilometres around their shop, so broker reports provide spatially resolved information about local market conditions. For each broker the survey records reported landlord preferences, the broker's own expectations about tenant acceptability, the rules imposed by local Resident Welfare Associations, identity-specific match probabilities, and the broker's experience with past tenancies. These identity-specific match probabilities are the basis for the spatial access measures used in Section 6.

4.4 Spatial Amenity Data

I assemble neighbourhood-level data on local amenities, including schools, hospitals, parks, religious institutions, and prevailing rents.[3] These measures serve two purposes. First, they construct the amenity index A_j and the rent measure r_j that enter household utility in equation (1), allowing observed sorting to be separated from amenity-driven sorting. Second, the location of religious institutions is used to classify neighbourhoods into area types, mosque versus temple areas, which underpins both the sampling strata and the area-type definitions used throughout the analysis.

4.5 Survey Instruments and Experimental Modules

Both surveys share a common instrument organised into three blocks. A *background* block records demographics, neighbourhood composition, housing satisfaction, and measures of exposure to neighbourhood disamenities. A *market-beliefs* block records landlord preferences, broker expectations, perceived matching probabilities by identity, and tenancy experiences. An *experiments* block contains the randomized modules used to measure the mechanisms of interest: an audit-style vignette experiment on supply-side discrimination, a neighbourhood-choice experiment on tenant homophily, and a public-goods game.[4]

4.6 Conflict Oversampling

Within each stratum, the design oversamples neighbourhoods with a documented history of communal conflict: two conflict-exposed neighbourhoods are drawn per stratum. Because conflict exposure is balanced against composition *within* stratum, this supports a within-stratum comparison of gatekeeping between conflict-exposed and other neighbourhoods, holding religious composition approximately fixed. I treat this as a design-based measure of neighbourhood conflict exposure, distinct from the self-reported conflict-driven moves recorded in the resident survey, which are subject to recall and selection and which I use only descriptively.

4.7 List Experiments

To measure discrimination free of social-desirability bias, the surveys include list (item-count) experiments. Respondents report only the *number* of statements on a list that apply to them; a

[3] Amenities data were scraped using keywords such as 'school' and 'mosque/masjid' from Google Maps; for prevailing rents I scraped MagicBricks.com, one of the prominent housing websites used in Delhi.

[4] The public-goods module is not central to the present draft and is reserved for future analysis.

randomly assigned treatment group sees one additional sensitive item. For tenants, the sensitive item records having been discriminated against in the rental market on the basis of identity; for landlords, it records having declined a prospective tenant on the basis of caste or religion. The difference in mean counts between treatment and control recovers the prevalence of the sensitive behaviour without any respondent having to admit it directly.

5 Measurement

This section describes how each object is measured. The framework (Section 3) requires two quantities that observed sorting confounds: own-group preference and identity-specific access. I recover each from a dedicated, randomized source of variation, and measure discrimination a second way, through a list experiment, to guard against social-desirability bias.

5.1 Own-Group Residential Preference

Own-group preference is measured with a neighbourhood-choice experiment, a stated-choice design in the spirit of conjoint analysis (Hainmueller et al., 2014). Respondents choose among hypothetical neighbourhoods that vary in rent, commuting time, and group composition, choosing as if every option were equally available, so the experiment isolates preference from the access constraints that bind in the field.

Table 4: Neighbourhood choice game

Area	**Rent (₹)**	**Commute Time (minutes)**	**Neighbour Demographics**
A	15,000	20	Mostly Upper-Caste Hindu
B	14,000	35	Mixed – Hindu, Muslim, Christian
C	13,000	25	Mostly Muslim

Table 4 shows an example. The respondent was told they have a choice between the three neighbourhoods, can afford the rent in each, and must decide which suits them best. The game was repeated four times with different rents, features, amenities, and demographics. On comprehension questions, accuracy was over 98%. Appendix B has details on all rounds.

Photo	A	B	C
	**Name:** Aarti Bhasin **Occupation:** Doctor **Rent:** ₹22,000 **Family:** Married, two children **Languages:** Hindi and English	**Name:** Atiya Rahman **Occupation:** Doctor **Rent:** ₹22,000 **Family:** Married, two children **Languages:** Hindi and English	**Name:** Atiya Rahman **Occupation:** Doctor **Rent:** ₹22,000 **Reference:** "Ideal tenant" according to previous landlord **Family:** Married, two children **Languages:** Hindi and English

Figure 2: Vignette experiment example

Letting respondent i choose among neighbourhoods j in task t, the deterministic utility is

$$\begin{aligned} V_{ijt} &= \beta_R \,\text{Rent}_{jt} \;+\; \beta_C \,\text{Commute}_{jt} \;+\; \theta_{MUS} \,\text{MuslimArea}_{jt} \;+\; \theta_{SC} \,\text{SCArea}_{jt} \\ &\quad +\; \eta_{HUC} \,\text{SameGroupHUC}_{ij} \;+\; \eta_{MUS} \,\text{SameGroupMuslim}_{ij} \;+\; \eta_{SC} \,\text{SameGroupSC}_{ij}, \end{aligned} \tag{10}$$

estimated by conditional logit with task fixed effects and standard errors clustered by respondent. The η terms are the group-specific own-group preferences. I report them as odds ratios $\exp(\eta_g)$.

Two design limitations must be stated plainly. First, the choice tasks vary attributes in fixed bundles rather than orthogonally, so while the own-group parameters η_g are identified from the identity×composition contrast, the rent and commute coefficients β_R, β_C are *not* cleanly identified. I therefore lead with odds ratios and treat willingness-to-pay magnitudes as illustrative at most. Second, each group's own-area option appears in few tasks, Muslim composition in two, Hindu upper caste and Scheduled Caste in one each, so η_{MUS} is the best-identified parameter. I corroborate the experimental estimates with respondents' stated ranking of own-group composition as a neighbourhood attribute.

5.2 Supply-Side Gatekeeping

I measure gatekeeping three ways and treat them as mutually corroborating.

Vignette experiment. Respondents evaluate hypothetical rental applicants and report whether they would accept each. The design is a factorial survey experiment in the sense of Auspurg and Hinz (2015): profiles are randomized across three arms, a benchmark identity (A), the same profile with identity changed (B), and the identity change plus a tenant-reliability signal (C).

Figure 2 shows an example. Respondents were asked to imagine they were landlords of a house, and were randomly allocated to seeing profile A, B, or C. The picture remained identical across the three arms. The prompt told them the applicant had applied to rent a house they advertised and is willing to pay the rent they set. They could 1. accept, 2. accept at a higher

rent (specifying the difference), 3. reject, or 4. wait x months and accept if no other application arrives. Each respondent saw eight profiles, ensuring a mix of applicant types. On comprehension questions accuracy was above 95%. For each profile respondents were asked to review afresh, so as not to assume the house was already let. Because respondents evaluate multiple vignettes in sequence, order effects are a known concern in this design (Auspurg et al., 2017); profile order was held fixed across arms, so any order effect is differenced out by the arm contrast. Appendix B has details on the profiles shown.

For the extensive-margin analysis, accept and accept-at-a-higher-rent are taken as acceptance and the other two as rejection. The results change little if only outright accepts and rejects are used. I estimate, for each clean identity transition and separately by respondent type,

$$\text{Accept}_i \;=\; \beta_0 \;+\; \beta_1\, \mathbf{1}\{B_i \text{ or } C_i\} \;+\; \beta_2\, \mathbf{1}\{C_i\} \;+\; \varepsilon_i, \tag{11}$$

where β_1 is the identity penalty and β_2 the incremental effect of the reliability signal. I describe this as a *vignette* rather than an audit: respondents answer as hypothetical landlords on hypothetical profiles, so the design carries hypothetical-bias and social-desirability concerns that a field audit would not, and that the list experiment below is included to address. Although in the Indian context, people are surprisingly open about their prejudices and openly talk about caste and religion biases, one might question that the respondents might want to appear secular and given that this is a hypothetical exercise, report higher acceptance for minorities than they would in practice. Even if that were true, the estimates can be considered a lower bound on the true discrimination.

Two features of the fielded instrument bound what (11) can say. The arm-C signal is *not* constant across profiles: it ranges from a landlord reference to police verification to a lease-length commitment to spousal employment, so β_2 is the effect of a heterogeneous signal, not of a single standardized reliability cue. I therefore report β_2 profile by profile rather than pooled, and distinguish profiles whose arm-C cue is a pure reliability signal from those whose cue also reveals ability to pay (Section 6.2). The clean identity transitions available in the instrument are between Hindu and Muslim profiles.

Broker-reported access. Brokers report, out of 100 representative properties in their service area, how many would accept a tenant of each of a range of identity and household types. These figures provide a spatially resolved measure of identity-specific access that maps into the access

probabilities of the framework, and, being reported by the gatekeepers themselves, are less subject to applicant social-desirability bias.

RWA rules. Brokers also report, for each of a set of tenant categories, whether local RWAs permit, prohibit, or have no rule on that category. I summarise the share of areas in which each category is reported "not allowed", comparing Muslim, Hindu upper-caste, and Scheduled-Caste applicants.

5.2.1 The intensive margin: reluctance and price

Equation (11) measures whether an applicant is refused. It discards the information in *how* a respondent accepts. The vignette instrument records four mutually exclusive responses, naturally ordered in the reluctance they express: accept at the offered rent; accept only at a higher rent; wait; reject. Collapsing these to a binary indicator throws away the distinction that Result 1 says is diagnostic, because a respondent who quotes a markup is discriminating on price while a respondent who refuses is discriminating on quantity.

The reluctance ladder. Let $y_i \in \{1, 2, 3, 4\}$ index the four responses in the order above, with higher values denoting greater reluctance. I model y_i as generated by a latent index

$$y_i^* \;=\; \beta_1 \,\mathbf{1}\{B_i \text{ or } C_i\} \;+\; \beta_2 \,\mathbf{1}\{C_i\} \;+\; \mathbf{W}_i'\boldsymbol{\delta} \;+\; \epsilon_i, \qquad \epsilon_i \sim N(0,1), \tag{12}$$

observed through thresholds $c_0 = -\infty < c_1 < c_2 < c_3 < c_4 = +\infty$, with $y_i = k$ if and only if $c_{k-1} < y_i^* \le c_k$. Estimated by ordered probit, $\hat{\beta}_1$ is the cleanest single summary of the identity effect: it uses every response category, imposes no conditioning on the sample, and is therefore not subject to the selection problem discussed below.

Decomposing the shift into price and quantity. Because (12) places all four responses on one latent scale, it also identifies *where* the identity shift sends respondents. Let

$$\Delta_k \;\equiv\; \frac{\partial \Pr(y_i = k)}{\partial\, \mathbf{1}\{B_i \text{ or } C_i\}}, \qquad k = 1, \ldots, 4, \tag{13}$$

denote the marginal effect of the identity change on each response probability, evaluated at sample means. Since probabilities sum to one, $\sum_k \Delta_k = 0$: the identity change moves mass out of acceptance at the offered rent ($\Delta_1 < 0$) and redistributes it, either to the price margin

($k = 2$) or to the quantity margin ($k = 3, 4$). I summarise the split by

$$\Lambda \;=\; \frac{\Delta_3 + \Delta_4}{\Delta_2 + \Delta_3 + \Delta_4} \;=\; \frac{\Delta_3 + \Delta_4}{-\Delta_1} \;\in\; [0, 1], \tag{14}$$

the share of the identity-induced shift that operates on quantity rather than price. $\Lambda \to 0$ describes a market that prices the disamenity and houses every applicant; $\Lambda \to 1$ describes one that rations. Equation (14) is the empirical counterpart of Result 1, and it is unit-free, which matters because a 19 percentage point refusal effect and a markup denominated in rupees cannot otherwise be compared. I report $\hat{\Lambda}$ by respondent type, with standard errors from a nonparametric bootstrap clustered on the respondent.[5]

The markup, and why it is a lower bound. Conditional on not refusing, I measure the price margin directly. Let $m_i \geq 0$ be the markup over the offered rent, observed only for respondents with $y_i \in \{1, 2\}$. Because m_i has a mass point at zero, a linear model is inappropriate (Tobin, 1958), and I estimate the two-part specification of Cragg (1971): a probit for whether any markup is quoted,

$$\Pr(m_i > 0 \mid y_i \leq 2) \;=\; \Phi\big(\beta_1^P \, \mathbf{1}\{B_i \text{ or } C_i\} + \beta_2^P \, \mathbf{1}\{C_i\} + \mathbf{W}_i'\boldsymbol{\delta}^P\big)\,, \tag{15}$$

and a Tobit for its magnitude, with $m_i = \max\{0, m_i^*\}$ and

$$m_i^* \;=\; \beta_1^T \, \mathbf{1}\{B_i \text{ or } C_i\} \;+\; \beta_2^T \, \mathbf{1}\{C_i\} \;+\; \mathbf{W}_i'\boldsymbol{\delta}^T \;+\; \nu_i. \tag{16}$$

The two-part form is preferred to a selection-corrected model because no exclusion restriction is available that would shift refusal without shifting the markup (Duan et al., 1983).

Equations (15) and (16) condition on $y_i \leq 2$, and that conditioning is not innocuous: it selects on the outcome of (11). Respondents with the strongest identity penalty refuse outright and never quote a price, so the sample surviving into the markup regressions is precisely the *less* discriminating part of the distribution. Writing $s_i = \mathbf{1}\{y_i \leq 2\}$ for selection, identity reduces $\Pr(s_i = 1)$ and does so most among respondents with large β_1, so $\mathbb{E}[\nu_i \mid s_i = 1, B_i] < \mathbb{E}[\nu_i \mid s_i = 1, A_i]$ and $\hat{\beta}_1^T$ is attenuated toward zero. **The estimated markup is therefore a lower bound on the price penalty a Muslim applicant would face in a market where refusal**

[5] Λ is a ratio of marginal effects from an ordered probit. Delta-method standard errors are available but depend on the internal naming of post-estimation results; the bootstrap is used because it does not.

were not available. This is not a defect of the design so much as a direct consequence of what the paper argues: in a market that rations, the price margin is mechanically hard to observe, because the respondents who would price most aggressively refuse instead.

Two things follow for how I report the results. I lead with $\hat{\beta}_1$ from the ladder (12), which conditions on nothing. And I report Lee (2009) trimming bounds on the price effect, estimated using the implementation of Tauchmann (2014). These bounds require no exclusion restriction and no parametric selection model, only the monotonicity condition that the identity change weakly reduces selection for every respondent, a condition Result 1 implies and that the reverse-transition estimate in Table 6 supports empirically. They are in the spirit of the partial-identification approach of Manski (1990): rather than point-identify an effect under assumptions the design cannot support, I report the range consistent with the data. Two caveats attach. The bounds assume a continuous outcome, and quoted markups cluster heavily on round numbers, so ties at the trimming quantile are a concern (Tauchmann, 2014). And because the estimator admits no covariates, it presumes random assignment of the treatment, which the arm randomisation delivers here by construction.

5.3 Discrimination Without Self-Report: List Experiments

The list experiments described in Section 4.7 provide a fourth measurement, and the one most robust to social-desirability bias. The treatment–control difference in mean item counts identifies the share of tenants who report having faced identity-based discrimination, and the share of landlords who admit having declined a tenant on caste or religious grounds, without any respondent stating so directly. Agreement between these indirect estimates and the direct vignette penalties is the key internal validity check on the gatekeeping results.

5.4 Perceived and Experienced Discrimination

From the resident survey I measure experienced discrimination (whether a respondent was denied housing and the reason they attribute to it) and perceived discrimination (the restrictions respondents expect, the landlord priorities they perceive, and whether they confine their search to areas where they anticipate acceptance). Comparing the two captures the anticipatory channel.

5.5 Amenities, Conflict, and the Sorting Outcome

Amenities enter through both stated rankings and the spatial measures of Section 4, and serve to separate amenity-driven from identity-driven sorting. Conflict exposure is the design-based, neighbourhood-level measure of Section 4.6. The sorting outcome classifies each household by residence in an own-group area.

6 Reduced-form results

6.1 Own-Group Preferences Are Real

I begin with model-free evidence from the neighbourhood-choice experiment. Because every respondent faces the same hypothetical menus, the share choosing the own-group alternative can be compared directly between own-group and out-group respondents. Based on equation (10), Panel A of Table 5 reports these shares. All three groups choose their own-group neighbourhood more often than others do. Upper-caste Hindus show the largest *raw* own-group preference (62.6 vs 47.1%), while Muslim and Scheduled Caste own-group areas are rarely chosen by outsiders (9.3% and 16.8%), producing the largest own-versus-other *gaps* (33.5 and 23.7 points). The two orderings carry different content: majority households have the strongest pull toward their own areas, whereas minority areas are the most avoided by non-members.

The conditional-logit estimates in Panel B confirm strong, significant own-group preference in every specification, but illustrate why I lead with the model-free shares. The Muslim odds ratio swings from 1.9 to 20.8 depending on whether area-type controls are included. I therefore report both specifications and rely on the significance and ordering rather than the literal magnitude. Stated rankings corroborate the experiment: 73% of Muslim, 70% of Scheduled Caste, and 50% of upper-caste Hindu respondents rank own-group composition among their three most important neighbourhood attributes.

6.2 Gatekeeping Against Muslim Applicants Is Severe

6.2.1 The Extensive Margin: Refusal

Table 6, based on equation (11), reports the vignette identity penalty for the clean Hindu→Muslim transitions. Pooling all respondents, switching an otherwise identical profile to a Muslim name lowers acceptance by 18.8 percentage points for a married woman and 22.9 percentage points

Table 5: Own-group residential preference

Panel A. Choice shares facing identical menus (%)			
Group	Own-group resp.	Out-group resp.	Gap (pp)
Hindu upper caste	62.6	47.1	15.5
Muslim	42.9	9.3	33.5
Scheduled Caste	40.4	16.8	23.7
Panel B. Conditional logit, own-group odds ratios			
	Interactions only	Area controls	
Own-group × Hindu UC	6.53***	2.35***	
	(0.38)	(0.17)	
Own-group × Muslim	1.90***	20.77***	
	(0.12)	(1.85)	
Own-group × SC	1.38***	3.02***	
	(0.10)	(0.27)	
Area-type controls	No	Yes	
Choice observations	22,752	22,752	

Notes: Panel A from the separation check facing identical menus. Panel B: conditional logit on the explicit-composition tasks of equation (10), one choice set per respondent-task, odds ratios with standard errors in parentheses. The Muslim odds ratio is highly sensitive to whether average area desirability is controlled, because Muslim areas are strongly disfavoured on average (estimated area main effect ≈ 0.08). The choice-share gaps in Panel A are the preferred, assumption-free summary. $^{***}p < 0.01$.

for a married man (both $p < 0.01$). The reverse transition (Muslim→Hindu) raises acceptance by 10.6 points, an internal-consistency check the design passes.

The penalty rises sharply across the chain of intermediaries. Figure 3 plots the Hindu→Muslim penalty by respondent type: it is -0.07 to -0.08 for respondents who are tenants, -0.24 to -0.30 for landlords, and -0.41 to -0.48 for brokers. Brokers, who screen applicants before landlords ever see them, are the most exclusionary link in the market, the result that anchors the enforcement argument of Section 8.

A caste benchmark. The instrument identifies caste only through surnames, and imperfectly. The one available caste-surname contrast replaces a Muslim name with "Manjhi", a Scheduled-Caste-associated surname. Acceptance *rises* by 14.0 points ($p < 0.01$) under this switch, indicating the Muslim name is penalised more heavily than the caste-marked one. This is consistent with Muslim identity being the most penalised, but it does not isolate a clean caste-versus-upper-caste penalty; the credible caste evidence comes from the broker and RWA

Table 6: Vignette identity penalty (Hindu → Muslim), by respondent type

	Pooled (1)	Tenant (2)	Landlord (3)	Broker (4)
Panel A. Married woman				
Identity penalty β_1	-0.188^{***}	-0.072^{**}	-0.242^{***}	-0.406^{***}
	(0.016)	(0.030)	(0.062)	(0.039)
Information effect β_2	0.002	0.010	−0.018	−0.048
	(0.018)	(0.031)	(0.049)	(0.049)
Panel B. Married man				
Identity penalty β_1	-0.229^{***}	-0.076^{**}	-0.299^{***}	-0.483^{***}
	(0.018)	(0.035)	(0.068)	(0.043)
Information effect β_2	−0.017	−0.053	0.022	−0.038
	(0.019)	(0.036)	(0.077)	(0.048)
Observations	3,868	1,032	263	615

Notes. OLS estimates of equation (11). The outcome is acceptance, coded one if the respondent accepts at the offered rent or at a higher rent and zero if the respondent waits or rejects. β_1 is the effect of altering applicant identity alone (arm A to arms B or C); β_2 is the incremental effect of adding a tenant-reliability signal (arm B to arm C). All respondents answer as hypothetical landlords. Robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

measures below.

6.2.2 The Intensive Margin: Rationing versus Pricing

Refusal is not the only way to discriminate. A landlord who dislikes a tenant can decline him, or he can quote him a higher rent. Which of the two he does is testable: Result 1 predicts that a market in excess demand discriminates on quantity, while a market that clears on price discriminates on price. The vignette instrument records four ordered responses, so the split can be measured directly.

Table 7 reports the reluctance ladder of equation (12). The identity shift raises reluctance sharply (+0.71 and +0.65, both $p < 0.01$), and the decomposition in Panel B shows where that shift goes. Of the 19 to 22 percentage points of probability mass displaced from acceptance at the offered rent, 18 to 21 points land on outright rejection and less than one point on acceptance at a higher rent. The rationing share is $\Lambda = 0.978$ and 0.965 for the two profiles, and between 0.957 and 0.967 across every respondent type (Table 8). Identity operates on this market almost entirely as a rationing device.

Table 7: Identity Penalties, Decomposition, and Price Margins

Panel A. Reluctance ladder, ordered probit

	(1) Married woman	(2) Married man
Identity penalty β_1	0.714***	0.654***
	(0.060)	(0.053)
Information effect β_2	-0.005	0.065
	(0.051)	(0.049)
Observations	3868	3868

Panel B. Decomposition of the identity shift

	(1) Married woman	(2) Married man
D1	-0.192***	-0.223***
	(0.016)	(0.017)
D2	0.004***	0.008***
	(0.001)	(0.001)
D3	0.008***	0.004***
	(0.001)	(0.001)
D4	0.180***	0.211***
	(0.015)	(0.016)
Dprice	0.004***	0.008***
	(0.001)	(0.001)
Dquant	0.188***	0.215***
	(0.015)	(0.016)
Lambda	0.978***	0.965***
	(0.005)	(0.005)
β_1	0.714***	0.653***
	(0.061)	(0.054)
Observations	2473	2473

Panel C. Price margin, conditional on not refusing

	(1) P(markup) f	(2) Markup Rs f	(3) Markup % f	(4) P(markup) m	(5) Markup Rs m	(6) Markup % m
Identity penalty β_1	0.386**	8080.780**	36.731**	0.063	457.076	4.571
	(0.179)	(3996.744)	(18.167)	(0.114)	(1809.041)	(18.090)
Information effect β_2	0.027	1061.169	4.823	0.250**	4394.544**	43.945**
	(0.148)	(3317.853)	(15.081)	(0.110)	(1796.220)	(17.962)
Observations	2967	2967	2967	2528	2528	2528

Notes: Panel A estimates equation (12) by ordered probit, with responses ordered accept-at-offered < accept-higher < wait < reject, so that higher values denote greater reluctance. β_1 and β_2 are defined as in table 6. Panel B reports the marginal effect of the identity change on each response probability, evaluated at sample means; the Δ_k sum to zero by construction. $\Lambda = (\Delta_3 + \Delta_4)/(-\Delta_1)$ is the share of the identity-induced shift operating on quantity rather than price, with bootstrap standard errors. Panel C conditions on the respondent not refusing, and estimates equations (15) and (16). Because respondents with the largest identity penalty refuse rather than quote a price, the Panel C sample is selected toward the less discriminating part of the distribution and the estimates are lower bounds on the price penalty (Section 5.2.1). Lee (2009) trimming bounds are reported beneath the Tobit estimate and in table 24. The married-man markup loads on a later experimental arm and is correspondingly more heavily selected. The expected-waiting-time margin is omitted pending correction of out-of-range values in that variable. Waiting enters Λ only through the response category Δ_3, which is identified independently of the reported duration. Robust Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

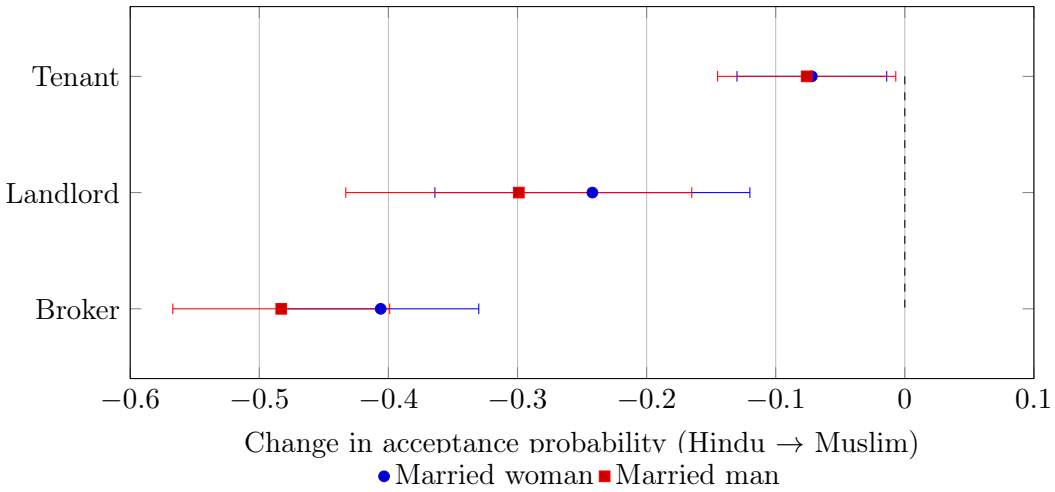


Figure 3: The Hindu→Muslim acceptance penalty by respondent type, with 95% confidence intervals. Brokers are the most exclusionary link.

Two qualifications discipline that reading. First, the price option is rarely used by *anyone*: only 0.9 and 3.0 percent of respondents choose acceptance at a higher rent, in treatment and control alike, and the estimated cutpoints separating the interior categories span 0.08 of the latent scale. Λ near one therefore reflects both that respondents seldom price and that the identity shift goes to refusal. The two facts point the same way, but they are distinct, and I report the response distribution alongside Λ so the reader can see which is doing the work.

Second, the price margin is thinly populated and its estimates should not be read as the paper's evidence on pricing. Only 35 respondents quote a markup for the married-woman profile and 115 for the married man. Among that small group the Muslim profile is more likely to face a markup (probit $\hat{\beta}_1^P = 0.386$, $p = 0.031$, a marginal effect of roughly [ ] percentage points on a base rate of 1.2 percent), with a Tobit coefficient of about ₹8,000. But because the respondents most inclined to discriminate refuse rather than quote, this sample is selected toward the least discriminating part of the distribution. Lee (2009) trimming bounds, reported in Appendix Table 24, place the effect between ₹48 and ₹116 for the married woman and between −₹40 and ₹342 for the married man, with confidence intervals that include zero in both cases. **After accounting for selection, the sign of the price penalty is not identified.** That is not a failure of the design but a consequence of what the paper argues: where refusal is available and costless, the price margin is mechanically unobservable, because the gatekeepers who would price most aggressively never name a price.

Table 8: Does the market ration or price? Rationing share by respondent type

	(1) Pooled	(2) Tenant	(3) Landlord	(4) Broker
beta1	0.677***	0.319***	0.999***	1.369***
	(0.045)	(0.085)	(0.150)	(0.119)
Dprice	0.007***	0.003***	0.012***	0.017***
	(0.001)	(0.001)	(0.004)	(0.004)
Dquant	0.200***	0.097***	0.294***	0.372***
	(0.013)	(0.025)	(0.039)	(0.025)
Lambda	0.967***	0.965***	0.961***	0.957***
	(0.004)	(0.008)	(0.013)	(0.012)
Observations	4946	1390	496	772

Each column estimates equation (12) on the indicated subsample and reports the decomposition of equation (13). $\Lambda \in [0, 1]$ is the share of the identity-induced shift out of acceptance-at-offered-rent that lands on the quantity margin (wait or reject) rather than the price margin (accept at a higher rent), as defined in equation (14). $\Lambda \to 0$ characterises a market that prices the disamenity and houses every applicant; $\Lambda \to 1$ characterises one that rations. Delta-method standard errors in parentheses. Profiles are pooled; **??** reports the split by profile. Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Information moves prices but not reliability judgements. The two profiles differ in what their arm-C signal contains, and the contrast is informative. For the married woman the signal is a pure reliability cue, a previous landlord describing her as an ideal tenant, and it moves nothing: $\hat{\beta}_2 = -0.005$ on the ladder ($p = 0.93$) and 0.027 on the markup probit ($p = 0.85$). For the married man the signal instead reveals that his wife is head of human resources at a multinational company, which is an income signal rather than a reliability signal, and against a ₹10,000 rent a large one. It leaves overall reluctance unchanged ($\hat{\beta}_2 = 0.065$, $p = 0.18$) but raises the probability of a markup (0.250, $p = 0.023$) and its size by about ₹4,400 ($p = 0.014$).

Respondents therefore do update on economically relevant information: told the applicant's household income is high, they charge more. They do not update on reliability information when the applicant is Muslim. This rules out the leading alternative reading of the null in Table 6, that respondents skim the profiles and ignore whatever is added to them, and it is the cleanest evidence in the paper that the reliability null reflects the mechanism rather than inattention. As Result 3 makes clear, however, it still does not separate taste from norm.

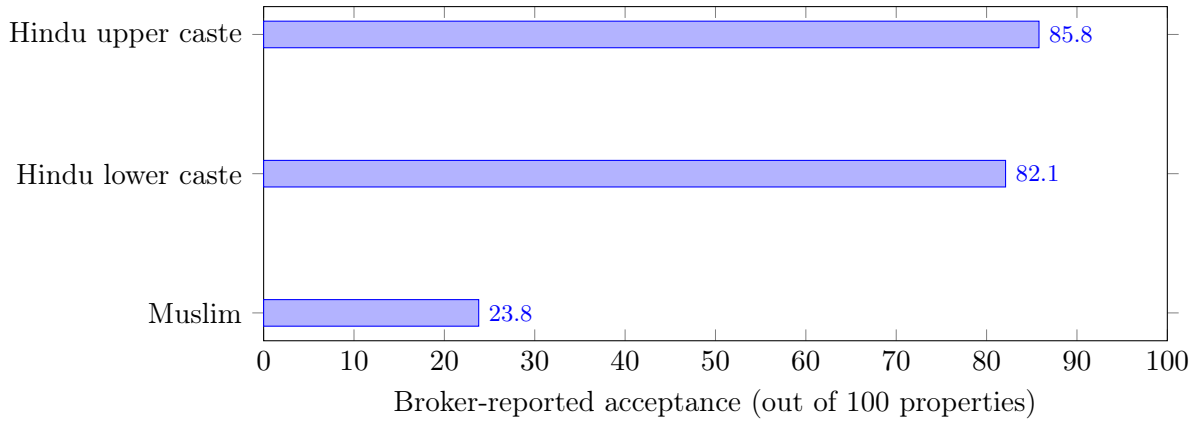


Figure 4: Broker-reported acceptance by applicant type (out of 100 properties). Muslim applicants face by far the lowest access.

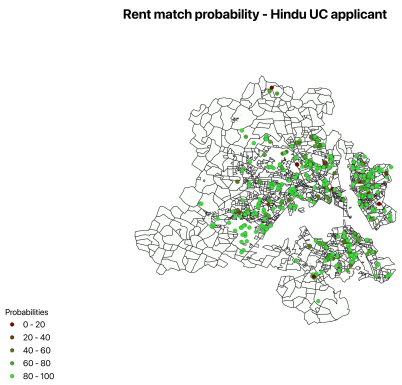
6.2.3 Broker Access and RWA Rules

The vignette evidence is corroborated by the gatekeepers' own reports. Figure 4 shows broker-reported acceptance out of 100 properties by applicant type. The Muslim figure (23.8) is the lowest by a wide margin, against 85.8 for upper-caste and 82.1 for lower-caste Hindus. The Muslim gap relative to upper-caste Hindus is 62.0 points ($p < 0.01$); the lower-caste Hindu gap is only 3.7 points, so the broker data locate the overwhelming exclusion on religion, not caste. Resident associations tell the same story: they are reported to bar Muslim tenants in 26.0% of areas, against 0.7% for upper-caste and 0.8% for lower-caste Hindus.

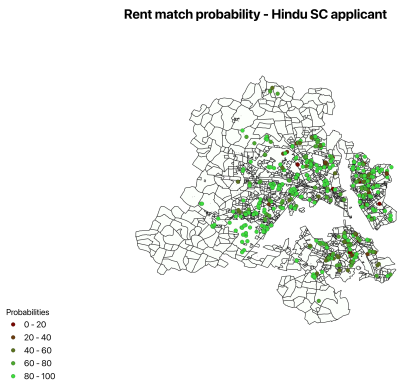
The probabilities are also area-specific. Figure 5 shows the broker-reported probabilities of acceptance for Hindu upper caste, Hindu Scheduled Caste, and Muslim applicants. Figure 6 shows the same for broker-reported RWA rules. Both point to significant access constraints facing Muslim applicants.

6.3 The Wedge: Exclusion Where Minorities Would Move

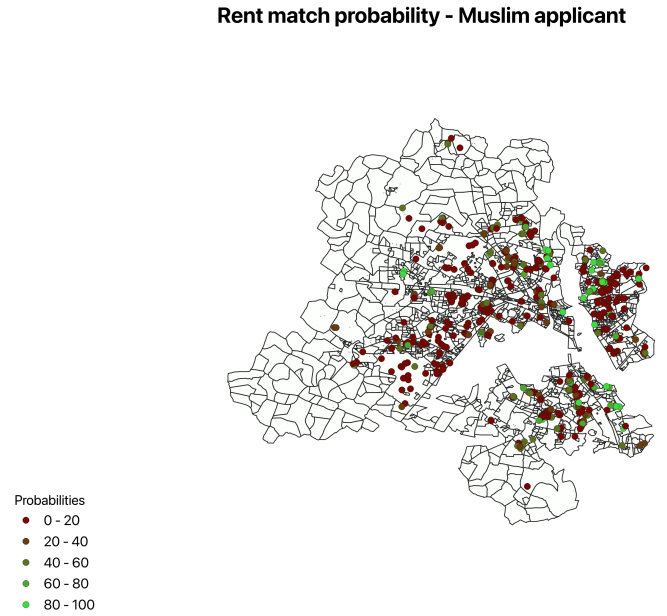
The central result combines preferences and access. In the choice experiment, Muslim respondents select a non-Muslim (mixed or majority-Hindu) neighbourhood 57% of the time, a clear revealed willingness to integrate. Yet brokers report a Muslim applicant can secure housing in only about 24 of 100 properties. The gap between a majority willing to live in mixed or majority areas and access to roughly a quarter of the market is the wedge between desired and feasible sorting. A portion of observed Muslim concentration therefore reflects not preference but the foreclosure of accessible alternatives: segregation as a constrained equilibrium, not revealed preference.



(a) Hindu Upper Caste

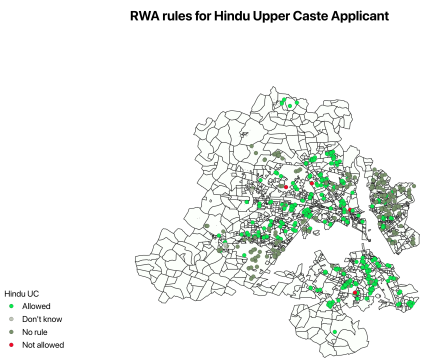


(b) Hindu Scheduled Caste

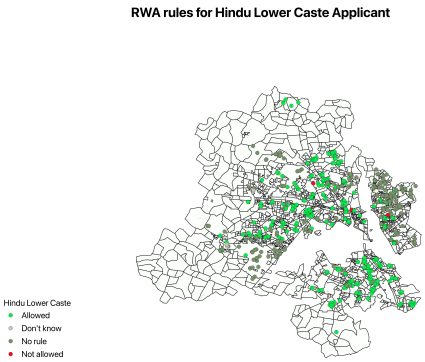


(c) Muslim

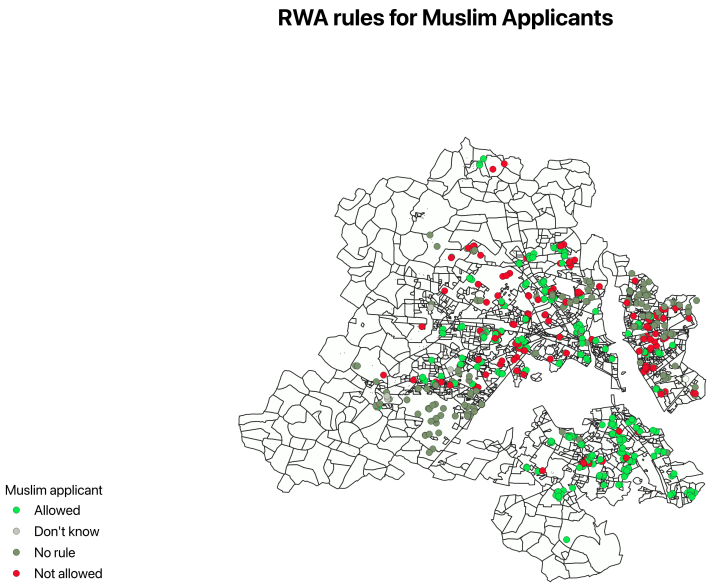
Figure 5: Broker-reported probabilities of finding a house in their area.



(a) Hindu Upper Caste



(b) Hindu Scheduled Caste



(c) Muslim

Figure 6: Broker-reported RWA rules on tenant acceptance, by identity.

Table 9: Perceived, experienced, and admitted discrimination

Panel A. Survey measures, by respondent religion		
	Non-Muslim	Muslim
Experienced a denial	15.9%	17.2%
of which attributed to religion	20.8%	42.5%
Ranks religion in landlords' top-3 factors	59.8%	62.6%
Panel B. List experiments (prevalence of sensitive item)		
Tenant: faced identity discrimination	0.011	(0.083)
Landlord: declined a tenant on identity	0.095	(0.170)
Attitudinal items (whole sample):		
Supports caste/religion reservation	0.227***	(0.063)
Believes in the Manusmriti caste hierarchy	0.154**	(0.062)
Believes India should be a Hindu nation	0.501***	(0.068)

Notes: Panel B reports the treatment−control difference in mean item counts (a prevalence), robust standard errors in parentheses. Treated cells are small (tenants 486, landlords 114), so the discrimination nulls rule out only *large* admitted rates. Attitudinal items pool all respondents; restricting to Hindu respondents is a planned refinement. $^{**}p < 0.05$, $^{***}p < 0.01$.

6.4 Perceived, Experienced, and Admitted Discrimination

Three further measures characterise the *nature* of the discrimination (Table 9). Experienced denial is common and similar across groups in the raw (16–17%), but among those denied, Muslims attribute the denial to religion far more often (42.5 vs 20.8%). Perceived discrimination is widespread: about 60% of respondents rank religion among the three most important factors landlords use. Yet the list experiments show almost no *admitted* discrimination: the prevalence of tenants reporting having faced identity discrimination is a precise zero (0.01, s.e. 0.08), and of landlords admitting to having declined a tenant on identity grounds is likewise insignificant (0.10, s.e. 0.17).

The contrast is a finding in itself. Direct measures show severe discrimination; the social-desirability-robust measure shows that few will *admit* to it even anonymously. Discrimination here is not owned as a personal choice but enacted as a norm. The same indirect method, applied to attitudes, uncovers what respondents do suppress: about half the sample indirectly endorses the proposition that India should be a Hindu nation, and sizeable minorities the Manusmriti caste hierarchy, a reservoir of majoritarian sentiment that is the attitudinal backdrop to the gatekeeping documented above.

7 Structural Estimation and Welfare (In progress)

7.1 From access constraints to damages

The reduced-form results establish that access differs sharply by identity. They do not say what that costs, because the cost of losing a property depends on what the household would have obtained from it. A household excluded from neighbourhoods it did not want loses little; a household excluded from precisely the neighbourhoods it values most loses a great deal. Converting access constraints into welfare therefore requires utility weights over neighbourhood attributes, which the reduced form does not supply.

Christensen and Timmins (2023) solve this problem for the U.S. rental market. They embed experimentally identified access constraints in a residential sorting model in the tradition of Bayer et al. (2007), treat the constraint as a restriction on the set of alternatives a household can consider, and compute the income compensation that would make a constrained renter as well off as an unconstrained one. For U.S. renters they estimate damages of 4.6 and 4.7 percent of annual income for African American and Hispanic/LatinX renters respectively. The framework I use here follows theirs in its object and its identification strategy, and departs from it in implementation for reasons set out in Section 7.3.

What is borrowed, and from where. The idea of treating discrimination as a restriction on the choice set rather than a shift in preferences, and of monetising it through the resulting change in expected utility, is due to Christensen and Timmins (2023); so is the decision to report the *differential* constraint faced by a minority household relative to a majority one, rather than the total constraint. The sorting framework into which the constraint is embedded is the equilibrium residential-sorting tradition of Bayer et al. (2007). The choice model itself is the conditional logit of McFadden (1974), and the treatment of access as latent availability of alternatives follows Manski (1977). The welfare formula is the log-sum compensating variation of **?**. What is new here is the source of the access measure, which is elicited from gatekeepers rather than inferred from correspondence responses, and the validation exercise in Section 7.6, which the U.S. data cannot support.

The elicitation difference is worth stating precisely. Christensen and Timmins (2023) recover the access probability from responses to inquiries about *observed listings*, aggregated to the census tract. Mine comes from the intermediaries who hold the inventory and is elicited directly

as a share of the local stock. It therefore captures properties an applicant would never learn exist, the non-showing margin of Section 3, which a correspondence design cannot reach by construction because a listing must be visible before an inquiry can be sent.

7.2 Utility and access

Let $n = 1, \ldots, K$ index neighbourhoods and i index tenant households, each belonging to a group $g(i) \in \{\text{HUC}, \text{MUS}, \text{SC}\}$. Household i derives utility from neighbourhood n:

$$U_{in} \;=\; \underbrace{-\alpha R_n + \mathbf{X}_n'\boldsymbol{\gamma} + \eta_{g(i)}\, s_{g(i),n}}_{V_{in}} + \varepsilon_{in}, \qquad \varepsilon_{in} \sim \text{i.i.d. type-I EV}, \tag{17}$$

where R_n is the prevailing rent, $\mathbf{X}_n$ collects neighbourhood attributes, $s_{g,n}$ is the share of n's residents in group g, and η_g is group-specific own-group preference. Equation (17) is (1) with rent entering with an estimated coefficient and unobserved quality absorbed into x_n, for reasons given in Section 7.7.

Rent enters linearly rather than through $\ln(I_i - R_n)$. This departs from Christensen and Timmins (2023) and is deliberate. Income in my data is observed only as an interval, recovered from reported rent and a coarse rent-burden category, and for a quarter of tenants the interval is unbounded above. Since the marginal utility of income enters the denominator of every willingness-to-pay and every welfare figure, estimating it from an interval-censored variable would propagate that imprecision through the entire section. Entering rent linearly makes α the marginal utility of income directly, requires no income measure for estimation, and has a further advantage noted below: constant marginal utility of income is precisely the condition under which the **?** welfare measure is exact rather than approximate. Income re-enters only when damages are expressed as a share of it, where the interval can be reported transparently.

Access. Let $p_{g,n} \in (0, 1]$ denote the probability that a household of group g can secure a rental in neighbourhood n, estimated by the broker access measure of (7). Following Manski (1977), I treat $p_{g,n}$ as the latent availability of alternative n to group g, so that the probability household i resides in n is the availability-weighted logit

$$P_{in} \;=\; \frac{p_{g(i),n}\exp(V_{in})}{\sum_m p_{g(i),m}\exp(V_{im})}, \tag{18}$$

which is (8) of the framework and collapses to the standard conditional logit when $p \equiv 1$.

7.3 Estimation

Equation (18) has a convenient form. Since $p_{g,n} \exp(V_{in}) = \exp\big(V_{in} + \ln p_{g,n}\big)$, the model is a conditional logit in which $\ln p_{g,n}$ enters as an *offset*: a regressor whose coefficient is constrained to unity. Estimation therefore requires no simulation of consideration sets, no choice of a number of draws, and no simulated-likelihood machinery.

This is where I depart from Christensen and Timmins (2023), who simulate consideration sets directly and weight them by their probabilities, following the applied consideration-set literature (Gaynor et al., 2016). Their implementation is the more literal representation of a random choice set; the availability-weighted form in (18) is the closed-form counterpart under the latent-availability interpretation of Manski (1977). With roughly [. . .] households choosing among $K = [\ldots]$ neighbourhoods, the closed form is the appropriate choice: the simulated version would add computational and inferential burden without adding identifying variation. I report the simulated version as a robustness check in Appendix [. . .].

A specification test the simulated version does not offer. The offset imposes a coefficient of one on $\ln p_{g,n}$. Relaxing that restriction and estimating

$$V^*_{in} \;=\; V_{in} + \lambda \ln p_{g,n} \tag{19}$$

provides a direct test of whether the elicited access measure enters choices in the way the model requires. Under (18), $\lambda = 1$. A $\hat{\lambda}$ near one supports both the availability structure and the broker elicitation; a $\hat{\lambda}$ far from one indicates that the access measure is mismeasured, or that access affects sorting through a channel other than availability. I report $\hat{\lambda}$ and the test of $\lambda = 1$ alongside the main estimates.

7.4 Damages

Expected utility from the unrestricted choice set is the standard inclusive value, and from the restricted set the availability-weighted inclusive value:

$$IV_i^{\text{unc}} = \ln \sum_n \exp(V_{in}), \qquad IV_i^{\text{con}} = \ln \sum_n p_{g(i),n} \exp(V_{in}). \tag{20}$$

Because rent enters linearly, the marginal utility of income is constant, and the compensating variation of **?** applies exactly:

$$CV_i \;=\; \frac{1}{\alpha}\left[IV_i^{\text{unc}} - IV_i^{\text{con}}\right] \;\geq\; 0. \tag{21}$$

Since $p_{g,n} \leq 1$ for all n, (21) is non-negative by construction, and it is denominated in rupees per month. No numerical solution is required.

Isolating discrimination. Following Christensen and Timmins (2023), I report the *differential* constraint rather than the total. For each Muslim household I compute (21) twice, once with $p_{M,n}$ and once with the access probabilities an upper-caste Hindu household would face, $p_{H,n}$, holding preferences and neighbourhood attributes fixed:

$$CV_i^{\text{disc}} \;=\; \frac{1}{\alpha}\left[\ln \sum_n p_{H,n}\exp(V_{in}) - \ln \sum_n p_{M,n}\exp(V_{in})\right]. \tag{22}$$

This is the damage attributable to discrimination over and above the search frictions any household in this market faces. Reporting the total constraint instead would conflate discrimination with ordinary market tightness.

Units. Reporting CV^{disc} in rupees per month, rather than as a share of annual income, has two advantages here. It requires no income measure, so the interval-censoring described above does not enter. And it is directly comparable to the reduced-form price penalty of Section 6.2.2: if the welfare cost of the access constraint is large relative to the markup a Muslim applicant faces conditional on acceptance, that is the rationing argument of Result 1 restated in money. I also report CV^{disc} as a share of income, evaluated at the lower bound, midpoint and upper bound of the implied income interval, and present it as a range.

7.5 Distortion in revealed preferences

The second structural result of Christensen and Timmins (2023) is that a model ignoring choice-set constraints yields biased willingness-to-pay estimates: systematic exclusion from high-amenity neighbourhoods makes an excluded group appear to place little value on those amenities. In the present specification that comparison is immediate, since the constrained and unconstrained models differ only by the offset. Marginal willingness to pay for attribute x is

γ_x/α in both, and I report it under each.

This has a specific implication for the Indian setting. If Muslim households are excluded from high-amenity neighbourhoods, a revealed-preference reading of their observed locations will understate their valuation of schools, health facilities and infrastructure. Since such estimates are used to justify the allocation of local public goods, exclusion in the housing market can propagate into under-provision in the neighbourhoods where excluded households end up, reinforcing the descriptive pattern Asher et al. (ming) document through a mechanism distinct from the political-economy explanations usually offered for it.

7.6 Validating the constraint correction

This paper is in an unusual position relative to Christensen and Timmins (2023). Their consideration-set estimates cannot be benchmarked against a constraint-free measure of preferences, because none exists in their data. Here one does: the neighbourhood-choice experiment of Section 5.1 elicits preferences under an explicit instruction to choose as if every option were available, and is therefore constraint-free by construction.

This permits a direct external validation. Estimating (17) on observed residence twice, with and without the offset, and comparing the implied own-group parameters η_g to the stated-choice estimates of Table 5, asks whether the correction recovers what an unconstrained elicitation reports. If the corrected estimates align with the stated-choice estimates while the naive ones do not, that is evidence the correction does what it claims. To my knowledge no existing application of a choice-set correction to discrimination has been able to run this test, and a null result would be equally informative about the limits of the method.

7.7 Caveats

Commuting. Equation (1) includes a commuting term. The survey records commuting time only for the household's current residence, so c_{in} is not observed for counterfactual neighbourhoods and the term cannot be identified. Commuting cost is therefore absorbed into $\mathbf{X}_n$ through accessibility attributes, and η_g should be read as own-group preference net of any correlation between composition and commuting cost.

Stock versus flow. Christensen and Timmins (2023) observe households that *moved* in a given year, so choices are made under contemporaneous constraints. My resident survey records current residence, which for long-tenure households reflects constraints prevailing at the time

of the move. I restrict the estimation sample to households that moved within [...] years and report the full sample as a robustness check.

The access measure. $p_{g,n}$ is a broker report, not an experimental estimate. It has the coverage advantage described above and the credibility disadvantage that it is stated rather than revealed. I report welfare under the raw and shrunk versions of the measure, and separately under access constructed from the vignette penalties of Table 6, which are independently randomized. Agreement across the three is the relevant check; $\hat{\lambda}$ from (19) is a fourth.

Rents. R_n is a quality-adjusted index of transacted rents, and it is endogenous to the exclusion being studied (Section 3). Equalising access would shift demand across neighbourhoods and rents would adjust, partly offsetting the gain (Bayer et al., 2007). The counterfactuals are therefore partial equilibrium and should be read as an upper bound on the welfare gain from removing the constraint.

Composition. The survey-based and institution-based measures of neighbourhood religious composition agree only weakly in these data. I use the survey measure throughout and report the alternative in Appendix [...].

8 Policy: Discrimination as a Tractable Enforcement Problem

8.1 Extending the existing agent-registration regime

The evidence locates the most exclusionary conduct in brokers, who screen applicants before landlords ever see them and whose penalty exceeds the landlord's despite their bearing none of the risk. Brokers are also the most regulable of the three gatekeepers: they operate from fixed premises, serve compact and identifiable areas, and are already a regulated category in an adjacent market.

Section 9 of the Real Estate (Regulation and Development) Act 2016 requires real estate agents to register with the state authority before facilitating a transaction. Delhi notified its rules in November 2016; registration runs five years, renewable, and operating unregistered attracts a penalty of ₹10,000 per day of default. But the obligation attaches to transactions in *registered projects*, the primary sale market. It does not reach secondary-market rental brokerage. Delhi's rental brokers are in consequence essentially unregulated, while a functioning registration, renewal and penalty architecture already exists one market over.

The recommendation follows directly: extend RERA agent registration to rental interme-

diaries, and condition the licence on non-discrimination in tenant screening. This requires no new statute, no new authority, and raises none of the constitutional difficulty that attaches to regulating housing societies' internal bye-laws (Section 2.5). Enforcement is practical because brokers are spatially fixed and few per neighbourhood, which makes the audit design used in the research literature directly available as a monitoring tool.

8.2 On RWA rules

Zoroastrian forecloses the easy argument that identity-based society bye-laws are simply unlawful. Two narrower routes remain. Rules operating on *tenancies* rather than on membership sit on weaker ground than the membership restriction upheld in *Zoroastrian*, since the tenant is not a party to the bye-laws he is being excluded by. And where an RWA exercises functions delegated by a municipal authority, its conduct may be more readily amenable to review. Both are contestable, and I present them as open questions rather than settled law.

8.3 Where enforcement binds hardest

The wedge result locates the highest-value target: mixed and majority neighbourhoods where minority households reveal willingness to live but gatekeeping is most intense. Returns to enforcement are not uniform across space, and the data identify where they are highest.

9 Conclusion

This paper asks how much of residential segregation in an urban rental market reflects what households want, and how much reflects what the supply side will permit. By measuring own-group preferences and identity-specific access on the same households and neighbourhoods, through randomized choice and vignette experiments, gatekeeper elicitation, and a social-desirability-robust list experiment, it separates two forces that the prior literature has only been able to study one at a time.

Three findings emerge. Own-group preferences are real and, for Muslim households, substantial. Supply-side gatekeeping against Muslim applicants is severe, and it operates almost entirely through refusal rather than price: of the response shifted by an identity change, roughly 97 percent lands on rejection and under 3 percent on a higher quoted rent. That pattern is what a market in excess demand looks like, and it is corroborated by brokers' own reports, RWA rules,

and an indirect list experiment. And the two forces do not coincide: exclusion is concentrated precisely in the mixed and majority neighbourhoods where minority households would be willing to live, so observed segregation is a constrained equilibrium, more separated than preferences alone would produce.

The contribution is both methodological and practical. Methodologically, the paper shows how to identify the wedge between desired and feasible sorting by measuring preferences and access separately rather than inferring one from the other, a strategy that travels beyond this setting. Practically, it locates the binding constraint in identifiable institutions, brokers and RWAs, whose conduct is a tractable target for anti-discrimination enforcement.

Several questions remain open. The welfare costs of the misallocation this paper documents, through commutes, amenity access, and the spatial concentration of skill, are measured only partially and warrant dedicated estimation. The structural model invites richer counterfactuals once the preference parameters are identified from stronger attribute variation. And whether gatekeeping responds to local conflict salience, which the sampling design is built to test, remains to be estimated.

References

Adida, C. L., Laitin, D. D., and Valfort, M.-A. (2010). Identifying barriers to Muslim integration in France. *Proceedings of the National Academy of Sciences*, 107(52):22384–22390.

Adukia, A. (2024). What ghettoisation has cost Indian Muslims. The Indian Express. Accessed: 2026-06-07.

Ahmed, A. M., Andersson, L., and Hammarstedt, M. (2010). Can discrimination in the housing market be reduced by increasing the information about the applicants? *Land Economics*, 86(1):79–90.

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

Alesina, A., Baqir, R., and Easterly, W. (1999). Public goods and ethnic divisions. *Quarterly Journal of Economics*, 114(4):1243–1284.

Asher, S., Jha, K., Novosad, P., Adukia, A., and Tan, B. (forthcoming). Residential segregation

and unequal access to local public services in India: Evidence from 1.5m neighborhoods. *American Economic Review*. NBER Working Paper 34818.

Asplund, J., Eslami, M., Sundaram, H., Sandvig, C., and Karahalios, K. (2020). Auditing race and gender discrimination in online housing markets. In *Proceedings of the Fourteenth International AAAI Conference on Web and Social Media*.

Auspurg, K. and Hinz, T. (2015). *Factorial Survey Experiments*, volume 175 of *Quantitative Applications in the Social Sciences*. Sage Publications, Thousand Oaks, CA.

Auspurg, K., Hinz, T., and Schmid, L. (2017). Contexts and conditions of ethnic discrimination: Evidence from a field experiment in a German housing market. *Journal of Housing Economics*, 35:26–36.

Bayer, P., Ferreira, F., and McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638.

Berry, M. C. and Wiener, R. L. (2020). Ex-offender housing stigma and discrimination. *Psychology, Public Policy, and Law*, 26(2):213–232.

Bertrand, M. and Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? a field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013.

Black, D. A. (1995). Discrimination in an equilibrium search model. *Journal of Labor Economics*, 13(2):309–334.

Bowlus, A. J. and Eckstein, Z. (2002). Discrimination and skill differences in an equilibrium search model. *International Economic Review*, 43(4):1309–1345.

Chandran, R. (2017). No Muslims, no single women: Housing bias turning Indian cities into ghettos. The Indian Express. Accessed: 2026-06-07.

Chetty, R., Hendren, N., and Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment. *American Economic Review*, 106(4):855–902.

Christensen, P., Sarmiento-Barbieri, I., and Timmins, C. (2021). Racial discrimination and housing outcomes in the united states rental market. NBER Working Paper 29516, National Bureau of Economic Research.

Christensen, P. and Timmins, C. (2022). Sorting or steering: The effects of housing discrimination on neighborhood choice. *Journal of Political Economy*, 130(8):2110–2163.

Christensen, P. and Timmins, C. (2023). The damages and distortions from discrimination in the rental housing market. *The Quarterly Journal of Economics*, 138(4):2505–2557.

Cragg, J. G. (1971). Some statistical models for limited dependent variables with application to the demand for durable goods. *Econometrica*, 39(5):829–844.

Datta, S. and Pathania, V. (2016). For whom does the phone (not) ring? discrimination in the rental housing market in Delhi, India. WIDER Working Paper 2016/55, UNU-WIDER, Helsinki.

Di Stasio, V., Lancee, B., Veit, S., and Yemane, R. (2021). Muslim by default or religious discrimination? results from a cross-national field experiment on hiring discrimination. *Journal of Ethnic and Migration Studies*, 47(6):1305–1326.

Duan, N., Manning, W. G., Morris, C. N., and Newhouse, J. P. (1983). A comparison of alternative models for the demand for medical care. *Journal of Business and Economic Statistics*, 1(2):115–126.

Flage, A. (2018). Ethnic and gender discrimination in the rental housing market: Evidence from a meta-analysis of correspondence tests, 2006–2017. *Journal of Housing Economics*, 41:251–273.

Galonnier, J. (2015). The enclave, the citadel and the ghetto: The threefold segregation of upper-class Muslims in India. *International Journal of Urban and Regional Research*, 39(1):92–111.

Gandhi, S., Green, R. K., and Patranabis, S. (2021). India's housing vacancy paradox: How rent control and weak contract enforcement produce unoccupied units and a housing shortage at the same time. Csep working paper, Centre for Social and Economic Progress. Revised 2023.

Gayer, L. and Jaffrelot, C. (2012). Muslims in Indian cities: Trajectories of marginalisation. Columbia University Press.

Gaynor, M., Propper, C., and Seiler, S. (2016). Free to choose? reform, choice, and consideration sets in the English National Health Service. *American Economic Review*, 106(11):3521–3557.

Gbadegesin, J. T. and Ojo, O. (2013). Ethnic bias in tenant selection in metropolitan ibadan private rental housing market. *Property Management*, 31(2):159–178.

Hainmueller, J., Hopkins, D. J., and Yamamoto, T. (2014). Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis*, 22(1):1–30.

Hsieh, C.-T. and Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39.

Human Rights Watch (2024). India: Increased abuses against minorities, critics. Human Rights Watch World Report 2024. Accessed: 2026-06-07.

Ifeanacho, K., Ewurum, N., and Aguome, N. (2023). Discrimination in the housing market: A literature review. In *Proceedings of the European Real Estate Society (ERES) Annual Conference*. ERES eres2023_186.

Knight Frank (2019). India real estate: Residential and office market report. Knight Frank Research.

Lang, K. and Lehmann, J.-Y. K. (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature*, 50(4):959–1006.

Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies*, 76(3):1071–1102.

Logan, T. D. and Parman, J. M. (2025). Racial residential segregation in the united states. *Journal of Economic Literature*, 63(3):964–1010.

Mander, H. (2025). Why Indian Muslims live apart. Scroll.in. Accessed: 2026-06-07.

Manski, C. F. (1977). The structure of random utility models. *Theory and Decision*, 8(3):229–254.

Manski, C. F. (1990). Nonparametric bounds on treatment effects. *American Economic Review*, 80(2):319–323.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, pages 105–142. Academic Press, New York.

Moretti, E. (2011). Local labor markets. *Handbook of Labor Economics*, 4:1237–1313.

Office of the Registrar General and Census Commissioner, India (2011). Census of india 2011: Religion. Government of India, Ministry of Home Affairs.

Ondrich, J., Stricker, A., and Yinger, J. (1999). Do landlords discriminate? the incidence and causes of racial discrimination in rental housing markets. *Journal of Housing Economics*, 8(3):185–204.

Pager, D. and Shepherd, H. (2008). The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets. *Annual Review of Sociology*, 34:181–209.

Pew Research Center (2026). More countries had elevated levels of social hostilities involving religion in 2023. Technical report, Pew Research Center. 16th annual report on restrictions on religion; covers 198 countries and territories. Accessed [date].

Reuters (2024). In Modi's Delhi, Indian Muslims segregate to seek security. Reuters. Accessed: 2026-06-07.

Reuters (2025). Anti-minority hate speech in India rose by 74% in 2024, research group says. Reuters. Accessed: 2026-06-07.

Rosen, E., Garboden, P. M. E., and Cossyleon, J. E. (2021). Racial discrimination in housing: How landlords use algorithms and home visits to screen tenants. *American Sociological Review*, 86(5):787–822.

Shah, A. (2024). On measuring Muslim segregation in urban India. *Urban Studies*. Online first.

Singh, G. (2019). Residential segregation by caste over time in urban India. *Environment and Urbanization*, 31(1):179–196.

Susewind, R. (2017). Muslims in Indian cities: Degrees of segregation and the elusive ghetto. *Environment and Planning A*, 49(6):1286–1304.

Tauchmann, H. (2014). Lee (2009) treatment-effect bounds for nonrandom sample selection. *The Stata Journal*, 14(4):884–894.

The Caravan (2019). No Kashmiris please: Housing discrimination in Delhi's Muslim ghettos. The Caravan. Accessed: 2026-06-07.

The Indian Express (2018). The segregated city. The Indian Express. Accessed: 2026-06-07.

Thorat, S., Banerjee, A., Mishra, V. K., and Rizvi, F. (2015). Urban rental housing market: Caste and religion matters in access. *Economic & Political Weekly*, 50(26–27):47–53.

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, 64(5):416–424.

Tobin, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica*, 26(1):24–36.

Turner, M. A. and Ross, S. L. (2003a). Discrimination in metropolitan housing markets: National results from phase i hds 2000. Technical report, Urban Institute.

Turner, M. A. and Ross, S. L. (2003b). Discrimination in metropolitan housing markets: Phase iii – native americans. Technical report, Urban Institute.

TwoCircles.net and Newslaundry (2024). Membership records of registered housing societies in Delhi. Investigative report.

U.S. Department of State (2024). 2023 report on international religious freedom: India. U.S. Department of State. Accessed: 2026-06-07.

Zhao, B., Ondrich, J., and Yinger, J. (2006). Why do real estate brokers continue to discriminate? evidence from the 2000 housing discrimination study. *Journal of Urban Economics*, 59(3):394–419.

A Caste-Specific Results

B Survey Experimental Materials

This appendix presents the experimental materials used in the survey. The materials include the landlord applicant-choice experiments, tenant neighbourhood-choice experiments, and the experimental scenarios presented to government officials.

B.1 Landlord Game 1: Applicant Profiles

Respondents were shown applicant profiles and asked to evaluate the prospective tenants. The profiles varied applicant identity and, in some cases, additional information about tenant reliability.

Table 10: Landlord Game 1: Applicant Profiles (1–4)

Photo	A	B	C
	**Name:** Aarti Bhasin **Occupation:** Doctor **Rent:** ₹22,000 **Family:** Married, two children **Languages:** Hindi and English	**Name:** Atiya Rahman **Occupation:** Doctor **Rent:** ₹22,000 **Family:** Married, two children **Languages:** Hindi and English	**Name:** Atiya Rahman **Occupation:** Doctor **Rent:** ₹22,000 **Reference:** "Ideal tenant" according to previous landlord **Family:** Married, two children **Languages:** Hindi and English
	**Name:** Sanjay Kumar **Occupation:** Store Manager **Monthly Salary:** ₹45,000 **Rent:** ₹13,000 **Family:** Single **Language:** Hindi	**Name:** Sanjay Kumar **Occupation:** Store Manager **Monthly Salary:** ₹45,000 **Rent:** ₹13,000 **Family:** Single **Language:** Hindi	**Name:** Sanjay Kumar **Occupation:** Store Manager **Monthly Salary:** ₹45,000 **Rent:** ₹13,000 **Contract Duration:** Willing to sign a two-year contract **Family:** Single **Language:** Hindi
	**Name:** Shakeel Ahmed **Occupation:** Software Engineer **Monthly Salary:** ₹1,10,000 **Rent:** ₹24,000 **Family:** Single **Languages:** Hindi and English	**Name:** Soham Manjhi **Occupation:** Software Engineer **Monthly Salary:** ₹1,10,000 **Rent:** ₹24,000 **Family:** Single **Languages:** Hindi and English	**Name:** Soham Manjhi **Occupation:** Software Engineer **Monthly Salary:** ₹1,10,000 **Rent:** ₹24,000 **Reference:** Police verification successful **Family:** Single **Languages:** Hindi and English
	**Name:** Arun Kumar **Occupation:** School Teacher **Rent:** ₹18,000 **Family:** Married, two children **Languages:** Proficient in Hindi and English	**Name:** Atul Pandey **Occupation:** School Teacher **Rent:** ₹18,000 **Family:** Married, two children **Languages:** Proficient in Hindi and English	**Name:** Atul Pandey **Occupation:** School Teacher **Rent:** ₹18,000 **Reference:** According to previous landlord, rent payments were very rarely delayed **Family:** Married, two children **Languages:** Proficient in Hindi and English

Table 11: Landlord Game 1: Applicant Profiles (5–8)

Photo	A	B	C
	**Name:** Aakriti Patel **Occupation:** Marketing Executive **Rent:** ₹16,000 **Family:** Divorced, one child, one pet **Languages:** Hindi, Gujarati and English	**Name:** Aakriti Patel **Occupation:** Marketing Executive **Rent:** ₹16,000 **Family:** Divorced, one child, one pet **Languages:** Hindi, Gujarati and English	**Name:** Aakriti Patel **Occupation:** Marketing Executive **Rent:** ₹16,000 **Reference:** Recently moved to the city **Family:** Divorced, one child, one pet **Languages:** Hindi, Gujarati and English
	**Name:** Gulnaz Bano **Occupation:** Police Constable **Monthly Salary:** ₹35,000 **Rent:** ₹11,000 **Family:** Single **Languages:** Hindi and English	**Name:** Pooja Bhatia **Occupation:** Police Constable **Monthly Salary:** ₹35,000 **Rent:** ₹11,000 **Family:** Single **Languages:** Hindi and English	**Name:** Pooja Bhatia **Occupation:** Police Constable **Monthly Salary:** ₹35,000 **Rent:** ₹11,000 **Reference:** Police verification completed **Family:** Single **Languages:** Hindi and English
	**Name:** Akbar Khan **Occupation:** Banker **Monthly Salary:** ₹1,45,000 **Rent:** ₹40,000 **Family:** Married, two children **Languages:** Hindi, Urdu and English	**Name:** Akbar Khan **Occupation:** Banker **Monthly Salary:** ₹1,45,000 **Rent:** ₹40,000 **Family:** Married, two children **Languages:** Hindi, Urdu and English	**Name:** Akbar Khan **Occupation:** Banker **Monthly Salary:** ₹1,45,000 **Rent:** ₹40,000 **Reference:** According to previous landlord, never delayed rent payments **Family:** Married, two children **Languages:** Hindi, Urdu and English
	**Name:** Rakesh Pandey **Occupation:** Researcher **Rent:** ₹10,000 **Family:** Married, one child **Language:** Hindi	**Name:** Arshad Alam **Occupation:** Researcher **Rent:** ₹10,000 **Family:** Married, one child **Language:** Hindi	**Name:** Arshad Alam **Occupation:** Researcher **Rent:** ₹10,000 **Reference:** Wife works as Head of Human Resources at a multinational company **Family:** Married, one child **Language:** Hindi

B.2 Landlord Game 2: Pairwise Tenant Choice

Respondents were presented with pairs of prospective tenants and asked to choose between Tenant A and Tenant B.

Table 12: Landlord Game 2: Tenant Pair 1

Tenant A	**Tenant B**
	
Name: Jayant Maheshwari **Occupation:** Student **Monthly Rent:** ₹12,000 **Family:** Unmarried **Contract Duration:** 12 months	**Name:** Aftab Hussain **Occupation:** Student **Monthly Rent:** ₹14,000 **Family:** Unmarried **Contract Duration:** 12 months

Table 13: Landlord Game 2: Tenant Pair 2

Tenant A	**Tenant B**
	
Name: Sanjay Kumar **Occupation:** Store Manager **Monthly Rent:** ₹25,000 **Family:** Unmarried **Contract Duration:** 18 months	**Name:** Jayant Pandey **Occupation:** Businessman **Monthly Rent:** ₹25,000 **Family:** Unmarried **Contract Duration:** 18 months

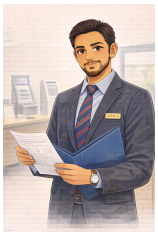
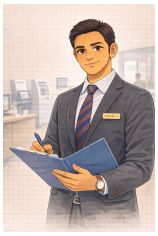
Table 14: Landlord Game 2: Tenant Pair 3

Tenant A	Tenant B
	
Name: Imran Khan **Occupation:** Doctor **Monthly Rent:** ₹22,000 **Family:** Married, one child **Contract Duration:** 24 months	**Name:** Akash Adhikari **Occupation:** Software Engineer **Monthly Rent:** ₹20,000 **Family:** Married, two children **Contract Duration:** 24 months

Table 15: Landlord Game 2: Tenant Pair 4

Tenant A	Tenant B
	
Name: Aarti Bhasin **Occupation:** Doctor **Monthly Rent:** ₹19,000 **Family:** Unmarried **Contract Duration:** 12 months	**Name:** Putul Gupta **Occupation:** Doctor **Monthly Rent:** ₹19,000 **Family:** Married, one child **Contract Duration:** 12 months

Table 16: Landlord Game 2: Tenant Pair 5

Tenant A	**Tenant B**
	
Name: Sameer Sharma **Occupation:** Banker **Monthly Rent:** ₹35,000 **Family:** Unmarried **Tenancy History:** No previous record **Contract Duration:** 12 months	**Name:** Akbar Khan **Occupation:** Banker **Monthly Rent:** ₹35,000 **Family:** Unmarried **Tenancy History:** Good reference from previous landlord **Contract Duration:** 12 months

B.3 Tenant Neighbourhood-Choice Games

Respondents were presented with hypothetical residential areas that varied in rent, commute time, neighbourhood demographics, and available amenities.

Table 17: Tenant Game: Neighbourhood Demographics I

Area	Rent (₹)	Commute Time (minutes)	Neighbour Demographics
A	15,000	20	Mostly Upper-Caste Hindu
B	14,000	35	Mixed – Hindu, Muslim, Christian
C	13,000	25	Mostly Muslim

Table 18: Tenant Game: Neighbourhood Demographics II

Area	Rent (₹)	Commute Time (minutes)	Neighbour Demographics
A	16,000	15	Mostly Muslim
B	15,000	30	Mixed – Hindu, Muslim, Christian
C	14,500	20	Mostly SC

Table 19: Tenant Game: Neighbourhood Amenities I

Area	Rent (₹)	Commute Time (minutes)	Amenities Available in the Area
A	25,000	40	Schools, parks, mosques, metro connectivity
B	28,000	30	Schools, metro connectivity, mosques, temples
C	22,500	35	Metro connectivity, large hospital, parks

Table 20: Tenant Game: Neighbourhood Amenities II

Area	Rent (₹)	Commute Time (minutes)	Amenities Available in the Area
A	15,000	20	Metro connectivity, parks, good schools, temples
B	14,000	35	Markets, lake, religious institutions, stadium
C	13,000	25	Affordable food options, historical monuments

B.4 Government Official Experimental Scenarios

Respondents were presented with hypothetical areas that differed in their demographic composition, characteristics, and public-service needs, and were asked to allocate funds to the areas they felt the funds would be better spent in, behaving like government officials.

Table 21: Government Official Game: Scenario 1

Area A	**Area B**	**Area C**
Demographics 40% General-Caste Hindu 40% Scheduled Caste 10% Muslim 5% Sikh 5% Other	**Demographics** 50% General-Caste Hindu 30% Muslim 10% Sikh 5% Christian 5% Other	**Demographics** 30% Hindu 60% Muslim 5% Sikh 5% Other
Characteristics • Well connected to the metro but extremely congested • Known for good schools and hospitals	**Characteristics** • Property values are increasing due to proximity to a lake • Most lawyers and government employees live here	**Characteristics** • Cheapest rents in the city, attracting students and migrants • Fresh fruits and vegetables are available in local markets
Needs • Public parks • Waste management • Traffic regulation	**Needs** • Flood protection • Road repairs • Drainage system	**Needs** • Road repairs • Banking services • Night-time security

Table 22: Government Official Game: Scenario 2

Area A	**Area B**	**Area C**
Demographics 60% Scheduled Caste 25% General-Caste Hindu 10% Christian 5% Other	**Demographics** 40% Hindu 40% Muslim 10% Sikh 5% Christian 5% Other	**Demographics** 50% Muslim 20% General-Caste Hindu 15% Scheduled Caste 5% Sikh 5% Other
Characteristics • Strong local business community • Large park	**Characteristics** • Newly developed area with modern apartments • Home to many employees in the IT sector	**Characteristics** • Old neighbourhood with several historic buildings • High property prices
Needs • Construction of government schools • Better public transport connectivity	**Needs** • Upgrading of the electricity grid • Police station	**Needs** • Improved air quality • Parking facilities • Cultural preservation fund

Table 23: Government Official Game: Scenario 3

Area A	**Area B**	**Area C**
Demographics Unknown	**Demographics** 55% Scheduled Caste 30% Muslim 10% Sikh 5% Christian	**Demographics** 65% Muslim 25% Hindu 5% Christian 5% Other
Characteristics • High-income area with gated residential complexes • Several bank branches and financial institutions	**Characteristics** • Government schools have poor infrastructure but perform well • Extremely congested area	**Characteristics** • Affordable housing is available, but the crime rate is increasing • Lack of community centres for young people
Needs • Green spaces • Environmental sustainability initiatives	**Needs** • Funding for government schools • Wider roads	**Needs** • Police station • Youth centre • Night-time security

C Additional Broker-Reported Acceptance Categories

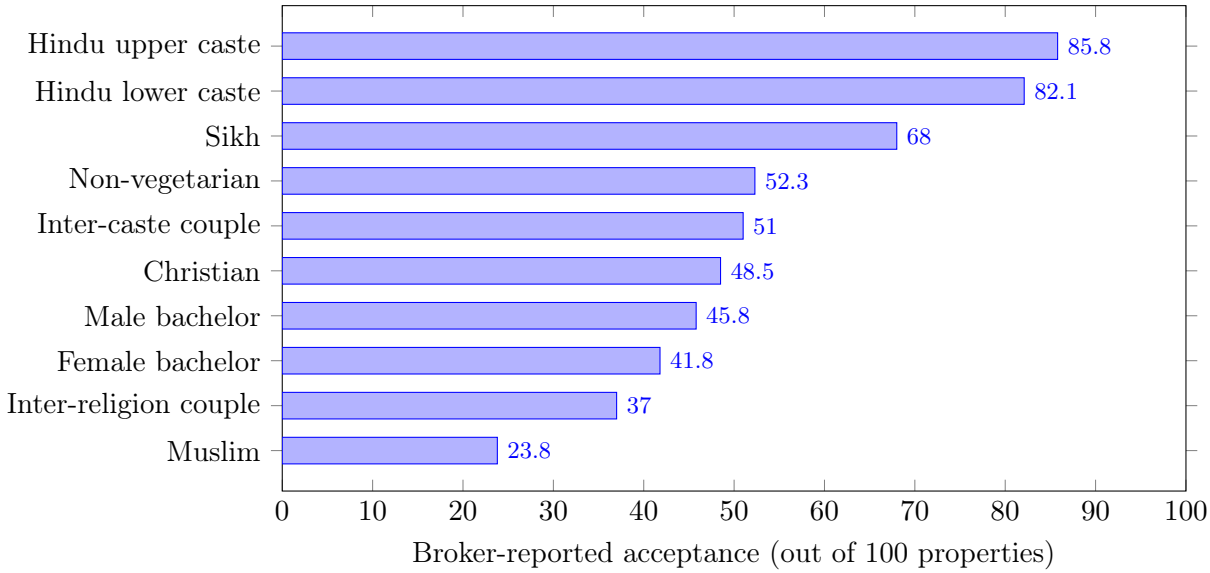


Figure 7: Broker-reported acceptance by applicant type, all categories.

D Lee Bounds on the Price Penalty

Table 24: Lee (2009) trimming bounds on the price penalty

	(1) Rs, f	(2) Rs, m	(3) % rent, f	(4) % rent, m
Identity penalty β_1				
lower	48.337	-40.135	0.220	-0.401
	(46.356)	(102.501)	(0.211)	(1.025)
upper	115.926***	341.811***	0.527***	3.418***
	(30.667)	(46.232)	(0.139)	(0.462)
Observations	3909	3909	3909	3909

Standard errors in parentheses

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