Discrimination, Preferences, and Residential Sorting: Evidence from Delhi's Rental Housing Market

Hamza Syed

September 14, 2026

Preliminary working paper. Please do not circulate.

Abstract

Residential segregation is routinely read as revealed preference. In a rental market, that reading requires an assumption which hasn't been tested: that the households who live apart could have chosen otherwise. Location reflects both what households want and what landlords, brokers and resident associations will permit, and these have never been measured on the same population. Using randomized neighbourhood-choice and vignette experiments fielded on 3,253 residents and 615 rental brokers in Delhi, together with direct elicitation from the intermediaries who hold the rental inventory, I measure both. Own-group residential preferences are substantial for every group. Supply-side exclusion of Muslim applicants runs the other way and is severe: brokers report that a Muslim applicant can secure housing in 24 of 100 properties in their service area against 86 for an upper-caste Hindu, and resident associations are reported to bar Muslim tenants in 26 percent of areas against under 1 percent for Hindus. Most Muslim households would choose a mixed or majority-Hindu neighbourhood; they can reach about a quarter of the market. Exclusion operates almost entirely on quantity rather than price: roughly 97 percent of the identity-induced shift in response is refusal rather than a markup. Observed sorting is therefore a constrained equilibrium rather than free choice, and the binding constraint sits in countable, regulable institutions.

1 Introduction

Two Muslim households and two Hindu households live in different parts of the same city. The standard reading of that fact is that they preferred it: people sort, and where they end up reveals what they wanted. In a rental market the reading requires something more. It requires that each household could have taken the other's place, and chose not to. If a landlord, a broker,

[1]I thank the IGC and STICERD for funding the full and pilot surveys respectively. I am grateful to my supervisors Tim Besley, Maitreesh Ghatak and Michael Callen, and to Sam Asher, Gharad Bryan, Catherine Szermann, Gilat Levy and Michael Greenstone for comments throughout the project. Sanjana Nahar, Ayush Dwivedi and Zoheen Shehzad provided excellent research assistance. The study received ethics clearance from the LSE Research Ethics Committee (Ref: 193409, 24 April 2023, amendment approved 24 May 2025).

or a residents' association would have refused, then the map is not a preference ordering. It is a preference ordering intersected with a permission set, and the two cannot be recovered from the map alone.

This paper measures both on the same population, and finds that they point in opposite directions.

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; Cutler and Glaeser, 1997), which makes the housing market one of the principal mechanisms through which economic opportunity is allocated within a city. Audit and correspondence methods have established that discrimination in housing access is real and persistent (Bertrand and Mullainathan, 2004; Turner et al., 2002; Flage, 2018; Neumark, 2018), and recent work shows that it operates not only through outright refusal but through steering (Christensen and Timmins, 2022, 2023). That literature measures access but is silent on preferences. Equilibrium sorting models recover preferences but assume households choose freely across neighbourhoods (Bayer et al., 2007; Schelling, 1971; Card et al., 2008). Neither measures both. The field therefore lacks an empirical handle on the counterfactual it most often invokes. Documenting Muslim and Scheduled Caste segregation across 1.5 million Indian neighbourhoods, Asher et al. (2026) state explicitly that their analysis cannot disentangle discrimination from homophily, and identify that separation as the natural next step.

Why Delhi. Two features of the setting make the separation feasible. First, religion is observable without its usual companions. 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 and isolating it has required elaborate design (Adida et al., 2010; Di Stasio et al., 2021). Hindus and Muslims in Delhi are, for the most part, co-national, co-ethnic, co-linguistic and long-settled in the same city. A name signals religion without signalling foreign birth, migration status or race. Scheduled Caste applicants supply a benchmark unavailable elsewhere: a stigmatised minority who are neither religiously distinct nor distinguishable by any of the markers that separate Muslims from Hindus, so comparing the two penalties separates religion-specific exclusion from generic minority disadvantage.

Second, the discretion that shapes tenant access is concentrated in institutions one can

interview. Brokers hold and generate the listing inventory for a compact service area and screen applicants before any landlord contact; resident welfare associations administer tenant-eligibility norms. Rather than infer exclusion from applicant-side callback rates, I ask the gatekeepers directly how much of their local stock each identity can reach.

What the paper finds. I establish two results and devote the rest of the paper to testing them and to asking what produces them.

Exclusion operates on quantity, not price. The vignette instrument records not just whether a respondent accepts an applicant but how: at the offered rent, only at a higher rent, or not at all. Changing an otherwise identical profile from a Hindu to a Muslim name displaces 19 to 22 percentage points of probability mass out of acceptance at the offered rent. Essentially all of it lands on refusal; under one point lands on acceptance at a markup. The rationing share i.e., the fraction of the identity-induced shift operating on quantity rather than price, is 0.97 and is flat across tenants, landlords and brokers. This is not what a market clearing on price would look like, and it is a measured quantity rather than an assumption.

Desired and feasible sorting diverge, and they diverge for Muslim households specifically. Facing hypothetical menus in which every option is stipulated to be available, Muslim respondents choose a mixed or majority-Hindu neighbourhood 57 percent of the time. Brokers report that a Muslim applicant can secure housing in about 24 of 100 properties in their service area, against 86 for an upper-caste Hindu and 82 for a lower-caste Hindu; resident associations are reported to bar Muslim tenants in 26 percent of areas against under 1 percent for Hindus. Figure 1 puts the two side by side. For upper-caste and Scheduled Caste Hindus, access comfortably exceeds stated willingness to live outside the own-group area. For Muslims it does not, and the gap is large.

The two results are connected. A market in which identity operates almost entirely on the quantity margin is one in which being unacceptable to a gatekeeper is not something an applicant can pay to overcome, so the wedge in Figure 1 cannot be closed by willingness to pay. It is also a market in which the competitive erosion that Becker (1957) predicts is muted: a landlord who declines an applicant and lets to someone else at the same rent forgoes nothing observable. Section 6 develops a simple model of selection from a pool of applicants that rationalises both results and delivers a third, testable implication, that supplying information about tenant reliability should not move the identity penalty, because acceptance turns on rank within the

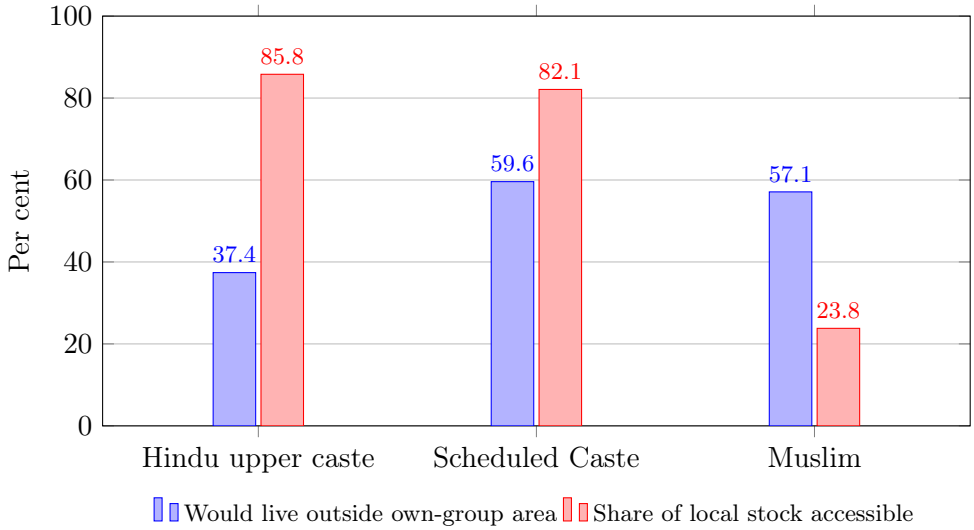


Figure 1: Desired and feasible sorting. The first bar is the share of respondents of each group who choose a neighbourhood that is not their own group's in the neighbourhood-choice experiment, in which every option is stipulated to be affordable and available (Table 5). The second is broker-reported acceptance out of 100 representative properties in the broker's service area for an applicant of that group (Figure 3). Scheduled Caste respondents are matched to the broker measure for lower-caste Hindu applicants. For Hindu groups, access exceeds stated willingness to move outside the own-group area; for Muslim households it falls well short of it.

pool rather than on the level of expected surplus. It does not move it, while an income signal embedded in the same experimental arm moves quoted rents sharply, which rules out inattention as the explanation for the null.

A further pattern runs through the evidence and is harder to accommodate. The penalty rises monotonically along the chain of intermediation: it is about -0.07 for respondents who are tenants, -0.24 to -0.30 for landlords, and -0.41 to -0.48 for brokers. Brokers bear none of the tenant's default risk and earn a commission on a completed match, so their exclusion is difficult to read as risk management or as a taste they personally hold. At the same time, list experiments that shield respondents from social-desirability pressure recover almost no *admitted* discrimination, even as the same indirect method uncovers widespread majoritarian sentiment on attitudinal items. The behaviour is severe; what people will own is not. The reading I find most consistent with the evidence, and which I state as an interpretation rather than a finding, is that exclusion here is enacted as an expectation about what others will accept rather than chosen as a personal preference.

Contributions. *Joint measurement.* I field randomized neighbourhood-choice experiments and randomized supply-side vignettes on the same respondents and elicit access constraints

directly from gatekeepers, so preferences and access are identified from separate, purpose-built sources of variation rather than from the residual of one another.

Gatekeeper elicitation. Asking the 615 brokers who hold the local rental inventory how much of it each identity can reach yields spatially resolved access measures that a correspondence design cannot produce. It also reaches a margin correspondence designs cannot: properties an applicant never learns exist, because the broker withholds them before any listing is seen.

Measuring the margin on which discrimination operates. The rationing share (Λ) is unit-free and comparable across settings, and it converts a claim about market structure into a number. The literature on discrimination in labour markets has concentrated on the price (wage) margin because wages are observed and hiring rejections are not; housing markets with gatekeeper elicitation reverse that, and this paper reports both.

Policy. India has no general statutory prohibition on discrimination in private housing, and the leading authority affirms the autonomy of housing societies to restrict membership by religion. The policy question is therefore not how to enforce an existing right but how to create a regulatory hook where none exists. 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. Extending it is a small legal step onto the actor the evidence identifies as the most exclusionary and the most regulable.

Road map. Section 2 describes the setting. Section 3 describes the data and design. Section 4 and Section 5 present the two main results, Section 6 interprets them, Section 7 reports further evidence on mechanisms, Section 8 the robustness analysis, Section 9 the cost of the access constraint, and Section 10 the enforcement implications.

2 Setting

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 (Moretti, 2011; Hsieh and Moretti, 2019; Glaeser and Luttmer, 2003). 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 disparities rather than merely reflect them.

Residential segregation along religious and caste lines is a prominent feature of urban India, and is documented at national scale by Asher et al. (2026) and across metropolitan areas by Bharathi et al. (2021) and Singh (2019), who finds caste segregation persisting or increasing in a majority of Indian cities between 2001 and 2011. Work specifically on Muslim segregation shows it to be widespread but highly variable in extent and form across and within cities (Susewind, 2017; Shah, 2024). Galonnier (2015) argues that Muslim residential clustering can simultaneously reflect involuntary segregation from discrimination and insecurity, voluntary clustering on cultural grounds, and exclusionary practices by dominant groups, which is precisely the identification problem this paper addresses; Gayer and Jaffrelot (2012) document the historical trajectory. Reporting suggests movement of Muslim households into Muslim-majority neighbourhoods in Delhi has increased following episodes of communal violence (Reuters, 2024).

Concern about anti-Muslim hostility in India has grown over the past decade (Reuters, 2025; Human Rights Watch, 2024; U.S. Department of State, 2024; Pew Research Center, 2026). I record this 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 i.e., selection from a pool of applicants, ex-ante screening under weak contract enforcement, and gatekeeping by intermediaries, are structural features of the market rather than functions of the political moment.

2.2 Religion as an axis of economic exclusion

The economics of discrimination has developed largely around race in the United States and ethnicity in Europe (Becker, 1957; Arrow, 1973; Phelps, 1972; Coate and Loury, 1993). Religion has been studied less, and in Western settings cannot easily be studied cleanly, because religious minority status is bundled with immigration, national origin and race. Adida et al. (2010) hold Senegalese origin fixed and vary only religious signalling in French applications, finding the Muslim applicant roughly 2.5 times less likely to be called for interview. Di Stasio et al. (2021) generalise the design across five European labour markets and separate an anti-Muslim penalty from a "Muslim by default" origin stigma. In housing, Ahmed and Hammarstedt (2008) document substantially fewer callbacks to Arab/Muslim-named applicants in Sweden, and Flage (2018), pooling 25 OECD correspondence studies from 2006 to 2017 comprising more than 300 estimates and over 110,000 inquiries, finds majority applicants almost twice as likely to be

selected as minority applicants generally and more than twice as likely relative to Arab or Muslim applicants specifically.

Evidence from India is thinner but pointed. Banerjee et al. (2009) find caste and religion affect callbacks in Indian white-collar labour markets; Thorat and Attewell (2007) and Siddique (2011) report analogous results. In housing, Datta and Pathania (2016) run a correspondence study in Delhi and find landlords are about 13% less responsive to Muslim than to comparable upper-caste Hindu applicants; Thorat et al. (2015) document caste and religion barriers in urban rental access. Gbadegesin and Ojo (2013) report ethnic bias in tenant selection in Ibadan. Reviewing the literature more broadly, Ifeanacho et al. (2023) find it concentrated almost entirely in the United States, Germany and Canada (also see Neumark (2018) and Flage (2018)).

Two findings from this literature bear directly on the design here, and I return to both. First, the *information* margin: Ahmed et al. (2010) and Auspurg et al. (2017) find that supplying landlords with additional applicant information, or signalling high occupational status, reduces discrimination in Swedish and German rental markets, consistent with a statistical mechanism operating there. My vignette design includes a directly comparable reliability-signal arm, and it is informative that the signal moves acceptance in Germany but not in Delhi. Second, the *tightness* margin: Christensen and Timmins (2023) find discriminatory constraints 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.

2.3 Delhi's rental market: prices, contracts, and enforcement

Delhi is the national capital and draws people from across the country. Rental housing accounts for roughly 28 to 30 percent of urban residence in India (Gandhi et al., 2021) and Delhi accounts for roughly 5 percent of the national rental stock (Knight Frank, 2019). Three features of this market shape what follows.

Rents do not do the allocating. Rental housing fell from 53 percent of the occupied urban residential stock in 1971 to 28 percent in 2011 (Gandhi et al., 2021), while the 2011 Census records 11.1 million vacant urban units, 12.4 percent of the stock, against 1.8 million in 1971. Gandhi et al. (2021) attribute the coexistence of shortage and vacancy 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. The Delhi government's Economic Survey for 2020–21 records 6.75 million people in low-income settlements, 695 slum settlements, 1,797 unauthorised colonies and 362 urban villages, and an additional requirement of 2.4 million housing units.

I flag what this evidence does and does not establish, because the paper has been read as assuming more than it needs. It establishes that posted rents in this market are not a free instrument: they are constrained by statute, adjusted infrequently, and set property-by-property rather than applicant-by-applicant. It does *not* establish that every vacancy carries a queue, and a large vacant stock is on its face evidence in the other direction. The paper's empirical claim is about the *margin on which identity operates*, which I measure. Excess demand is one micro-foundation for what I find, and Section 6 works it through, but the results in Section 4 do not depend on it.

Contracts are thin and enforcement is slow. 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 for recovering 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, being under twelve months and unregistered, falls outside rent control. Roughly 19 percent of households in formal urban rental housing hold a written contract, and roughly 14 percent in slums (Gandhi et al., 2021). Contract enforcement in India took an average of 1,445 days as of 2019 (World Bank, 2020). 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.

Screening therefore happens at the door. A landlord who cannot rely on ex post remedies must resolve his exposure ex ante. 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. This generates a testable prediction - because identity is being used where verification is absent, supplying verification directly should reduce the identity penalty, and Section 4 shows that it does not.

2.4 The gatekeepers and the legal status of exclusion

How a rental transaction works. Brokers, community networks and online portals are the main search avenues in Delhi, with over 90 percent of tenants in the survey reporting having used one or more.[1] Brokers pre-screen applicants before taking a shortlist to the landlord; identity, occupation, family composition and dietary preferences are established pre-viewing, in the first meeting with the broker. The broker then arranges the viewing, conducts the first round of rent negotiation, and drafts the agreement. He typically charges half a month's rent from each side on completion and services an area of roughly 1.5 to 2 kilometres around his shop. A fuller narrative of the transaction is in Appendix B.

The feature that matters analytically 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 screening outcomes for properties the applicant already knows about; it cannot observe the inventory withheld. Gatekeeper elicitation can.

Who exercises discretion. *Landlords* hold the residual right to select and bear the consequences of selection. *Brokers* hold and generate the listing inventory, pre-screen, vouch for the tenant, and draft the agreement but do not own the property, set the rent, sign the contract, or bear the tenant's default risk. That last point carries weight: the broker's income is a commission on a completed match and his continuing business depends on his standing with local landlords, so his screening question is not "will this tenant perform?" but "will this tenant be accepted?". *Resident welfare associations* administer eligibility norms that are often unwritten but enforced in practice, and that operate on the landlord rather than the tenant. Reporting based on the membership records of all 1,960 registered housing societies in Delhi found 68 percent with no Muslim member, with Muslim membership across societies barely exceeding 3 percent (TwoCircles.net and Newslaundry, 2024); I use this as an external benchmark for the elicited RWA measure in Section 4.6.

Both brokers and RWAs are formal entities and are *countable.* Brokers operate from fixed premises and serve compact areas, so the population of relevant gatekeepers in any neighbourhood is small and observable. This is what makes the discrimination measured here a tractable

[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.

target for policy rather than an immutable feature of preferences, and it parallels the argument of Kline et al. (2022) that discrimination concentrated in a small number of identifiable organisations is far more amenable to enforcement than discrimination diffused across many.

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. 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 a single religious community, holding that the governing Cooperative Societies Act did not prohibit such a restriction, that it was not contrary to public policy, and that the associational autonomy protected by Article 19(1)(c) was engaged. 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 10 argues that the broker, not the society, is that point.

3 Data and design

The analysis combines a survey of resident households, a survey of rental brokers, neighbourhood-level spatial data, and randomized experiments embedded in both surveys. The design observes, for the same neighbourhoods, both the preferences households reveal over locations and the access constraints they face.

3.1 Sampling

Residents. 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, in particular areas with mosques versus those with temples, alongside mixed 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, which limits enumerator discretion in selecting respondents. All resident and broker locations were geocoded.

Brokers. The broker sampling frame is a list of rental brokers scraped from Google Maps across the city. Brokers were drawn at random from this frame, controlling for area type as defined by the same temple- and mosque-concentration classification used for the resident strata, so that broker coverage tracks the composition strata rather than commercial density alone. Figure 2 shows the locations of both surveys.

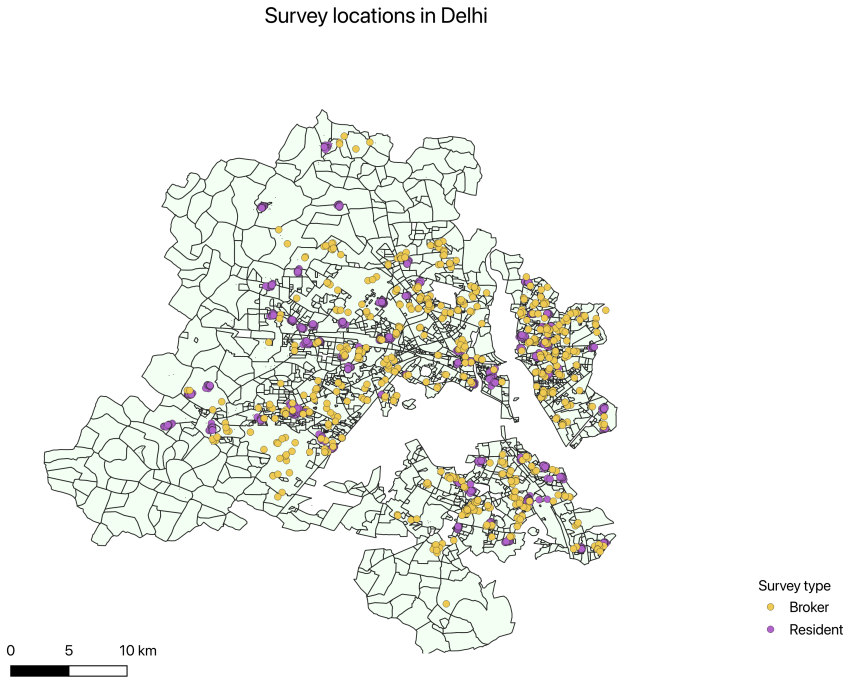


Figure 2: Survey areas. Points show the approximate locations of surveyed resident households and rental brokers across Delhi.

3.2 Fieldwork, incentives, framing, and ethics

The study received ethics clearance from the LSE Research Ethics Committee [2]. Participation was voluntary and informed consent was obtained before each interview.

Respondents received a household gift set - typically a set of mugs or a comparable household item worth approximately ₹200–250 for residents and approximately ₹500 for brokers. The response rate was approximately 98 percent In some richer areas, the survey team was not allowed entry, so those areas were replaced by a backup area from the sampled list. Non-response is therefore unlikely to be a material source of bias, which matters here because differential refusal on a sensitive topic would be a first-order threat.

The surveys were introduced to respondents as a study of rental housing decisions in Delhi,

[2]Ethics review approval ref: 193409, 24 April 2023, amendment approved 24 May 2025

not as a study of discrimination, and enumerators were trained to maintain that framing throughout. The sensitive modules are embedded among modules on search, satisfaction, amenities and tenancy experience. This reduces the risk that respondents infer the object of measurement and adjust their answers, and it is one reason to treat the direct measures as informative rather than as pure expressions of what respondents think the researcher wants to hear.

3.3 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 below. Religion and caste are collected directly and used to construct the three identity groups that recur throughout: Hindu upper caste, Muslim, and Scheduled Caste. Table 1 compares the sample composition with the 2011 Census.

Table 1: Religious composition of survey respondents

Religion	Frequency	Per cent	Delhi (2011 Census)
Buddhist	4	0.12	0.12
Christian	8	0.25	0.96
Hindu	2,708	83.25	80.21
Muslim	496	15.25	12.78
Jain	9	0.28	1.39
Sikh	26	0.80	4.43
Other / not stated	2	0.06	0.10
Total	3,253	100.00	100.00

Notes. One observation per surveyed resident household. 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). The sample slightly over-represents Muslim households relative to the city, by construction: the sampling design stratifies on local religious composition in order to obtain adequate coverage of minority-concentrated neighbourhoods (Section 3.1). Sikh households are under-represented but are not material to the analysis. Columns may not sum to 100 due to rounding.

3.4 Broker survey

The broker survey covers 615 rental brokers and is the paper's distinctive instrument, so it is worth setting out precisely what it delivers and why a broker report is the right measurement

tool here rather than a second-best.

What brokers are asked. For each broker the survey records reported landlord preferences in his area; his own expectations about tenant acceptability; the rules imposed by local resident welfare associations; identity-specific match probabilities; his experience with past tenancies; and his responses to the same vignette module that residents receive. Three of these do distinct work and should not be conflated. The *identity-specific match probabilities* - out of 100 representative properties in your service area, how many would accept a tenant of each type - give the access measure $p_{g,n}$. The *RWA rules* give an institution-level measure of rule-based exclusion. The *vignette responses* give a broker-specific identity penalty directly comparable to the landlord's and the tenant's.

Why elicitation rather than audit. First, coverage. A correspondence design must observe a listing before it can send an inquiry, so it cannot see properties withheld at the pre-screen. The broker report is a share of the local stock and therefore includes them. Second, spatial resolution. A broker serves a compact area, so his report is a neighbourhood-level measure, which is what the sorting analysis in Section 9 requires and what an applicant-side callback rate cannot supply. Third, incidence. Applicant-side measures capture the composition of refusal and non-response without locating it; asking the intermediary locates it in the intermediary.

Why brokers answer. The obvious objection is that a broker has every reason to understate his own exclusionary conduct. Three things cut against it. The survey is framed as a study of rental housing decisions (Section 3.2), and brokers describe tenant screening as a service they provide to landlords rather than as discrimination they practise, which is also the reading the list experiments in Section 7.1 support. The reported figures are, if anything, implausibly candid: brokers report Muslim access at under a quarter of that for upper-caste Hindus. And the elicited RWA measure can be benchmarked externally against the housing-society membership records of TwoCircles.net and Newslaundry (2024), which is done in Section 4.6. Any residual bias from social desirability would attenuate the measured gaps, so the estimates are conservative in the direction that matters.

3.5 Spatial data

I assemble neighbourhood-level data on local amenities, including schools, hospitals, parks, religious institutions and prevailing rents.[3] These serve two purposes: they construct the amenity measures used in Section 7.2, and the location of religious institutions is used to classify neighbourhoods into area types, which underpins both the sampling strata and the area-type definitions used throughout.

3.6 Experimental modules

Both surveys share an instrument organised into three blocks. An *experiments* block contains the randomized modules: a vignette experiment on supply-side gatekeeping (Section 4), a neighbourhood-choice experiment on residential preference (Section 5), a pairwise tenant-choice module (**??**), and list experiments (Section 7.1). Full materials are in Appendix C. A *market-beliefs* block records landlord preferences, broker expectations, perceived matching probabilities by identity, and tenancy experiences. A *background* block records demographics, neighbourhood composition, housing satisfaction and exposure to neighbourhood disamenities. The background block is at the end so as not to prime the respondents to think about their identity.

Vignette experiment. Respondents are asked to imagine they are the landlord of a house they have advertised, and evaluate eight hypothetical applicant profiles. 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). The design is a factorial survey experiment in the sense of Auspurg and Hinz (2015). For each profile the respondent may (1) accept at the offered rent, (2) accept only at a higher rent, specifying the difference, (3) wait x months and accept if no other application arrives, or (4) reject. Comprehension-check accuracy exceeded 95 percent. Because respondents evaluate multiple vignettes in sequence, order effects are a known concern (Auspurg et al., 2017); profile order was held fixed across arms, so any order effect is differenced out by the arm contrast.

Three features of the fielded instrument matter for what follows. *The photograph is held constant across arms within a profile.* Photographs were AI-generated with prompt text instructing that the depicted person carry no religious markers, and were field-tested with survey

[3] Amenity data were scraped from Google Maps by keyword; prevailing rents were scraped from MagicBricks.com.

staff to confirm that they did not read as identifiably Hindu or Muslim. Identity therefore varies through the name alone. *Arm C is not a constant signal*: it ranges from a landlord reference to police verification to a lease-length commitment to spousal employment, so I report its effect profile by profile and distinguish cues that are pure reliability signals from those that also reveal ability to pay. *Three profiles carry identical names in arms A and B* and are therefore placebo transitions, reported in Section 8.1.

Neighbourhood-choice experiment. Respondents choose among three hypothetical neighbourhoods varying in rent, commuting time, and either group composition or amenity bundle, in the spirit of conjoint analysis (Hainmueller et al., 2014). They are told they can afford the rent in each option and should choose the one that suits them best, so the exercise fixes availability by instruction and isolates preference from the access constraints that bind in the field. The game was repeated four times; comprehension accuracy exceeded 98 percent.

List experiments. To measure discrimination free of social-desirability pressure, both surveys include list (item-count) experiments (Imai, 2011; Blair and Imai, 2012). Respondents report only the *number* of statements on a list that apply to them; a 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 stating it directly.

4 Result 1: exclusion operates on quantity, not price

A landlord who does not want a particular tenant has two instruments. He can decline him, or he can quote him a higher rent. Which he uses is not a matter of taste alone but of what the market permits, and it is measurable, because the vignette instrument records the response on an ordered scale rather than as a binary. This section reports that measurement. It is the paper's first main result and, I will argue in Section 6, the one that organises everything else.

Table 2 lists the objects estimated in this section and the words I use for them throughout.

Table 2: Objects estimated in Section 4, and where

Symbol	In words	Estimated by	Reported in
β_1	Identity penalty	Arm A $\rightarrow$ arms B/C contrast	Tables 3 and 9
β_2	Information effect	Arm B $\rightarrow$ arm C contrast	Tables 3 and 9
Δ_1	Shift out of acceptance at offered rent	Ordered probit marginal effect	Table 3
Δ_2	Shift onto the price margin	Ordered probit marginal effect	Table 3
Δ_3	Shift onto the quantity margin (refusal)	Ordered probit marginal effect	Table 3
Λ	Rationing share, $\Delta_3/(-\Delta_1)$	Ratio of the above	Tables 3 and 4
m_i	Quoted markup over the offered rent	Two-part model, conditional on not refusing	Table 3
$p_{g,n}$	Access: share of local stock reachable by group g	Broker elicitation	Figure 3

Notes. All objects in this table are defined formally in the text of Section 4. Section 6 maps them onto the model.

4.1 The reluctance ladder

The four recorded responses are naturally ordered in the reluctance they express: accept at the offered rent; accept only at a higher rent; wait; reject. Almost no respondent chooses to wait, so I collapse *wait* and *reject* into a single category, *refuse*, and work with three ordered outcomes:

$$y_i = \begin{cases} 1 & \text{accept at the offered rent,} \\ 2 & \text{accept only at a higher rent} \quad (\text{the } \textit{price} \text{ margin}), \\ 3 & \text{refuse} \quad (\text{the } \textit{quantity} \text{ margin}). \end{cases} \tag{1}$$

I model y_i as generated by a latent reluctance index

$$\underbrace{y_i^*}_{\text{reluctance}} = \underbrace{\beta_1}_{\substack{\text{identity} \\ \text{penalty}}} \mathbf{1}\{B_i \text{ or } C_i\} + \underbrace{\beta_2}_{\substack{\text{information} \\ \text{effect}}} \mathbf{1}\{C_i\} + \mathbf{W}_i'\boldsymbol{\delta} + \epsilon_i, \qquad \epsilon_i \sim N(0,1), \tag{2}$$

observed through thresholds $c_0 = -\infty < c_1 < c_2 < c_3 = +\infty$, with $y_i = k$ if and only if $c_{k-1} < y_i^* \leq c_k$, and estimated by ordered probit with standard errors clustered on the respondent. The estimate $\hat{\beta}_1$ is the cleanest single summary of the identity effect: it uses every response category, and it conditions on nothing, so it is not subject to the selection problem that afflicts the price margin below.

4.2 Where the identity shift goes

Because equation (2) places all responses on one latent scale, it identifies not only how much reluctance identity generates but *where* that reluctance sends respondents. Let

$$\Delta_k \equiv \frac{\partial \Pr(y_i = k)}{\partial \mathbf{1}\{B_i \text{ or } C_i\}}, \qquad k = 1, 2, 3, \tag{3}$$

be the marginal effect of the identity change on each response probability, evaluated at sample means. Probabilities sum to one, so $\sum_k \Delta_k = 0$: the identity change moves mass out of acceptance at the offered rent ($\Delta_1 < 0$) and redistributes it either to the price margin (Δ_2) or to the quantity margin (Δ_3). The split is summarised by the *rationing share*

$$\underbrace{\Lambda}_{\substack{\text{rationing}\\ \text{share}}} = \frac{\Delta_3}{\Delta_2 + \Delta_3} = \frac{\Delta_3}{-\Delta_1} \in [0, 1]. \tag{4}$$

Rationing share $\Lambda \to 0$ describes a market that prices the disamenity and houses every applicant; $\Lambda \to 1$ describes one that refuses. Λ is unit-free, which is what makes the comparison possible at all: a refusal effect in percentage points and a markup denominated in rupees cannot otherwise be set against each other. I report $\hat{\Lambda}$ with standard errors from a nonparametric bootstrap clustered on the respondent.

4.3 The price margin, and why it is a lower bound

Conditional on not refusing, I measure the price margin directly. Let $m_i \geq 0$ be the quoted markup over the offered rent, observed only for respondents with $y_i \leq 2$. 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,

$$\Pr(m_i > 0 \mid y_i \leq 2) = \Phi\left(\beta_1^P \mathbf{1}\{B_i \text{ or } C_i\} + \beta_2^P \mathbf{1}\{C_i\} + \mathbf{W}_i'\boldsymbol{\delta}^P\right), \tag{5}$$

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

$$m_i^* = \beta_1^T \mathbf{1}\{B_i \text{ or } C_i\} + \beta_2^T \mathbf{1}\{C_i\} + \mathbf{W}_i'\boldsymbol{\delta}^T + \nu_i. \tag{6}$$

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).

Why the estimated markup understates the true price penalty. The conditioning in equations (5) and (6) is not innocuous, and the argument is worth stating in words before the algebra because it recurs throughout the paper.

Only a respondent who does not refuse can quote a price. Changing the applicant's identity makes refusal more likely, and it does so most among exactly those respondents who object to

the applicant most strongly. The group that survives into the markup regression is therefore not a random subset of the sample: it is the subset least inclined to discriminate. Whatever markup they quote understates what the full sample would charge if refusal were unavailable to them. The estimated price penalty is a *floor*, not an estimate.

Formally, write $s_i = \mathbf{1}\{y_i \leq 2\}$ for selection into the markup sample. Identity reduces $\Pr(s_i = 1)$, and reduces it most among respondents with large identity effect 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. This is not a defect of the design so much as a direct consequence of what the paper argues: where refusal is available, the price margin is mechanically hard to observe, because the gatekeepers who would price most aggressively never name a price.

Two things follow for reporting. I lead with $\hat{\beta}_1$ from the ladder equation (2), which conditions on nothing. And I report Lee (2009) trimming bounds on the price effect, estimated using Tauchmann (2014). These 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 the reverse-transition estimate in Section 8.2 supports empirically. They are in the spirit of the partial-identification approach of Manski (1990). 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 the estimator admits no covariates, so it presumes random assignment, which the arm randomisation delivers by construction.

4.4 Results

Table 3 reports the ladder, the decomposition and the price margin for the two cases where there is a shift from a Hindu profile to a Muslim profile. The identity shift raises reluctance sharply (+0.71 and +0.65, both $p < 0.01$). Of the 19 to 22 percentage points of probability mass displaced out of acceptance at the offered rent, 18.8 and 21.5 points land on refusal and 0.4 and 0.8 points on acceptance at a higher rent. The rationing share is $\Lambda = 0.978$ and 0.965.

Table 4 shows that this is not a feature of one respondent type. The identity penalty rises steeply along the chain of intermediation - $\hat{\beta}_1$ of 0.32 for respondents who are actually tenants, 1.00 for respondents who are actually landlords, 1.37 for respondents who are actually brokers, but rationaing share Λ does not move: it is between 0.957 and 0.965 for every type. Respondents differ enormously in how much they exclude and not at all in *how* they exclude.

Whatever is producing the quantity margin is a feature of the market rather than of who is doing the screening.

Two qualifications discipline the 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. Λ near one therefore reflects both that respondents seldom price at all and that the identity shift goes to refusal. The two facts point the same way but 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: 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 ($\hat{\beta}_1^P = 0.386$, $p = 0.031$, a marginal effect of [insert marginal effect] percentage points on a base rate of 1.2 percent), with a Tobit coefficient of about ₹8,000. But by the argument of Section 4.3 this sample is selected toward the least discriminating part of the distribution. Lee (2009) bounds (Table 23) place the effect between ₹48 and ₹116 for the married woman and between −₹40 and ₹342 for the married man, with confidence intervals including zero in both cases. **After accounting for selection, the sign of the price penalty is not identified.**

4.5 Information moves prices but not the identity penalty

The reliability signal in the information arm C has no effect on reluctance: $\hat{\beta}_2 = -0.005$ ($p = 0.93$) for the married woman and 0.065 ($p = 0.18$) for the married man, and the same null appears in every column of Table 9. The obvious alternative reading is that respondents skim the profiles and ignore whatever is added to them. The two profiles rule that out, because their information arm-C cues differ in kind.

For the married woman the cue is a pure reliability signal - a previous landlord describing her as an ideal tenant, and it moves nothing, on the ladder or on the markup probit (0.027, $p = 0.85$). For the married man the cue instead reveals that his wife is head of human resources at a multinational company. That is an income signal rather than a reliability signal, and against a ₹10,000 rent a large one. It leaves overall reluctance unchanged 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 read the added information and do update on it when it is economically relevant: told the applicant's household income is high, they charge more. They do not update on reliability information when the applicant is Muslim. This is the cleanest evidence in

Table 3: Identity penalty, margin decomposition, and the price margin

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	3,868	3,868

Panel B. Where the identity shift goes

		(1) Married woman	(2) Married man
Δ_1	Out of acceptance at offered rent	-0.192^{***}	-0.223^{***}
		(0.016)	(0.017)
Δ_2	Onto the price margin	0.004^{***}	0.008^{***}
		(0.001)	(0.001)
Δ_3	Onto the quantity margin (refusal)	0.188^{***}	0.215^{***}
		(0.015)	(0.016)
Λ	Rationing share	0.978^{***}	0.965^{***}
		(0.005)	(0.005)
Observations		[N]	[N]

Panel C. Price margin, conditional on not refusing

	(1) Pr(markup)	(2) Markup, ₹	(3) Markup, %	(4) Pr(markup)	(5) Markup, ₹	(6) Markup, %
	Married woman			Married man		
Identity penalty β_1	0.386^{**}	8080.78^{**}	36.73^{**}	0.063	457.08	4.57
	(0.179)	(3996.74)	(18.17)	(0.114)	(1809.04)	(18.09)
Information effect β_2	0.027	1061.17	4.82	0.250^{**}	4394.54^{**}	43.95^{**}
	(0.148)	(3317.85)	(15.08)	(0.110)	(1796.22)	(17.96)
Observations	2,967	2,967	2,967	2,528	2,528	2,528
Respondents quoting a markup	35	35	35	115	115	115

Notes. The unit of observation is a respondent × vignette profile. Each respondent evaluates eight applicant profiles as a hypothetical landlord; the two columns correspond to the two profiles for which a clean Hindu → Muslim name change is available, a married woman with two children (a doctor, offered rent ₹22,000) and a married man with one child (a researcher, offered rent ₹10,000). Arm A shows the Hindu-named applicant; arm B shows an otherwise identical profile with a Muslim name; arm C shows arm B plus a tenant-reliability signal. β_1 is therefore the effect of changing the applicant's name from Hindu to Muslim holding every other attribute and the photograph fixed, and β_2 the incremental effect of the reliability signal.

Panel A estimates ladder equation (2) by ordered probit on the three ordered outcomes of equation (1): accept at offered rent < accept at a higher rent < refuse, where *refuse* pools the "wait" and "reject" responses, which are combined because fewer than [x] percent of responses are "wait". Higher values denote greater reluctance, so a positive β_1 is a penalty. Panel B reports the marginal effect equation (3) of the identity change on each response probability, evaluated at sample means; the Δ_k sum to zero by construction. $\Lambda = \Delta_3/(-\Delta_1)$ from equation (4) is the share of the identity-induced shift operating on quantity rather than price. Panel C conditions on the respondent not refusing and estimates (5) and (6); columns (1) and (4) report probit coefficients, (2) and (5) Tobit coefficients in rupees, and (3) and (6) Tobit coefficients in percentage of the offered rent. 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 4.3); Lee (2009) trimming bounds are in Table 23. Standard errors clustered on the respondent in Panels A and C; bootstrap standard errors clustered on the respondent, 1,000 replications, in Panel B. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: The identity penalty rises along the chain; the margin it operates on does not

		(1) Pooled	(2) Tenant	(3) Landlord	(4) Broker
β_1	Identity penalty	0.677***	0.319***	0.999***	1.369***
		(0.045)	(0.085)	(0.150)	(0.119)
Δ_2	Shift onto the price margin	0.007***	0.003***	0.012***	0.017***
		(0.001)	(0.001)	(0.004)	(0.004)
Δ_3	Shift onto the quantity margin	0.200***	0.097***	0.294***	0.372***
		(0.013)	(0.025)	(0.039)	(0.025)
Λ	Rationing share	0.967***	0.965***	0.961***	0.957***
		(0.004)	(0.008)	(0.013)	(0.012)
Observations		4,946	1,390	496	772
Respondents		[N]	[N]	[N]	[N]

Notes. Each column estimates (2) by ordered probit on the indicated subsample, pooling the two clean Hindu → Muslim profiles, and reports the decomposition (3). The unit of observation is a respondent × profile. "Tenant" and "Landlord" are resident respondents who report currently renting and currently letting a property respectively; "Broker" are respondents to the broker survey. "Pooled" column includes tenants, landlords, and respondents who live in their own houses. All respondents answer the vignette as hypothetical landlords, so the contrast across columns isvin who the respondent is, not in what they were asked. $\Lambda = \Delta_3/(-\Delta_1) \in [0, 1]$ is the share of the identity-inducedvshift out of acceptance-at-offered-rent that lands on the quantity marginv(refusal) rather than the price margin (acceptance at a higher rent).v$\Lambda \to 0$ characterises a market that prices the disamenity and houses every applicant; $\Lambda \to 1$ characterises one that refuses. The broker–tenantvdifference in Λ is -0.009 ($p = 0.876$) and should not be read as a gradient; the gradient is in β_1. Bootstrap standard errors clustered on the respondent, 1,000 replications.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

the paper that the reliability null reflects the mechanism rather than inattention. It does not, by itself, establish that the discrimination is taste-based rather than statistical; that limitation is developed as Result 3 below.

4.6 Corroboration: broker access and RWA rules

The vignette evidence is corroborated by two independent elicited measures. Figure 3 reports broker-reported acceptance out of 100 representative properties in the broker's service area. The Muslim figure (23.8) is the lowest of ten categories 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 3.7 points. The exclusion the broker data locate is religion-specific and is not a proxy for low status: Muslim applicants rank below non-vegetarians, inter-caste couples, unmarried men and unmarried women.

Resident welfare associations tell the same story: brokers report that local RWAs bar Muslim tenants in 26.0 percent of areas, against 0.7 percent for upper-caste and 0.8 percent for lower-caste Hindus. This elicited measure can be checked against an external record. TwoCircles.net

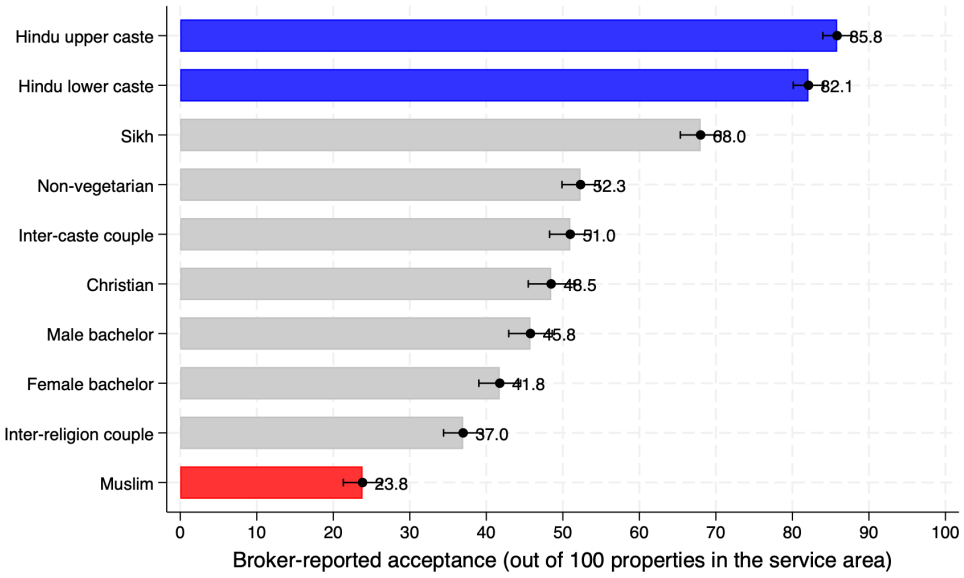


Figure 3: Broker-reported access by applicant type. Each of the 615 surveyed brokers was asked, of 100 representative properties in his service area, how many would accept a tenant of each type; 529 of them served the rental market. Bars are means across brokers and whiskers show 95% confidence intervals. The Muslim–upper-caste-Hindu gap is 62.0 points ($p < 0.01$); the lower-caste-Hindu gap is 3.7 points.

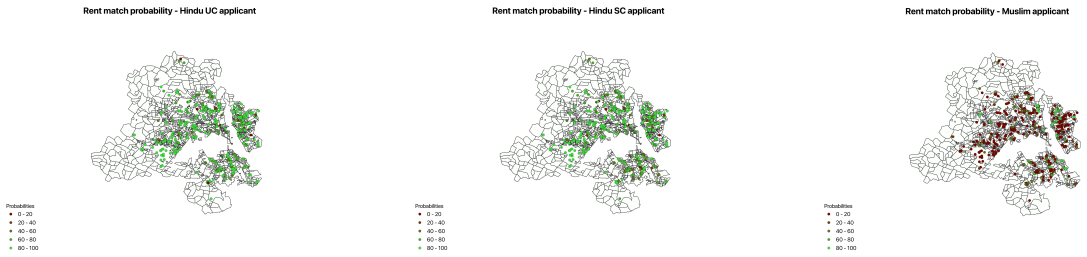
and Newslaundry (2024), working from the membership records of all 1,960 registered housing societies in Delhi, find 68 percent with no Muslim member and aggregate Muslim membership barely above 3 percent. The two measures are not the same object, one is a reported rule about tenancies, the other an administrative record of ownership, but they are of the same order and point the same way, which is the strongest external validation available for any of the elicited measures in this paper.

The access measures are also spatially resolved. Figure 4 maps broker-reported acceptance and Figure 5 broker-reported RWA rules, in each case for Hindu upper caste, Hindu Scheduled Caste and Muslim applicants.

5 Result 2: the wedge between desired and feasible sorting

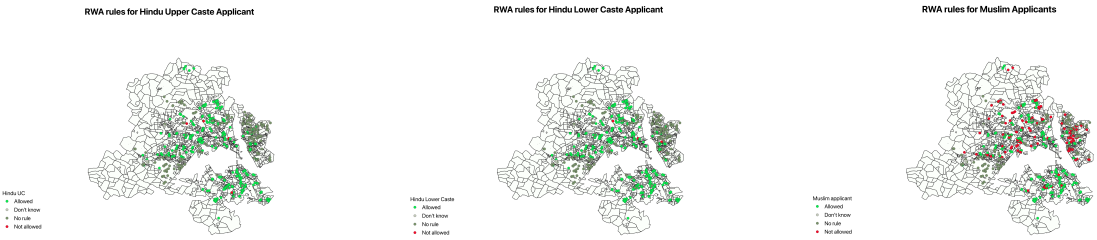
5.1 Own-group preferences are real

I begin with model-free evidence. Because every respondent faces the same hypothetical menus, and every option is stipulated to be affordable and available, the share choosing the own-group alternative can be compared directly between own-group and out-group respondents. Panel A of Table 5 reports these shares. All three groups choose their own-group neighbourhood



(a) Hindu upper caste (b) Hindu Scheduled Caste (c) Muslim

Figure 4: Broker-reported probability that an applicant of each type can find housing in the broker's service area. Each point is a surveyed broker, shaded by his reported acceptance rate out of 100 local properties.



(a) Hindu upper caste (b) Hindu Scheduled Caste (c) Muslim

Figure 5: Broker-reported resident welfare association rules on tenant acceptance, by applicant identity. Each point is a surveyed broker, shaded by whether local RWAs are reported to permit, prohibit, or have no rule on that category.

more often than others do. Upper-caste Hindus show the largest *raw* own-group preference (62.6 against 47.1 per cent), while Muslim and Scheduled Caste own-group areas are rarely chosen by outsiders (9.3 and 16.8 per cent), producing the largest own-versus-other *gaps* (33.6 and 23.6 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.

To put structure on this, let respondent i choose among neighbourhoods n in task t with deterministic utility

$$V_{int} = \beta_R \,\mathrm{Rent}_{nt} + \beta_C \,\mathrm{Commute}_{nt} + \theta_{MUS} \,\mathrm{MuslimArea}_{nt} + \theta_{SC} \,\mathrm{SCArea}_{nt} + \underbrace{\eta_{g(i)}}_{\substack{\text{own-group}\\ \text{preference}}} \mathrm{SameGroup}_{in}, \tag{7}$$

estimated by conditional logit with task fixed effects and standard errors clustered on the respondent, and reported as odds ratios $\exp(\eta_g)$ in Panel B. The estimates confirm strong, significant own-group preference in every specification, but they also 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, because Muslim-composition areas are disfavoured on average by all groups. I therefore rely on the significance and the ordering rather than on the literal magnitude. Stated rankings corroborate the experiment: 73 per cent of Muslim, 70 per cent of Scheduled Caste and 50 per cent of upper-caste Hindu respondents rank own-group composition among their three most important neighbourhood attributes.

5.2 The wedge

The central result combines preferences with access. In the choice experiment Muslim respondents select a mixed or majority-Hindu neighbourhood 57 per cent of the time - a clear stated willingness to live outside a Muslim-majority area. Brokers report that a Muslim applicant can secure housing in about 24 of 100 local properties. Figure 1 sets the two against each other for all three groups, and the contrast is asymmetric in a way that matters for interpretation: for Hindu respondents, both upper-caste and Scheduled Caste, reported access comfortably exceeds stated willingness to move outside the own-group area, so their observed concentration is consistent with preference. For Muslim households it does not. A portion of observed Muslim concentration therefore reflects not preference but the foreclosure of accessible alternatives.

Table 5: Own-group residential preference

Panel A. Choice shares facing identical menus (%)			
Group	Own-group respondents	Out-group respondents	Gap (pp)
Hindu upper caste	62.6	47.1	15.5
Muslim	42.9	9.3	33.6
Scheduled Caste	40.4	16.8	23.6
Panel B. Conditional logit, own-group odds ratios			
	(1) Interactions only	(2) Area-type controls	
Own-group × Hindu upper caste	6.53***	2.35***	
	(0.38)	(0.17)	
Own-group × Muslim	1.90***	20.77***	
	(0.12)	(1.85)	
Own-group × Scheduled Caste	1.38***	3.02***	
	(0.10)	(0.27)	
Area-type controls	No	Yes	
Choice observations	22,752	22,752	
Respondents	[N]	[N]	

Notes. Respondents choose among three hypothetical neighbourhoods varying in monthly rent, commuting time and neighbour composition, having been told that they can afford the rent in every option and should choose the one that suits them best. Availability is therefore fixed by instruction, so the exercise measures preference rather than realised access.

Panel A: for each group, the share of choice tasks in which the own-group-composition option is chosen, computed separately among respondents of that group (column 1) and among respondents of other groups (column 2), over the tasks that offer that group's own-group option. The gap is the difference. The unit of observation is a respondent × task.

Panel B: conditional logit estimates of (7) on the explicit-composition tasks, with task fixed effects, reported as odds ratios $\exp(\eta_g)$ with standard errors in parentheses. Column (2) adds indicators for the area type of each option. The Muslim odds ratio is highly sensitive to this control because Muslim-composition areas are disfavoured on average by all groups . Standard errors clustered on the respondent. The choice-share gaps in Panel A are the preferred, assumption-free summary.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

This is the paper's central object, and it is worth being precise about what it does and does not claim. It does not claim that Muslim households have no own-group preference: Table 5 shows that they have a substantial one, larger on the gap measure than upper-caste Hindus. It claims that the observed degree of separation exceeds what those preferences alone would produce, because the supply side forecloses the majority of the options that 57 per cent of Muslim respondents say they would take. Segregation here is a constrained equilibrium, not a revealed preference.

6 Interpreting the results

Two facts require explanation. Identity operates almost entirely on the quantity margin, uniformly across respondent types. And the penalty rises steeply along a chain of intermediation in which the actor who excludes most is the one who bears none of the risk. This section sets out a simple model that accommodates both and delivers a third implication that the data can test. The model is an interpretation, not an identification strategy: nothing in Section 4 or Section 5 depends on it.

6.1 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. It takes three observationally distinct forms, each measured here with a different instrument, set out in Table 6.

Table 6: Forms of exclusion and their measurement

Form	Mechanism	Measured by
Refusal	Landlord declines at any rent	Identity penalty β_1 (Table 3)
Non-showing	Broker withholds; applicant never learns	Broker access $p_{g,n}$ (Figure 3)
Rule-based	RWA norm makes the let infeasible	Reported RWA rules (Figure 5)

Notes. Two adjacent objects should not be conflated with exclusion. *Price discrimination* is a markup conditional on acceptance (Table 3, Panel C): a movement along the choice set rather than a restriction of it, and Section 4 turns the relative size of the two 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).

6.2 A model of selection from a pool

Environment. A city has K neighbourhoods, $n = 1, \ldots, K$. Neighbourhood n contains N_n rental properties indexed j. Each household i has type $g \in \{HUC, MUS, SC\}$ - Hindu upper caste, Muslim, Scheduled Caste - with population shares μ_g summing to one.

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

$$U_{ij} \;=\; \phi A_{n(j)} \;-\; r_j \;-\; \tau\, c_{i,n(j)} \;+\; \eta_{g(i)}\, s_{g(i),n(j)} \;+\; \xi_{n(j)} \;+\; \varepsilon_{ij}, \tag{8}$$

where A_n is an amenity index, r_j rent, c_{in} commuting cost, $s_{g,n}$ 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_g is own-group residential preference, and it is what the neighbourhood-choice experiment of Section 5.1 identifies.

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

$$\Pi_{ij} \;=\; r_j \;-\; \kappa_{ij} \;-\; d_{g(i),j}, \tag{9}$$

where κ_{ij} is the expected cost of default, damage and failure to recover possession which is large in this market and largely unverifiable ex ante (Section 2.3), and $d_{g,j}$ is the non-pecuniary cost of a type-g tenant, which decomposes as

$$d_{g,j} \;=\; \underbrace{d^{L}_{g,j}}_{\text{landlord's own taste}} \;+\; \underbrace{d^{N}_{g,n(j)}}_{\text{expected neighbour or RWA sanction}}. \tag{10}$$

The decomposition in (10) is not cosmetic. d^L is a preference; d^N is a norm. They have identical implications for behaviour and opposite implications for policy: a preference is a fact about a person, while a norm is sustained by expectations that a regulator can in principle move.

Selection from a pool. The model's only substantive assumption is about what the landlord's problem is.

Assumption 1 (Selection, not acceptance)**.** At the prevailing rent r_j, property j attracts a pool Q_j of applicants with $|Q_j| = q_j \geq 1$, and the posted rent does not vary with the identity of the applicant. The landlord's problem is therefore not to accept or reject a single applicant

but to select

$$i_j^* = \underset{i \in Q_j}{\text{argmax}}\ \Pi_{ij}. \tag{11}$$

Acceptance under (11) requires *winning the pool*, not merely being acceptable. Everything below follows from that substitution, and it is worth being explicit about how much the assumption asks for and what supports it.

It asks for two things: that a landlord facing an unwanted applicant has an alternative use of the property that does not require lowering the rent, and that he does not post applicant-specific rents. It does *not* require the market to be in aggregate excess demand, and I do not claim to have established that it is (Section 2.3). Several features of this market could deliver Assumption 1 without aggregate excess demand: rent control and short unregistered tenancies make posted rents sticky and adjusted at renewal rather than at application; the option value of holding a unit vacant is high when gross yields are thin, so the relevant alternative to letting is often not letting; and quoting a different rent to different applicants for the same property is socially awkward in a face-to-face, broker-mediated transaction in a way that declining is not. Chronic excess demand is one micro-foundation and the most direct one, but it is not the only one, and the results in Section 4 are consistent with all of them.

What the data do establish is the *consequence* that Assumption 1 implies, which is Result 1, and which I measure rather than assume. That is the right logical direction: the paper infers the structure of the market from the margin on which discrimination operates, rather than assuming the structure and predicting the margin.

6.3 Three implications

Result 1 (Discrimination shows up as refusal)**.** *If the landlord could post a type-specific rent he would set $r_j(g) = r_j + d_{g,j}$; every type would then be housed and discrimination would appear entirely as a price premium. Under Assumption 1 the rent is not an instrument he can apply to the applicant, and he has an alternative tenant, so he does not price the disamenity. He selects a different tenant, and discrimination appears as refusal.*

In the data: Rationing share $\Lambda = 0.97$ (Table 3, Panel B; Table 4). Of the response displaced by an identity change, 97 per cent lands on refusal. And it is 0.96 to 0.97 for tenants, landlords and brokers alike i.e., the margin is a property of the market, not of the gatekeeper.

Result 2 (Exclusion is convex in taste)**.** *Take* $d_{HUC,j} = d_{SC,j} = 0$*,* $d_{MUS,j} = d > 0$*, and* κ *type-independent, and suppose the other* $q-1$ *members of the pool are drawn independently from the population. From landlord's problem (equation* (11)*), a Muslim applicant is selected only if no non-Muslim is in the pool, so*

$$\Pr(\textit{accept} \mid MUS) \;=\; \mu_{MUS}^{q-1}. \tag{12}$$

Two things follow, and the first is the more surprising. Acceptance in (12) does not depend on the *magnitude* of d at all, only on its sign: once the landlord is selecting rather than accepting, any strictly positive distaste produces exclusion that falls geometrically in pool size. One does not need extreme animus to generate broker-reported access of 24 in 100.

This resolves what would otherwise be a contradiction in the evidence. The direct measures show severe discriminatory behaviour; the list experiments of Section 7.1 show that almost nobody will admit to discriminating even anonymously. Under (12) both can hold at once, because from any individual gatekeeper's point of view nothing that felt like discrimination occurred, only a mild preference exercised over a pool. The behaviour is severe; the taste need not be. **Discrimination here is an aggregation phenomenon.**

The independence of magnitude is, however, a knife-edge property of the deterministic payoff in (9). Add any idiosyncratic component to Π_{ij} and the magnitude of d re-enters: acceptance is then strictly decreasing in d rather than independent of it, though still far more responsive to pool size than a single-applicant model would suggest. The qualitative point that mild and widely shared distaste is sufficient, still survives; the exact independence does not, and I do not lean on it.

Second, $\Pr(\text{accept} \mid MUS)$ is strictly decreasing in q: exclusion rises with the number of competing applicants. 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 per cent) than for those listed more than seven days (77.5 per cent), 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* (11)*, acceptance depends on the* ranking *within the pool, 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 non-Muslim applicant still wins.

In the data: Information effect $\hat{\beta}_2 \approx 0$ throughout (Table 3, Table 9), while the same experimental arm containing an *income* rather than a reliability signal moves quoted rents sharply (Section 4.5).

Result 3 also bounds what the null can be used for. A zero on β_2 does *not* establish that the 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 (Coate and Loury, 1993; Bohren et al., 2022).

6.4 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 none of κ, his revealed penalty can be written

$$d^B = \underbrace{d_b^{\text{own}}}_{\text{his own taste}} + \underbrace{\mathbb{E}_b\left[d_j^L\right]}_{\text{expected landlord taste}} + \underbrace{\mathbb{E}_b\left[d_n^N\right]}_{\text{expected RWA norm}} . \tag{13}$$

Only the first term is a preference the broker holds; the remainder is derived demand for discrimination. If the broker screens on behalf of the landlord he expects to approach and, given asymmetric costs, plausibly on behalf of a relatively strict one, since 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, then over-exclusion is his optimal response. This is the analogue of the mechanism Christensen and Timmins (2022) document, in which agents cater to the prejudiced beliefs of their clients.

I offer (13) as a reading of the broker–landlord gradient rather than as a tested decomposition. I have attempted to separate the terms using within-broker variation in reported landlord preferences and local RWA restrictiveness, and the estimates are too imprecise to support a claim in either direction.

6.5 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 (Becker, 1957). 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: 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 (discrimination shows up as refusal) supplies a mechanism and Section 4 a measurement. And empirically, these models are typically estimated by inverting observed outcomes, whereas here the choice-set restriction is measured directly, from the gatekeepers who impose it. Table 7 sets out the correspondence; it also indicates why the two literatures are complementary exactly where each is weakest, since the labour literature observes wages but not hiring rejections and this design observes the reverse.

Table 7: Correspondence between the housing and labour markets

Housing	Labour
Landlord	Employer
Broker	Recruiter, temp agency, referral network
Resident welfare association	Closed shop, occupational licensing board
Selection from a pool of applicants	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

6.6 Access and the wedge, formally

Let $a_{g,j}$ denote the probability that an applicant of type g is selected at property j, that is $a_{g,j} = \Pr\Big(g(i_j^*) = g \mid g \in Q_j\Big)$. Neighbourhood-level access for type g is

$$\underbrace{p_{g,n}}_{\text{access}} \;=\; \frac{1}{N_n} \sum_{j \in n} a_{g,j}. \tag{14}$$

Equation (14) is precisely what the broker access measure estimates: the share of properties in a service area at which a given type would be accepted. Combining preferences and access, the probability that i resides in n is the availability-weighted logit

$$P_{in} = \frac{p_{g(i),n} \exp(V_{in})}{\sum_m p_{g(i),m} \exp(V_{im})}, \tag{15}$$

which follows the latent-availability interpretation of Manski (1977) and 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 object Section 5.2 documents and Section 9 prices.

One consequence should be flagged now. Equations (8) and (15) treat r_j as given, and it is not: exclusion redistributes demand, so if Muslim households are barred from a large share of the stock their demand concentrates on the accessible remainder. The counterfactuals in Section 9 hold rents fixed and are therefore partial equilibrium (Bayer et al., 2007).

7 Mechanisms and further evidence

7.1 Enacted rather than owned

Three further measures characterise the *nature* of the discrimination (Table 8). Experienced denial is common and similar across groups in the raw (16 to 17 per cent), but among those denied, Muslim respondents attribute the denial to religion far more often (42.5 against 20.8 per cent). Perceived discrimination is widespread: about 60 per cent of respondents rank religion among the three most important factors landlords use. Yet the list experiments recover almost no *admitted* discrimination.

The contrast is a finding in itself, and it is what Result 2 (Exclusion is convex in taste) predicts: direct measures show severe discriminatory behaviour while the social-desirability-robust measure shows that few will own it. Applied to attitudes rather than behaviour, the same indirect method 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. That is a reservoir of majoritarian sentiment, and its coexistence with a null on admitted discrimination is informative. The reading I favour is that exclusion is not experienced by those who practise it as a choice about a person but as compliance with what others will

accept.

Two caveats bound how hard this can be pushed, and I state them rather than reporting the nulls as findings. The treated cells for the behavioural items are small - 486 tenants and 114 landlords, so the landlord estimate of 0.095 with a standard error of 0.170 rules out only large admitted rates, not modest ones. And the behavioural items are worded in a self-incriminating way ("declined a tenant on the basis of caste or religion") while the attitudinal items are not, so part of the contrast may be a wording effect rather than a difference in what respondents will admit. That reading, if correct, reinforces rather than undermines the interpretation above: respondents will endorse the norm and not the act.

Table 8: Perceived, experienced, and admitted discrimination

Panel A. Direct survey measures, by respondent religion		
	Non-Muslim	Muslim
Experienced a denial of housing	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 the sensitive item)		
Behavioural items:		
Tenant: faced identity discrimination	0.011	(0.083)
Landlord: declined a tenant on identity	0.095	(0.170)
Attitudinal items:		
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 A: shares of resident respondents, one observation per respondent. "Experienced a denial" records whether the respondent reports having been refused housing; the second row conditions on that and reports the share attributing the refusal to religion. "Top-3 factors" records whether the respondent ranks religion among the three attributes they believe landlords weigh most heavily.
Panel B: list (item-count) experiments (Imai, 2011; Blair and Imai, 2012). Respondents report only the *number* of items on a list that apply to them; a randomly assigned treatment group sees one additional sensitive item. Each entry is the treatment−control difference in mean item counts, which identifies the prevalence of the sensitive item, with robust standard errors in parentheses. Treated cells are 486 for the tenant item and 114 for the landlord item, so the behavioural nulls rule out only *large* admitted rates. Attitudinal items pool all respondents
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.2 What the accessible set looks like

Exclusion confines households to a subset of neighbourhoods. Whether that is costly depends partly on how those neighbourhoods are served. I construct amenity densities from geocoded points of interest, normalised by colony area and transformed by the inverse hyperbolic sine, and separate them into services provided publicly (hospitals, parks, police posts, water infrastructure) and privately (clinics, sports facilities, monuments). All specifications condition on colony classification, the Delhi tenure ladder running from approved plotted development through regularised and unregularised unauthorised colonies to informal settlements, and on total point-of-interest density, which proxies how thoroughly an area is mapped. Construction, full results and robustness are in Appendix E; the summary is short.

Public provision shows no gradient in religious composition (-0.037, s.e. 0.269). Private provision does (-0.696, s.e. 0.161), and the difference between the two is itself significant (0.660, s.e. 0.317, $p = 0.039$). The gap is concentrated in medical clinics (-0.887, $p < 0.001$), smaller in sports facilities (-0.241, $p = 0.005$), and absent for historical monuments (-0.059, $p = 0.158$), which is the falsification one wants given that Delhi's Muslim-concentrated neighbourhoods are disproportionately located in the historic city. The sharpest version is within health provision: hospitals, overwhelmingly state-provided and sited by planning decision, show no gap; clinics, private and sited by practitioners choosing where to establish, show the largest one in the data. The same service is provided evenly by the state and unevenly by the market.

Two limits on the interpretation. These are counts of facilities, not measures of quality, capacity or whether a facility functions, and the limitation is asymmetric: for small, numerous, competitive private services a count is a reasonable proxy for what a household can obtain, while for public services it is a poor one, since a state hospital exists in the data whether or not it is staffed. The public-provision null is therefore an absence of evidence rather than evidence of absence. And the result is a cross-sectional association between composition and provision, not a causal effect of exclusion on provision.

8 Robustness and alternative measures

8.1 Placebo arms

The vignette instrument contains three profiles in which arms A and B carry *identical* names - Sanjay Kumar, Aakriti Patel and Akbar Khan (Tables 12 and 13). For these the A $\rightarrow$ B contrast

changes nothing at all, so β_1 should be zero by construction. It is: the estimated penalty is statistically indistinguishable from zero for all three profiles (**??**). This is the strongest internal validation the design offers. It rules out the possibility that the measured penalty reflects an arm effect, a position effect, or any feature of the instrument other than the name.

8.2 The reverse transition

One profile runs the identity change in the opposite direction, replacing a Muslim-named applicant (Gulnaz Bano, a single police constable) with an otherwise identical Hindu-named one (Pooja Bhatia). Acceptance *rises* by 10.6 percentage points. The design therefore passes an internal-consistency check, and the estimate supports the monotonicity condition required by the Lee (2009) bounds of Section 4.3.

8.3 Binary coding of acceptance

The ladder of (2) is the preferred specification because it uses every response category. For comparability with the correspondence literature I also report the binary specification

$$\text{Accept}_i \;=\; \beta_0 \;+\; \beta_1\,\mathbf{1}\{B_i \text{ or } C_i\} \;+\; \beta_2\,\mathbf{1}\{C_i\} \;+\; \varepsilon_i, \tag{16}$$

where acceptance pools accept-at-offered-rent and accept-at-a-higher-rent against refusal. Table 9 reports it. Switching an otherwise identical profile to a Muslim name lowers acceptance by 18.8 percentage points for the married woman and 22.9 for the married man, and the penalty rises monotonically along the chain of intermediation. Results are unchanged if only outright accepts and rejects are used.

8.4 Dropping the ordering assumption

The ladder imposes that acceptance at a higher rent lies between acceptance at the offered rent and refusal on a single latent scale. The interior cutpoints are close together, so the ordering of the two interior states is not strongly supported by the data, and a reader may reasonably doubt that quoting a markup is "less reluctant" than refusing. I therefore re-estimate the decomposition from a multinomial logit over the three unordered outcomes of (1), which imposes no ordering at all. to be added

Table 9: Binary specification: vignette identity penalty by respondent type

	(1) Pooled	(2) Tenant	(3) Landlord	(4) Broker
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
Respondents	[N]	[N]	[N]	[N]

Notes. OLS estimates of (16). The unit of observation is a respondent × vignette profile. 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. Arm A shows a Hindu-named applicant; arm B shows an otherwise identical profile with a Muslim name; arm C shows arm B plus a tenant-reliability signal. β_1 is therefore the effect of changing the applicant's name from Hindu to Muslim holding every other attribute and the photograph fixed; β_2 is the incremental effect of the reliability signal. All respondents, including brokers, answer as hypothetical landlords. Standard errors clustered on the respondent, since each respondent evaluates eight profiles. The landlord column rests on 263 respondents and is correspondingly imprecise.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

8.5 Inference

Each respondent evaluates eight vignette profiles and four choice tasks, so all standard errors in the paper are clustered on the respondent. Standard errors for Λ, which is a ratio of marginal effects, come from a nonparametric bootstrap clustered on the respondent with 1,000 replications.to be added

8.6 Two preference elicitations

The choice instrument contains four tasks: two describing neighbourhoods by religious composition and two by amenity bundle. They do not identify the same parameters, and the reason is a design limitation worth stating plainly rather than working around.

In the *composition* tasks the own-group option is also the most expensive and the least distant in every task in which it appears, so rent is positively correlated with own-group composition by construction. Estimated on those tasks the rent coefficient is +0.613 and the commute coefficient +0.109: respondents appear to pay more for longer commutes. Those tasks identify own-group preference but cannot identify a price coefficient.

In the *amenity* tasks the listed bundles differ completely across the three options, so the amenity indicators are option indicators. A variable that identifies the alternative exhausts the two degrees of freedom available in a three-option task and its coefficient is not separately identified. Willingness to pay for a metro station, a school or a park therefore cannot be recovered from this design. What these tasks *can* identify is the price coefficient, because rent varies within task while the religious institutions in the bundles function as composition signals rather than as facilities. Table 10 establishes that respondents read them that way: a mosque is not a facility a Hindu household can use, yet Hindu respondents avoid mosque-containing options at more than five standard errors while Muslim respondents are drawn to them and are indifferent to temples.

The two elicitations give different magnitudes and a different ordering across groups, and I report both rather than choosing. The common finding, and the one the analysis relies on, is that own-group preference is substantial for every group. The instrument should be corrected in a future round by varying one attribute at a time across options.

Table 10: Religious institutions are read as composition signals, not as amenities

	Mosque present in the bundle	Temple present in the bundle
Hindu upper caste	-0.633^{***}	0.610^{***}
	(0.119)	(0.131)
Muslim	0.519^{***}	0.032
	(0.137)	(0.146)
Hindu lower caste	-0.536^{***}	0.462^{***}
	(0.123)	(0.134)
Choice observations	20,826	
Respondents	[N]	

Notes. Conditional logit on the two amenity tasks of the neighbourhood-choice experiment, in which options are described by a bundle of local facilities rather than by the composition of the neighbourhood. Entries are coefficients on an indicator for whether the option's bundle includes a mosque (column 1) or a temple (column 2), estimated separately by respondent group. Rent, commuting time and task fixed effects are included throughout. Standard errors clustered on the respondent. A mosque is not a facility a Hindu household consumes, so the pattern is what a composition reading predicts and not what an amenity reading predicts.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

8.7 Is this generic minority disadvantage? The caste comparison

Scheduled Caste applicants are a stigmatised minority who are neither religiously distinct from the majority nor distinguishable by the markers that separate Muslims from Hindus, so the comparison separates religion-specific exclusion from generic minority disadvantage. Three pieces of evidence point the same way. In the broker data, the Muslim access gap relative to upper-caste Hindus is 62.0 percentage points while the lower-caste Hindu gap is 3.7 (Figure 3). In the RWA data, Muslim tenants are reported barred in 26.0 per cent of areas against 0.8 per cent for lower-caste Hindus. And in the vignette instrument the one available caste-surname contrast replaces a Muslim name with "Manjhi", a Scheduled-Caste-associated surname, and acceptance *rises* by 14.0 points ($p < 0.01$).

The vignette contrast is the weakest of the three, because the instrument identifies caste only through surnames and does not contain a clean upper-caste-versus-Scheduled-Caste transition. The credible caste evidence is the broker and RWA comparison. Whatever is operating in this market is religion-specific and is not a proxy for poverty, low status, or a general tendency of gatekeepers to prefer applicants like themselves. Full caste-specific analysis is in Appendix A.

8.8 Biases and bounds

The paper makes several bounding claims, in different directions and about different objects, and it is easy for a reader to think they conflict. Table 11 sets them out together.

Table 11: Biases and the direction of the resulting bound

Concern	Direction	Why	Evidence or check
Social desirability in the vignette	Lower bound on discrimination	Respondents may wish to appear secular in a hypothetical exercise and overstate acceptance of minority applicants, which would attenuate β_1.	Cross-instrument agreement with elicited access and RWA rules; survey framed as a study of housing decisions, not discrimination (Section 3.2).
Selection into the markup sample	Lower bound on the price penalty	Only non-refusers can quote a price, and identity induces refusal most among those who object most, so the markup sample is the least discriminating part of the distribution (Section 4.3).	Lee (2009) bounds, Table 23; the reverse transition supports monotonicity.
Broker self-report of access	Conservative: understates the gap	Any incentive to appear non-discriminatory pushes reported Muslim access up toward reported Hindu access.	External benchmark against housing-society membership records (TwoCircles.net and Newslaundry, 2024).
Amenity counts omitted from the welfare exercise	Lower bound on the cost of exclusion	Utility varies across areas only through composition, which is the one dimension on which exclusion does not disadvantage Muslim households (Section 9).	Section 7.2 shows a private-provision gap that the exercise cannot price.
Rents held fixed	Upper bound on the gain from equalising access	Equalising access would shift demand across areas and rents would adjust, partly offsetting the gain (Bayer et al., 2007).	Partial-equilibrium by construction; not estimated.

Notes. The last two rows bound *different* objects and therefore do not conflict: the amenity omission bounds the cost of the constraint that exists, while the fixed-rent assumption bounds the benefit of removing it.

8.9 Do the stated measures track behaviour?

Every estimate in this paper is elicited rather than revealed: preferences from a hypothetical choice experiment, gatekeeping from hypothetical vignettes, access and RWA rules from broker reports. This is the paper's principal vulnerability and it deserves a direct answer rather than a caveat.

Four things support the measures. First, an *external behavioural benchmark* exists for this city: Datta and Pathania (2016) run a correspondence study in the Delhi rental market and find landlords markedly less responsive to Muslim than to comparable upper-caste Hindu applicants. Second, *cross-instrument agreement*: the vignette penalty, the broker access measure and the RWA-rule measure are three instruments with different response formats, different respondents and different failure modes, and they agree. Third, *administrative corroboration*: the elicited RWA prohibition rate is of the same order as the housing-society membership records of TwoCircles.net and Newslaundry (2024). Fourth, the *methodological literature* finds conjoint and vignette designs to track real behaviour reasonably well (Hainmueller et al., 2015), and supplies bounds on experimenter demand (de Quidt et al., 2018) and design standards for survey experiments (Haaland et al., 2023).

None of this makes elicited measures equivalent to revealed ones. It makes the specific claim - that Muslim applicants face large exclusion on the quantity margin in this market, robust to the failure of any one instrument.

9 What the access constraint costs

The reduced-form results establish that access differs sharply by identity. They do not say what that costs, because the cost of losing a neighbourhood depends on what the household would have obtained from it. This section prices the constraint using a sorting model in the tradition of Bayer et al. (2007) and Christensen and Timmins (2023), with one departure: rather than recovering preferences by inverting observed residential choices, I take them from the choice experiment, in which rent is randomised across options and availability is fixed by instruction. This matters concretely. Estimated on observed residence in these data the coefficient on neighbourhood rent is -0.044 (s.e. 0.095) and the interaction with the top income tercile implies a *positive* price coefficient for the richest households: the rent index is proxying unobserved neighbourhood quality. Since the marginal utility of income is the denominator of every welfare

figure, that is not a defect one can report around.

Construction. Let $p_{g,n}$ be the access probability of neighbourhood level access (equation (14)), estimated from the broker measure. Following Manski (1977), treat it as the latent availability of alternative n, so that expected utility over the unrestricted and restricted menus is $IV_i^{\text{unc}} = \ln \sum_n \exp(V_{in})$ and $IV_i^{\text{con}} = \ln \sum_n p_{g(i),n} \exp(V_{in})$. Because rent enters utility linearly the marginal utility of income is constant and the compensating variation of Small and Rosen (1981) applies exactly, $CV_i = |\beta_R|^{-1}[IV_i^{\text{unc}} - IV_i^{\text{con}}] \geq 0$. I report the *differential* constraint, evaluated at Muslim access less the same expression at upper-caste Hindu access, which isolates the cost of discrimination from the search frictions any household faces. The menu is the 268 areas for which both broker-reported access and a composition classification are available: 246 Hindu-majority, 16 Muslim-majority and 6 with no clear majority. A *variety–targeted* decomposition separates what any restriction of that size would cost from what is attributable to *which* areas are foreclosed, using a placebo constraint that removes the same number of areas at random, averaged over fifty draws.

Two results, and what they are worth. The constraint facing Muslim households is an order of magnitude larger than that facing upper-caste Hindus: a loss in expected utility of 1.129 log points against 0.124, a ratio of 9.1. Under the stated-preference price coefficient the differential constraint facing the median Muslim household is worth about ₹12,000 per month.

I want to be exact about what this magnitude adds, because the answer is: not much. Utility varies across areas only through composition, and 92 per cent of the areas in the menu share the same classification, so the inclusive value is close to a count of alternatives. The variety term equals $-\ln \bar{p}_g$ to three decimal places for every group. The welfare figure is therefore a monotone transformation of the access measure already reported in Section 4.6, expressed in utility units. It is not independent evidence and I do not present it as such. Its purpose is to put the access constraint on a scale comparable to the price penalty of Section 4.3, and on that comparison the ordering is clear: the welfare cost of foreclosed access is large relative to a price penalty bounded between ₹48 and ₹116 and unsigned. That is Result 1 (discrimination shows up as refusal) restated in money.

The decomposition is what the exercise does deliver. The *targeted* component is negative for every group: the actual constraint costs *less* than a random restriction of the same size (-0.067 log points for Muslim households, -0.003 for both Hindu groups). Access is positively correlated

with own-group composition for everyone, because gatekeeping sorts households toward areas resembling them, and those are also the areas at which the composition term raises utility. Two things follow. The exclusion Muslim households face is close to *uniform* across neighbourhood types rather than concentrated on the areas they would most value; and to the limited extent that it is targeted, it pushes them toward areas they mildly prefer. The cost of exclusion here is the foreclosure of options, not relegation to worse ones.

Caveats. Six, stated compactly. *The model captures one channel*: rent is held fixed across alternatives and amenity coefficients are not identified, so utility varies only through composition, which is the one dimension on which exclusion does not disadvantage Muslim households. The omission is not neutral - Section 7.2 shows a private-provision gap that the exercise cannot price, so it drops a channel running against these households while keeping one running in their favour. *Composition is nearly degenerate*: 246 of 268 areas are Hindu-majority, which is why the inclusive value reduces to a count. *Commuting drops out*: the survey records commuting time for the current residence only, so it is unobserved for counterfactual areas. *The access measure is elicited, not experimental* (Section 8.9). *Preference parameters matter*: using the composition-task estimates of η_g in place of the amenity-task estimates gives a differential constraint of 0.777 log points for Muslim households rather than 1.005, and reverses the ordering across groups, so this is a sensitivity check rather than a bound. *Rents are held fixed*, so the exercise is partial equilibrium and the estimate is an upper bound on what removing the constraint would deliver (Bayer et al., 2007).

10 Policy: discrimination as a tractable enforcement problem

10.1 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 and is renewable, and operating unregistered attracts a penalty of ₹10,000 per day of default under section 62. 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: extend RERA agent registration to rental intermediaries, 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. 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 (Turner et al., 2002; Neumark, 2018).

10.2 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.

10.3 Where enforcement binds

The exclusion documented here is close to uniform across neighbourhood types (Section 9), which has a direct implication for enforcement design: there is no subset of neighbourhoods in which the constraint is concentrated and on which effort could be economically targeted. A registration-and-conduct regime applied across the rental brokerage market is therefore the appropriate instrument, rather than geographically targeted inspection.

11 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 neighbour-

hoods through randomized choice and vignette experiments, gatekeeper elicitation, and a social-desirability-robust list experiment, it separates two forces that the prior literature has been able to study only one at a time.

Two results carry the paper. Identity operates on this market almost entirely through refusal rather than price: of the response displaced by an identity change, roughly 97 per cent lands on rejection and under 3 per cent on a higher quoted rent, and that proportion is the same for tenants, landlords and brokers even though the level of exclusion differs by a factor of four across them. And desired and feasible sorting diverge, for Muslim households specifically: most would choose a mixed or majority-Hindu neighbourhood, and brokers report they can reach about a quarter of the market. 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, and how to measure which margin discrimination operates on rather than assuming it - both of which travel beyond this setting. Practically, it locates the binding constraint in identifiable institutions whose conduct is a tractable target for enforcement, in a jurisdiction where no general anti-discrimination right in private housing exists to enforce.

Several questions remain open. The welfare exercise shows that exclusion here forecloses options rather than relegating households to systematically worse ones, though the neighbourhoods minority households can reach do have measurably less private service provision, which the exercise cannot price. The mechanism that produces the quantity margin is identified by the model rather than by the data, and a direct measurement of local market tightness would test it. And the decomposition of the broker's penalty into his own taste and the demand for discrimination he faces is, on present power, out of reach; it is the natural target for a follow-up design.

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.

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.

Arrow, K. J. (1973). The theory of discrimination. In Ashenfelter, O. and Rees, A., editors, *Discrimination in Labor Markets*, pages 3–33. Princeton University Press, Princeton, NJ.

Asher, S., Jha, K., Novosad, P., Adukia, A., and Tan, B. (2026). Residential segregation and unequal access to local public services in India: Evidence from 1.5m neighborhoods. *American Economic Review*. Forthcoming. NBER Working Paper 34818.

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.

Banerjee, A., Bertrand, M., Datta, S., and Mullainathan, S. (2009). Labor market discrimination in Delhi: Evidence from a field experiment. *Journal of Comparative Economics*, 37(1):14–27.

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.

Becker, G. S. (1957). *The Economics of Discrimination*. University of Chicago Press, Chicago.

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.

Bharathi, N., Malghan, D., and Rahman, A. (2021). Isolated by caste: Neighbourhood-scale residential segregation in Indian metros. *Journal of Regional Science*.

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

Blair, G. and Imai, K. (2012). Statistical analysis of list experiments. *Political Analysis*, 20(1):47–77.

Bohren, J. A., Hull, P., and Imas, A. (2022). Systemic discrimination: Theory and measurement. *NBER Working Paper 29820*.

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

Card, D., Mas, A., and Rothstein, J. (2008). Tipping and the dynamics of segregation. *Quarterly Journal of Economics*, 123(1):177–218.

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. 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.

Coate, S. and Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review*, 83(5):1220–1240.

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

Cutler, D. M. and Glaeser, E. L. (1997). Are ghettos good or bad? *Quarterly Journal of Economics*, 112(3):827–872.

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.

de Quidt, J., Haushofer, J., and Roth, C. (2018). Measuring and bounding experimenter demand. *American Economic Review*, 108(11):3266–3302.

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.

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.

Glaeser, E. L. and Luttmer, E. F. P. (2003). The misallocation of housing under rent control. *American Economic Review*, 93(4):1027–1046.

Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing information provision experiments. *Journal of Economic Literature*, 61(1):3–40.

Hainmueller, J., Hangartner, D., and Yamamoto, T. (2015). Validating vignette and conjoint survey experiments against real-world behavior. *Proceedings of the National Academy of Sciences*, 112(8):2395–2400.

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.

Imai, K. (2011). Multivariate regression analysis for the item count technique. *Journal of the American Statistical Association*, 106(494):407–416.

Kline, P., Rose, E. K., and Walters, C. R. (2022). Systemic discrimination among large u.s. employers. *Quarterly Journal of Economics*, 137(4):1963–2036.

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.

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.

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

Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature*, 56(3):799–866.

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

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].

Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review*, 62(4):659–661.

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.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186.

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

Siddique, Z. (2011). Evidence on caste based discrimination. *Labour Economics*, 18:S146–S159. Supplement 1.

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

Small, K. A. and Rosen, H. S. (1981). Applied welfare economics with discrete choice models. *Econometrica*, 49(1):105–130.

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.

Thorat, S. and Attewell, P. (2007). The legacy of social exclusion: A correspondence study of job discrimination in India. *Economic & Political Weekly*, 42(41):4141–4145.

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., Ross, S. L., Galster, G. C., and Yinger, J. (2002). Discrimination in metropolitan housing markets: National results from phase i of hds 2000. Technical report, The Urban Institute, Washington, DC. Prepared for the U.S. Department of Housing and Urban Development.

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.

World Bank (2020). Doing business 2020: Comparing business regulation in 190 economies. Technical report, World Bank, Washington, DC. Enforcing Contracts indicator, India.

A Caste-Specific Results

B Anatomy of a Rental Transaction

Brokers, community networks and online portals are the main avenues for rental housing search in Delhi and in most major Indian cities, with over 90 per cent of the tenants in the survey reporting having used one or more of these. The remainder rely on door-to-door searches for 'To-let' signs. Brokers often post and manage advertisements on housing portals at the behest of landlords; some portals market themselves specifically on putting tenants in direct contact with owners and saving both sides the brokerage.

Brokers handle the administrative documents and pre-screen potential tenants before taking the shortlisted applicants to the landlords. This was confirmed in focus group discussions and informal conversations with brokers in Delhi. Landlords specify the type of tenant they want and brokers screen 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, usually taking the prospective tenant to the property himself. He typically charges half a month's rent from the tenant and the same from the landlord once the deal is formalised, though this is negotiable; there is normally no fee for showing a property, and brokers service an area of roughly 1.5 to 2 kilometres around their shop.

Post-viewing there may be negotiation over rent and deposit, in which the broker generally conducts the first round. Landlords specify an ideal rent and the amount they would be willing to concede, and applications may be rejected by the broker if they fall below the landlord's expectations. In some cases the landlord joins the negotiation at the broker's request.

Once rent and deposit are settled the applicant pays a token amount to hold the property pending checks such as police verification. The parties then sign an eleven-month leave-and-licence agreement, renewable before expiry and generally with some increase in rent, although a majority of agreements end up informal.

The feature that matters for measurement is that pre-screening precedes the viewing. Exclusion at the pre-screen is invisible to the applicant, who never learns which properties existed.

C Survey Experimental Materials

This appendix presents the experimental materials used in the survey: the landlord applicant-evaluation vignettes, the landlord pairwise tenant-choice module, and the tenant neighbourhood-choice tasks.

C.1 Landlord Game 1: Applicant Profiles

Respondents were shown eight applicant profiles and asked to evaluate each as if they were the landlord of an advertised property. Arms A and B differ only in the applicant's name; arm C adds a reliability or verification cue to arm B. The photograph is identical across the three arms within each profile. Photographs were AI-generated with prompt text specifying that the depicted person should carry no religious markers, and were field-tested with survey staff to confirm that they did not read as identifiably Hindu or Muslim.

Profiles 2, 5 and 7 (Sanjay Kumar, Aakriti Patel and Akbar Khan) carry *identical* names in arms A and B and are therefore placebo transitions; they are analysed in Section 8.1. Profile 6 (Gulnaz Bano → Pooja Bhatia) runs the identity change in reverse (Section 8.2). Profile 3 (Shakeel Ahmed → Soham Manjhi) is the Muslim-to-Scheduled-Caste contrast (Section 8.7), and profile 4 (Arun Kumar → Atul Pandey) is a within-Hindu caste-marked surname contrast [currently unanalysed].

Table 12: Landlord Game 1: applicant profiles 1–4

Photo	Arm A	Arm B	Arm 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 13: Landlord Game 1: applicant profiles 5–8

Photo	Arm A	Arm B	Arm 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

C.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. Results are reported in **??**.

Table 14: 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 15: 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

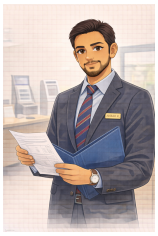
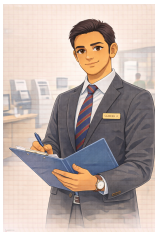
Table 16: Landlord Game 2: tenant pair 3

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 17: 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 18: 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

C.3 Tenant Neighbourhood-Choice Games

Respondents were presented with hypothetical residential areas varying in rent, commute time, and either neighbourhood demographics or available amenities, and were told they could afford the rent in every option.

Table 19: Tenant game: neighbourhood demographics I

Area	Rent (₹)	Commute (min)	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 20: Tenant game: neighbourhood demographics II

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

Table 21: Tenant game: neighbourhood amenities I

Area	Rent (₹)	Commute (min)	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 22: Tenant game: neighbourhood amenities II

Area	Rent (₹)	Commute (min)	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

D Lee Bounds on the Price Penalty

Table 23: Lee (2009) trimming bounds on the identity effect on the quoted markup

	(1) ₹, married woman	(2) ₹, married man	(3) % of rent, woman	(4) % of rent, man
Lower bound	48.34	−40.14	0.220	−0.401
	(46.36)	(102.50)	(0.211)	(1.025)
Upper bound	115.93***	341.81***	0.527***	3.418***
	(30.67)	(46.23)	(0.139)	(0.462)
Trimming proportion	0.221	0.303	0.221	0.303
Observations	3,909	3,909	3,909	3,909

Notes. Lee (2009) trimming bounds on the effect of changing the applicant's name from Hindu to Muslim on the markup quoted over the offered rent, estimated using Tauchmann (2014). The estimand differs from the Tobit coefficient in Table 3, Panel C: these bounds are on the effect among respondents who would not refuse under either identity, whereas the Tobit conditions on observed non-refusal. The bounds require no exclusion restriction and no parametric selection model, only monotonicity — that the identity change weakly reduces the probability of quoting rather than refusing for every respondent — which Section 8.2 supports empirically. Because the identity change reduces selection, the lower bound trims the upper tail of the treated distribution and the upper bound the lower tail; the trimming proportion is the share trimmed.

Two caveats. The estimator assumes a continuous outcome, and quoted markups cluster heavily on round numbers: 18 of the 35 markups for the married woman are exactly ₹3,000 and 42 of the 115 for the married man are exactly ₹5,000, so ties at the trimming quantile are a concern. And the estimator admits no covariates, so it presumes random assignment, which the arm randomisation delivers by construction. Confidence intervals for the effect include zero in both cases, so the sign of the price penalty is not identified after accounting for selection.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E Amenity Provision by Neighbourhood Composition

This appendix reports the provision results summarised in Section 7.2 and the checks behind them.

Construction

Amenity counts are scraped from Google Maps by keyword and attached to the colony in which they fall. Each count is divided by colony area in square kilometres and transformed by the inverse hyperbolic sine, which accommodates the zeros that a log cannot. Categories are grouped by provision mechanism: *public* (hospitals, parks, police posts, water infrastructure) and *private* (clinics, sports facilities, historical monuments).

Schools are excluded from both. Delhi's school stock mixes government, aided and private institutions that the scraped data cannot separate, so assigning them to either group would misclassify a large share; dropping them does not affect the results, since the public index including schools gives a coefficient of -0.009 against -0.037 without. Religious institutions are excluded throughout, because the sampling strata of Section 3.1 are defined by mosque and temple presence, so including them as amenities while conditioning on religious composition would be circular.

Each index is the row mean of its transformed components, standardised across areas. All specifications condition on *colony classification*, Delhi's administrative tenure ladder collapsed from eleven categories to six (approved plotted development and group housing; regularised unauthorised colonies; unregularised unauthorised colonies; informal settlements; villages; industrial areas), and on *point-of-interest density*, the total count of scraped locations of any type per square kilometre, which proxies how thoroughly an area is represented in the underlying data.

Results

Point-of-interest density enters both level regressions strongly but the difference regression weakly ($t = 1.16$, $p = 0.25$), so uneven coverage in the underlying data affects public and private categories similarly and largely differences out. Colony classification predicts provision *levels* strongly but the public–private *balance* not at all ($F = 1.28$, $p = 0.26$): tenure status determines how much provision an area receives, composition determines what kind.

Table 24: Amenity provision and neighbourhood religious composition

	Provision indices		Private categories		
	Public	Private	Clinics	Sports	Monuments
	(1)	(2)	(3)	(4)	(5)
Muslim share	−0.037	−0.696***	−0.887***	−0.241***	−0.059
	(0.269)	(0.161)	(0.243)	(0.086)	(0.042)
POI density	0.010***	0.006***	[ ]	[ ]	[ ]
	(0.002)	(0.002)			
Colony classification FE	Yes	Yes	Yes	Yes	Yes
R^2	0.223	0.151	[ ]	[ ]	[ ]
Observations	177	177	177	177	177
Difference, public − private					
Muslim share	0.660**				
	(0.317)				
POI density	0.004				
	(0.003)				

Notes. OLS with robust standard errors; the unit of observation is a colony. The dependent variable in columns (1) and (2) is the standardised provision index described above; in columns (3) to (5) it is the inverse hyperbolic sine of the category's density per square kilometre. "Muslim share" is the share of surveyed residents in the colony reporting Islam. "POI density" is the total count of scraped points of interest of any type per square kilometre, included to absorb differences in how thoroughly areas are mapped. The lower panel reports the same specification with the difference between the two standardised indices as the dependent variable, which tests whether composition predicts the public–private *balance* rather than the level. Column (5) is a falsification: Delhi's Muslim-concentrated neighbourhoods are disproportionately located in the historic city, so a monument gradient would have indicated that the private index was picking up that geography.

These are cross-sectional associations between composition and provision, not causal effects of exclusion on provision.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Robustness to local purchasing power

A natural alternative explanation is demand rather than composition: private services locate where households can pay for them, and Muslim-concentrated neighbourhoods may be poorer.

Table 25: Private provision and Muslim share, conditioning on local income

	(1)	(2)	(3)
Muslim share	-0.696^{***}	-0.716^{***}	-0.773^{**}
	(0.161)	(0.219)	(0.348)
Log mean income			-0.002
			(0.242)
Sample	All colonies	≥ 5 tenants	≥ 5 tenants
Income control	No	No	Yes
POI density	Yes	Yes	Yes
Colony class. FE	Yes	Yes	Yes
Observations	177	100	62

Notes. Dependent variable is the standardised private-provision index (clinics and sports facilities); the unit of observation is a colony. Income is imputed for each surveyed tenant as reported rent divided by the mid-point of the reported rent-burden category, then averaged over the colony. Columns (2) and (3) restrict to colonies with at least five surveyed tenants, below which the colony mean is too noisy to be informative; column (3) further restricts to colonies for which the imputed mean is defined. The two restrictions separate the two things that might have been responsible: the result is not driven by the colonies lost when the sample is restricted, and it is not driven by local purchasing power. Robust standard errors in parentheses.

The income coefficient should not be read as establishing that demand is irrelevant: it is imputed from a coarse categorical variable and averaged over few households, so it rules out a purchasing-power explanation only to the extent that colony mean income captures the relevant variation in local demand.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The coefficient is stable across the three columns. Restricting to better-surveyed colonies moves it from -0.696 to -0.716; adding the income control moves it to -0.773, well inside one standard error.

What these measures capture

The indices are counts of facilities per unit area. They are not measures of quality, capacity, or whether a facility functions, and the literature this connects to (Asher et al., 2026) operates under the same constraint. The limitation is asymmetric, which matters for the relative weight the two results can bear. For privately provided services a count is a reasonable proxy for what

a household can obtain: clinics are small, numerous and competitive, so their density plausibly tracks access and waiting times. For publicly provided services it is a poor proxy: a state hospital appears in the data whether or not it is staffed, a park whether or not it is maintained. The public-provision null should therefore be read as an absence of evidence rather than as evidence of absence, and is entirely consistent with an unmeasured gradient in the quality of public services.

A second limitation follows from the source. Point-of-interest data reflect what has been recorded as much as what exists, and commercial listings — clinics among them — are the category where that gap is likely to be widest in dense, informal neighbourhoods. The public–private contrast is the main protection against this, since both categories are scraped identically and the mapping-intensity proxy differences out between them. An administrative register of licensed clinics, matched to even a subset of colonies, would settle the question more convincingly than any further specification.

F Additional Robustness Tables