The Spatial Cost of Formality: Street Vendors' Preferences in India

Muskan Jain & **Lachi Singh** (Job Market Paper)*

November 6, 2026

Please click here for the latest version

Abstract

Formalisation policies in developing countries have proven hard to get right, and how much they achieve is sensitive to policy design. The problem is harder still for informal firms trading on public land, because a state cannot grant such a firm legal status without also deciding where it may trade. A large group of these firms is street vendors, who sell from carts, stalls and ground displays on streets and pavements and account for between 1.5 and 6.8 per cent of employment in low- and middle-income countries. States therefore offer them formality at a location rather than at a price, and deliver it through government-built vending zones that commonly stand empty. We price what relocation costs these firms, using a stated-preference experiment embedded in an original geolocated survey of 2,922 street vendors across 206 marketplaces in Patna, India, where the Street Vendors Act of 2014 makes certification free and attaches the protected spot to a location instead. Vendors value formality: at an unchanged commute the average vendor would pay up to 4.5 per cent of monthly sales for a zone, and 68 per cent accept one when the spot is free. Distance then consumes that value, and unevenly. The acceptable rent falls to 2.1 per cent of sales one kilometre further out and to nothing at about 3.8 kilometres, with more than half the value going in the first kilometre alone. Rent and information shift take-up less than distance does over the moves a zone is likely to require. Zones sit empty not because the state has failed to offer formality or because vendors do not want it, but because conferring it requires moving the firm, and moving it costs most of what it is worth.

JEL Classification: O17, R38, C25, O18, C93
Keywords: Informality, Firm Location, Relocation costs, Vending Zones, Stated Preferences

*Department of Economics, London School of Economics and Political Science
Any correspondence/feedback regarding the draft can be sent to l.singh7@lse.ac.uk. We thank Steve Machin, Swati Dhingra, and Tim Besley for their valuable feedback, and participants at seminars at the London School of Economics and Durham University and at the Royal Economic Society Annual Conference for helpful comments. We gratefully acknowledge financial support from the British Academy/Leverhulme Small Research Grants scheme, the International Growth Centre (IGC) Small Grants programme, and the HER grant from STICERD at the London School of Economics. We are

1 Introduction

Around 2 billion people, or 60 per cent of the global workforce, work in the informal sector (ILO 2018). Governments in developing countries commonly attribute low fiscal capacity and low aggregate productivity to persistent informality, and have introduced policies that grant informal firms legal rights and social protection upon registration. Registration rates nonetheless remain low, and how much a given policy achieves depends on how it is designed and on how much the firms it targets want what it offers.

One sub-group of informal firms sits awkwardly in this framework. Street vendors are own-account firms that sell goods and services from carts, stalls, and ground displays in public spaces to which they have no legal right; the International Labour Organization's definition covers both those trading from a fixed pitch and those trading while mobile (ILO 2018). Street vending accounts for 1.5 to 6.8 per cent of total employment in low- and middle-income countries (Roever and Skinner 2016). For these firms a government cannot confer legal status without also deciding where they may lawfully trade. Formality is therefore offered at a location rather than at a price, and accepting it means relocating the firm.

Where vendors trade is contested. They concentrate on main roads and pavements claimed simultaneously by vehicles and pedestrians, and over the past decade governments have legislated intensively to ease congestion, keep public spaces usable, and bring vendors onto a register. The debate has produced forceful policy, including Accra's decongestion by-laws, Bangkok's Return the Footpath to Pedestrians, Ho Chi Minh City's Civilized Streets, and Johannesburg's Operation Clean Sweep. Appendix A1 documents 54 such policies across 27 countries and six regions, of which just over 70 per cent designate or construct a physical site to which vendors are expected to move: a demarcated, legally protected area in which certified vendors receive fixed-size spots for a monthly rent. Relocating to that site is the condition on which the protected spot is granted, and it is a condition the standard instruments of formalisation policy do not address.

Those instruments have been studied closely. The administrative and monetary costs of becoming and remaining formal are well documented, and the weight of the evidence is that removing the fixed costs of registration raises formality by less than reducing the recurring costs of formality or raising the cost of informality through enforcement (Bruhn and McKenzie 2014; Rocha, Ulyssea, and Rachter 2018; Farias and Hsu Rocha 2026; Andrade, Bruhn, and McKenzie 2016; Ulyssea 2020; Campos, Goldstein, and McKenzie 2023), though which margin binds remains contested. Far less attention has been paid to the relocation costs that formality can carry, or to firms operating on public land. Where land appears at all, it appears as an eligibility condition rather than as a price. de Mel, McKenzie, and Woodruff (2013) offered informal firms in Sri Lanka payments worth up to two months of median profits to register, and firms operating on public land were 27 percentage points less likely to register, because they could not produce the paperwork establishing a claim to the premises they occupied. Firms screened out on documentation

grateful to the Patna Municipal Corporation (PMC) for granting approval to conduct this fieldwork, and to Shri Animesh Kumar Parashar, IAS, Divisional Commissioner, for his support in facilitating this research. The authors declare no other financial relationships or conflicts of interest relevant to this research. We acknowledge the use of Claude in writing for improving clarity of the paper.

never face a choice, so what formality is worth to them, and what they would give up for it, is never observed.

This paper measures the relocation costs that formalisation imposes on firms without premises, and the preferences of the firms the policy is aimed at. We separately identify and value four potential barriers to relocation: location and commuting costs, rent, coordination failure among vendors, and limited information about the policy. We thereby identify the binding constraint and show what scope redesign has. In doing so we create new and detailed data on a group of informal firms about which little is known.

India's Street Vendors (Protection of Livelihood and Regulation of Street Vending) Act, 2014 provides a useful setting. It requires every urban local body to survey and certify vendors and establish vending zones. Certification is free, issued upon survey, and does not require proof of a claim to the vendor's current site, so the Act removes the fixed registration costs on which much of the formalisation literature turns. What certification does not deliver, in practice, is protection where the vendor already stands: a fifth of our sample holds a certificate of vending, yet certified vendors are no less likely than uncertified ones to report harassment or eviction in the preceding six months. The right the Act effectively confers is a legally protected spot inside a zone, and that right is conditioned on location, so low take-up points to the terms of relocation rather than to the cost of registering. We study the Act's implementation in Patna, the capital of Bihar, where approximately 30,000 vendors operate in public spaces and the Municipal Corporation has built several vending zones, all but one of which have zero occupancy.

That near-total non-relocation is why observed behaviour cannot reveal how vendors value the terms on offer. Vendors rarely change marketplaces; the zones are too few and too empty for their attributes to vary in the field; and the share of peers who relocate is jointly determined with each vendor's own decision. We therefore use a stated-preference experiment, which recovers valuations of attributes that do not vary in the field (McFadden 1974). We conducted an original geolocated survey of 2,922 vendors across 206 marketplaces in Patna, recording each vendor's location, rent, commute, and trading conditions. Each vendor then chooses five times between remaining at the current site and moving to a hypothetical vending zone whose distance, rent and expected trading conditions are defined relative to that verified status quo, and half the vendors in every marketplace receive information about the Act and the credibility of its enforcement between rounds. Section 3 describes the design, the discrete-choice model, and our measures to limit hypothetical bias.

Vendors value formality. A municipality designating an existing site as a zone, so that no commute changes, could charge the average vendor up to 4.5 per cent of monthly sales for the protected spot, and 68 per cent of vendors accept such a zone when the spot is free. What they value is protection rather than productivity, which accords with evidence that registration raises neither earnings nor productivity for most firms and that its benefits arise from reduced risk (de Mel, McKenzie, and Woodruff 2013; Benhassine et al. 2018; Farias and Hsu Rocha 2026; Ulyssea 2020).

Relocation then consumes that value, and it does so unevenly. The maximum rent the average vendor would accept falls to 2.1 per cent of sales after a one-kilometre longer

commute and reaches zero at about 3.8 kilometres, beyond which no rent is acceptable at all. More than half of what the legal right is worth is lost to the first kilometre alone, and the two kilometres after it together cost less than that first one. The penalty is specific to location: no comparable penalty attaches to a change in the rent or in the trading conditions a zone offers, the number of spots being the one exception, and its step is smaller than distance's. It is what the vendor gives up by leaving the site, since the customer base does not move with the firm.

That average conceals most of the variation across vendors, and what it conceals is not something a municipality could anticipate. Once the 27 vendor characteristics we observe are in the model, allowing vendors to differ is worth a further 760 held-out log points against about 29 from the characteristics themselves. Grouping vendors by what those characteristics predict spreads take-up at the best design across a range of 36 points, against 92 points when vendors are grouped by their own estimated valuation. Part of the variation is therefore visible to a survey, and that part is large in take-up terms, but it is much the smaller part of the total, so a programme cannot be aimed with any precision at the vendors most likely to take it up. What a municipality can vary instead is the terms of the offer.

The three remaining explanations account for much less. Rent is a real instrument but a bounded one: it carries disutility from the first percentage point of sales and bites most steeply between one and two per cent, and charging two per cent at an unchanged commute costs eleven points of take-up, against six points for moving the zone one kilometre and twenty-two for moving it five. Over the range a municipality would realistically choose between, rent and distance are comparable instruments; it is the longer moves that distance dominates. We find no detectable effect of randomised variation in the expected share of relocating peers, the explanation most commonly advanced for unoccupied formal space (Cooper and John 1988; Brock and Durlauf 2001). Information raises what a zone is worth without relaxing the relocation constraint, lifting take-up by a little over three points at every design we evaluate. And the number of spots a zone offers does not attract vendors, so building larger zones is not a remedy. One margin we bound rather than price: the spot dimensions fixed by the Bihar rules are smaller than the space most vendors appear to occupy, and because our choice tasks hold operating scale fixed, stated willingness to move assumes an accommodation the zones as built cannot provide.

These results contribute to the formalisation and relocation literatures. To our knowledge, this is the first paper to value the barriers to formalisation facing firms without premises on a common scale, each vendor's monthly sales, allowing the barriers to be ranked rather than merely listed. The ranking is not the one the policy debate assumes: distance first, rent second, information third, and coordination indistinguishable from zero. The paper also prices location, which the formalisation literature generally treats only as an eligibility condition. Because the Act removes the documentary requirements that would otherwise exclude firms without premises, we can directly estimate the cost of exchanging location for legal status, and we observe the preferences of firms that such requirements would ordinarily screen out before any choice is made.

The relocation literature has primarily studied households. Evaluations of slum relocation find that destination location matters as much as housing quality and that a better

dwelling in a worse location may leave households no better off (Takeuchi, Cropper, and Bento 2008; Barnhardt, Field, and Pande 2017), a conclusion the slum-upgrading literature reaches from the other direction by showing how much of the long-run gain from secure tenure depends on where it is conferred (Harari and Wong 2026). The evidence is mixed, however: in Addis Ababa, subsidised relocation to peripheral housing had high take-up and no adverse labour-market effects (Franklin 2025; Agness and Getahun 2024). Studies using discrete-choice methods to identify what relocation sites must offer have relied on moves that occurred (Takeuchi, Cropper, and Bento 2008; Lall, Lundberg, and Shalizi 2008). For households, location is partly a consumption choice, and movers can change employers to offset some losses. For firms, location is an input into production, and no comparable substitution is available. By experimentally assigning destination attributes rather than observing them, we recover the shape of the relocation cost and not only its level, and the shape is what matters for policy: the cost is front-loaded, so the greater part of it is attached to moving at all and can be avoided but not compensated, while only the remainder, the later kilometres, can be bought back.

The results delimit what policy design can achieve. Where an existing vending site can itself be designated as a zone, vendors receive the protected spot without incurring the cost of moving, and our estimates support this as close to the only viable option rather than merely the cheapest. It is also the option the municipality has least room to choose, since the sites vendors occupy are the sites whose congestion the policy exists to relieve. Where relocation is unavoidable, the scope for buying the distance back through rent is narrow, and narrowest over exactly the moves a municipality is most likely to need. Zones sit empty not because the state has failed to offer formality or because vendors do not want it, but because conferring it on a firm without premises requires moving the firm, and moving it costs most of what the legal status is worth.

The remainder of the paper is organised as follows. Section 2 describes the institutional setting, sampling frame, and survey. Section 3 presents the stated-preference design, discrete-choice model, and information experiment. Section 4 reports the results, and Section 5 discusses their implications and limitations.

2 Data

2.1 Institutional Context

In India, street vending is governed by the Street Vendors (Protection of Livelihood and Regulation of Street Vending) Act, 2014 (SVA henceforth). The Act requires every urban local body to survey the street vendors in its jurisdiction, issue certificates of vending, and designate vending zones to which vendors are to be allocated; state governments frame the rules and schemes under which those obligations are discharged. Vendors accommodated in these zones receive legal recognition and protection from eviction, rights that are absent in their current informal locations, and recognition under the SVA also makes them eligible for complementary government schemes, most notably the Pradhan Mantri Street Vendor's AtmaNirbhar Nidhi (PM SVANidhi) scheme, which provides subsidised collateral-free loans. Implementation is overseen by Town Vending Committees (TVCs) at the municipal level, with state governments retaining flexibility over licensing terms, fee structures, and renewal conditions. This decentralised approach has produced

substantial variation in implementation across cities.

The Act's core mechanism is spatial, and it is worth being precise about which costs it removes and which it relocates. Certification carries no fee, and the vendor files nothing: the municipality conducts the survey at the marketplace and registers vendors there. Nor does eligibility rest on documenting a claim to the site occupied, a requirement that elsewhere excludes firms trading on public land. The fixed cost of entering the formal sector, on which a substantial part of the formalisation literature has focused, therefore falls on the state rather than on the firm. What remain are recurring obligations: the scheme framed under the Act provides for a monthly charge for a spot inside a vending zone, and a certificate of vending must be reactivated at five-yearly intervals against a small fee (India 2014). Legal status is thus conditioned on neither payment at entry nor documented tenure, but on where the vendor trades.

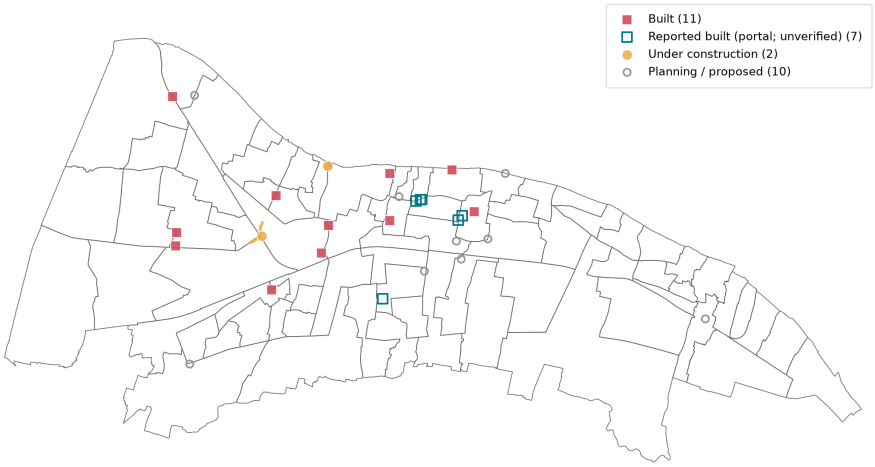
Where those zones are sited is not a residual administrative decision. The stated motivation of policymakers in relocating vendors is to relieve the congestion of carriageways and pedestrian walkways, so the sites on which vendors concentrate are precisely the sites a municipality is seeking to clear. The same density constrains the alternatives, since a zone requires a plot of a given size and few such plots are available in the built-up centre of a congested city. Both considerations push the location of a zone away from the corridors on which vendors currently trade, and a municipality implementing the Act is correspondingly constrained in the offer it can make.

We selected Patna, the capital city of Bihar, as the study site for three reasons. First, its entire urban area, spanning 75 wards, is governed by a single statutory authority, the Patna Municipal Corporation (PMC). All street vendor regulation, including licensing, enforcement, and policy implementation, is handled through one set of rules. This unified jurisdiction eliminates the confounding effects of inter-agency variation present in cities like Delhi, where multiple municipal corporations operate with potentially different enforcement practices, and simplifies both data collection and the interpretation of vendor behaviour.

Second, Patna's street vending landscape is shaped by commuter-led demand rather than by tourism. Vendors serve a stable local customer base of office workers, daily-wage earners, and households, selling necessities, predominantly produce and food for immediate consumption, priced for the city's purchasing power, and their stalls are long-established and embedded in kin-based social networks. This low-churn setting is well suited to studying the expectations and social ties that shape relocation decisions, without the confounding variation a transient vendor population would introduce.

Third, the SVA's implementation failures are particularly visible in Patna. Bihar is among India's most populous and least developed states, and the street vending sector in Patna is large, at approximately 30,000 vendors according to PMC records. Under the SVA, PMC has constructed multiple vending zones, yet implementation has stalled: all but one record zero occupancy, and the exception is only a quarter utilised. Figure 1 shows their locations. The zone designs also fix the space allocated to each vendor spot. Under the Bihar implementation of the Act, each spot measures either 6 by 4 feet or 5 by 3 feet, a dimension we return to in Section 4.8.

Figure 1: Proposed and Constructed Vending Zones across the City



Note: Officially recorded street vending zones in Patna, by construction stage. Thirty-three appear in a government record: a SULM portal entry, the 2021 Urban Development & Housing Department sanction (Memo 11913/2021), or a record of construction. Thirty are plotted, of which 11 are built, 2 are under construction and 10 are in planning; a further 7 are reported built on portal entries whose status field and status text disagree. The remaining three could not be located. Coordinates come from the SULM portal (10) and geocoded landmarks (20). One zone is designated as a stretch of road rather than a point and is drawn along the road.

2.2 Primary Survey

Data collection proceeded in three stages: marketplace listing, vendor recruitment, and the main survey. Patna is administratively divided into 75 wards. We selected 40 wards as the focus of the project based on vendor concentration, using the 2021 National Urban Livelihoods Mission (NULM) survey of vendors obtained from PMC. We geocoded potential vending locations, constructed vending zones, and overlaid these with ward boundaries to identify the 40 wards with the highest concentration of street vendor activity.

Marketplace Listing. We conducted a comprehensive listing exercise of marketplaces in these 40 wards in August 2024. The 40 study wards contain approximately 17,000 static vendors, the substantial majority of Patna's stationary vendor population. During the listing exercise, we recorded the geolocation of each marketplace, characteristics of the surrounding area, and an approximate count of vendors operating in each market. We define a marketplace as any public space or location where more than 10 vendors are clustered selling goods or services. We define a street vendor as a person, typically self-employed, selling goods or services and operating their business from a temporary built-up structure at a fixed location in public spaces such as streets, lanes, sidewalks, footpaths, pavements, public parks, or any other public place. Importantly, we exclude hawkers and mobile vendors from this definition, focusing instead on vendors with fixed locations. This is narrower than the ILO definition given in Section 1, which admits mobile vendors. The restriction follows from the experiment: the choice tasks price attributes of a fixed location, and rent, number of spots and commuting distance have no counterpart for a vendor who does not occupy one. Our project specifically targets produce and street food vendors, who together comprise 66% of the street vending sector according to EBRD informal enterprise surveys. The listing exercise identified 375 marketplaces across the 40 wards, which were reduced to 206 analytical marketplaces in two

steps. First, 43 marketplaces were excluded prior to the survey: their primary product was not produce or street food, and one market had been declared a commercial zone on a first-come-first-served basis, leaving 332 shortlisted marketplaces. Second, 126 of these were merged into nearby markets to protect the information experiment. Because treatment is assigned within marketplaces, contamination is governed by the time between a treated vendor being surveyed and a control vendor nearby being surveyed. Merging adjacent clusters into a single marketplace meant each was surveyed within a single day, leaving hours rather than days for information to travel; had they been retained as separate units, they would have been visited on different days, and vendors in the later cluster could have heard about the Act from those surveyed earlier.

Recruitment and Survey Administration We returned to the 206 identified marketplaces between June and October 2025 to conduct surveys. The survey process proceeded in two stages. First, we conducted recruitment visits during peak marketplace hours to identify eligible vendors and schedule survey appointments at times when vendors were relatively free to participate. We employed a systematic random walk recruitment strategy, approaching every third vendor in each marketplace who fit our project focus (produce or street food vendors operating from fixed locations). We recruited 3,499 vendors.

Second, we administered the main survey to recruited vendors within a few days of the recruitment. Surveys were conducted in-person and typically lasted 45 minutes. Out of those recruited vendors who agreed to participate, 2,935 participated in the main survey yielding an initial response rate of 84%. We exclude 13 observations with incomplete responses, leaving a final analytical sample of 2,922 vendors across 206 marketplaces. Descriptive statistics in this section are reported on that sample. The discrete-choice estimation in Section 4 uses 2,910 of them across 203 marketplaces, twelve vendors being dropped for a missing covariate and taking three marketplaces with them (Section 3.2); no descriptive reported here differs by more than a tenth of a percentage point between the two. Survey respondents are broadly similar to survey non-respondents (Appendix A2).

Figure 2, panel (a) maps the geo-locations of all 2,922 vendors in the sample across Patna, revealing clusters distributed across the city. The spatial distribution of vendors reveals strong clustering along transportation corridors. Panel (b) maps vendors' reported residences, geocoded from the landmarks they named to the centroid of the surrounding locality. Residences retain a dense core in the central wards but extend substantially further east, indicating that a sizeable share of vendors commute inward to their stalls. Panel (c) shows the markets within the city that vendors name as their source of supply. Including 14 markets that are outside city boundaries, wholesale (produce) markets sit largely on the periphery of the built-up area, while ingredient markets (retail markets) are concentrated in the centre alongside the densest vendor clusters. This implies that bulk produce enters the city at its edges and is being redistributed inward. 89% of vendors live within 3 km of their place of business (Figure 3 (a)) and 91% of them report that they don't pay any fees for occupying their current business spot (Figure 3 (b)).

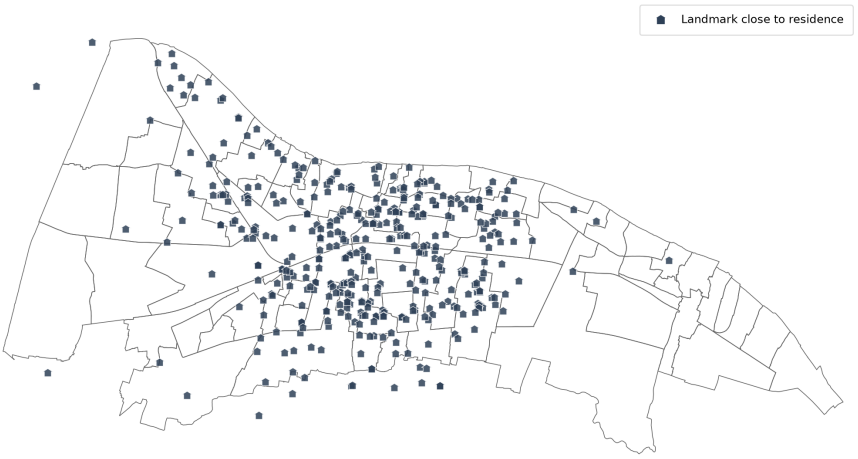
Table 1 presents descriptive statistics for the analytical sample. Vendors are on average 39 years old with 6 years of education; 88% are male and 96% belong to historically disadvantaged caste groups (87% Other Backward Castes, 10% Scheduled Caste/Tribe).

Figure 2: Geo-location of vendors, their homes and sourcing markets

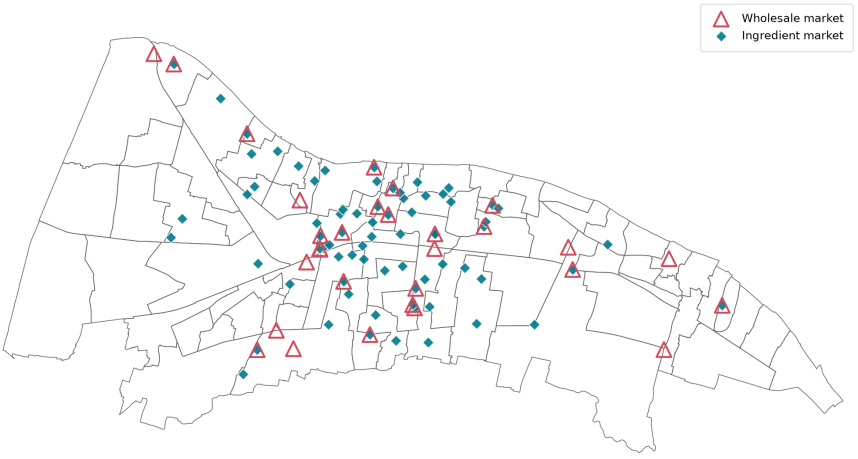
(a) Vendors' Business



(b) Vendors' Residence



(c) Wholesale and Inventory Markets

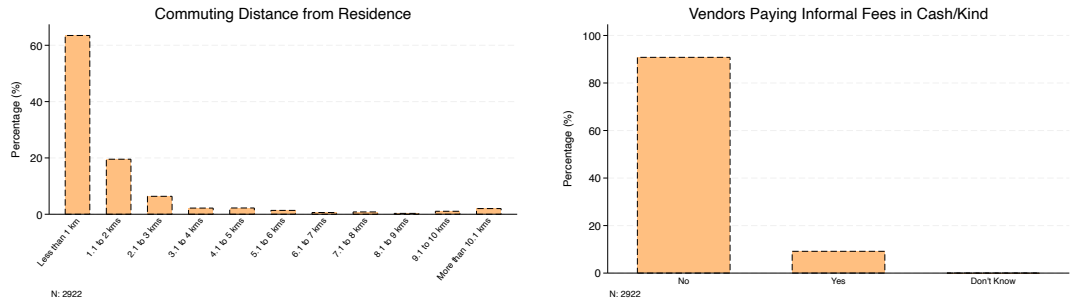


Note: (a) Surveyed vendors across the 40 study wards. (b) Each marker is a distinct landmark reported as near a vendor's home, geocoded to its locality centroid; 412 of the 719 landmarks reported were geocoded confidently and fall within the mapped area. Residence locations are available for 2,374 of 2,922 vendors (81.2 per cent), plus 118 who live near their stall and take its coordinates; the remainder gave no answer or a landmark that could not be tied to a place name. (c) Markers are distinct markets named by vendors: outlined triangles are wholesale (produce) markets, filled diamonds ingredient markets. Fourteen markets outside the mapped area are included in all counts but not drawn. Market locations are available for 1,805 vendors (61.8 per cent), covering 84.0 per cent of fruit and 79.5 per cent of vegetable vendors but only 32.9 per cent of cooked-food vendors, who mostly source from neighbouring shops or take deliveries.

Figure 3: Status-quo Environment

(a) Reported Commuting Distance between Residence and Marketplace

(b) Share of Vendors Paying Informal Fees at Current Business Location



For most individuals, street vending is not transitional work: 53% have been vending for over 10 years, and 95% work 7 days per week averaging 9 hours per day. Average daily profit is INR 619 with average daily sales of INR 2,991. 85% of our sample operates from stalls or carts, and 90% trade along roads or streets with the remaining 10% on pedestrian walkways, leaving them highly exposed to eviction.

Figure 4 describes the licensing and enforcement environment in which these vendors operate. Panel (a) shows that 20% hold a certificate of vending, so certification under the SVA has reached a fifth of the sample. Panel (b) reports how often eviction drives (*halla gadis*) occur in the vendor's marketplace: 10% report that they never do, while 70% report at least one a month and 23% at least one a week. Panels (c) and (d) record personal exposure over the six months before the survey, with 41% reporting harassment from officials or the police and 33% reporting an eviction or confiscation of goods. Enforcement is routine rather than exceptional, and it reaches vendors directly rather than through fines or paperwork.

Set against this, awareness of the policy that would protect them is close to zero: only 1.3% could identify the Street Vendors Act (Figure 5). Only 7.7% of vendors were aware of any policy concerning street vending, and 6.3% knew the name of any policy. This combination, active enforcement pressure alongside near-zero policy awareness, motivates both the study setting and the information experiment described in Section 3.3.

3 Empirical Strategy

Pricing the relocation the Act requires means recovering how vendors value the terms on which formality is offered: the rent charged, the distance from their residence, the number of spots, the competitive and demand conditions they expect to find, and the share of peers who also move. Three features of the setting prevent us from recovering these valuations from observed behaviour. First, 99% of vendors have never changed marketplace, so revealed preference data contain almost no variation in location choice. Second, the terms do not vary in the field: the zones PMC has constructed are few and almost entirely unoccupied (Section 2.1), so there is neither a spread of offers to compare nor take-up against which to price them. Third, the peer share is simultaneously determined, since the individual relocation decision depends on the share s of peers relocating while s is itself the aggregate of those decisions. Observational estimates would in any case

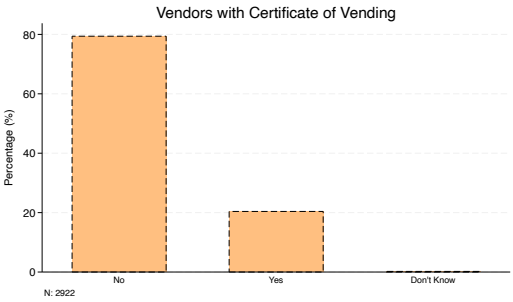
Table 1: Summary Statistics

Variable	Mean	SD
Age	38.64	12.76
Education	6.08	4.76
Male	0.88	0.32
Marital Status		
Married	0.82	0.38
Single	0.16	0.36
Widowed/Divorced	0.02	0.15
Caste Group		
Backward Caste	0.86	0.34
Scheduled Caste/Tribe	0.10	0.30
General	0.04	0.19
Product Category		
Fresh Produce	0.59	0.49
Cooked Food/Drinks	0.41	0.49
Premises Type		
On Main Roads or streets	0.90	0.30
On Pedestrian Walkways	0.10	0.30
Infrastructure Type		
Have Stalls/Carts	0.85	0.36
On the Floor	0.15	0.36
Other Business related		
Work 7 days a week	0.95	0.22
Hours of Operation	9.22	3.67
Have employees	0.32	0.46
Yesterday's Profits (INR)	618.53	626.31
Weekly Profits (INR)	4,259.65	4,056.31
Monthly Profits (INR)	18,458.50	17,577.33
Own Other Businesses	0.15	0.35
Never Changed Marketplaces	0.99	0.08
Years in Business		
Upto 1 year	0.06	0.24
1 to 5 years	0.21	0.41
5 to 10 years	0.20	0.40
More than 10 years	0.53	0.50

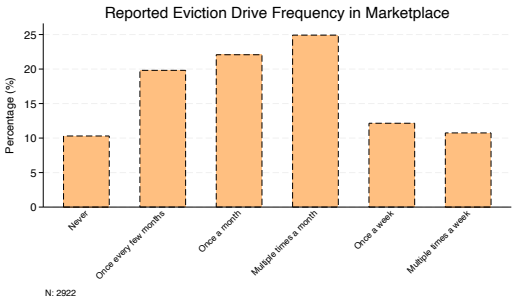
Note: N = 2922. All monetary units are expressed in Indian Rupees (INR). Monthly profits are constructed from reported weekly profits as weekly profits $\times$ 52/12. Profits were elicited for the previous day and the previous week only, since most vendors do not keep written accounts and longer recall windows are unreliable.

Figure 4: Licensing and Enforcement Environment

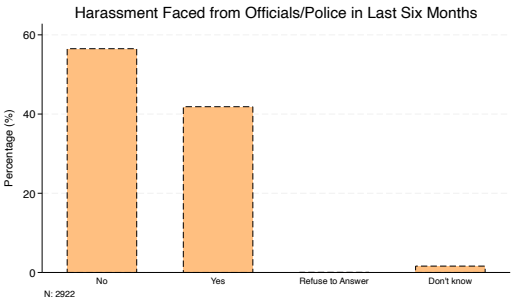
(a) Certificate of Vending



(b) Eviction Drives Frequency in Marketplace



(c) Harassment from Officials or Police



(d) Eviction or Confiscation of Goods

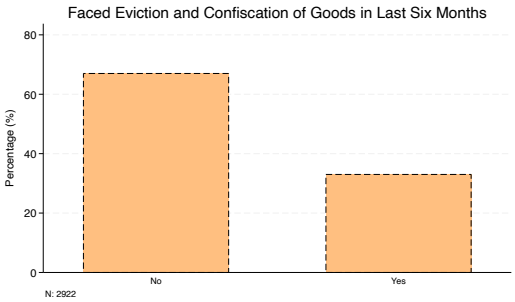
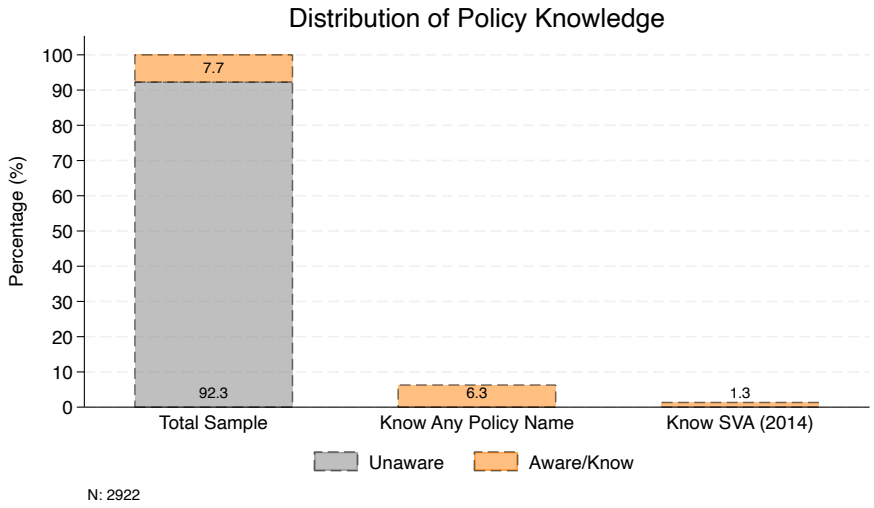


Figure 5: Awareness of Policies Among Street Vendors



confound the causal effect of peer behaviour with common unobservables driving both.

We address all three by assigning the terms of the offer rather than observing them. A stated preferences experiment presents each vendor with a sequence of hypothetical zones whose attributes are drawn from a designed distribution, so that rent, distance and the peer share vary independently of anything about the vendor. The approach is well established in transportation and environmental economics and has been applied to labour market choice by Mas and Pallais (2017). Two features of our implementation are intended to limit hypothetical bias: the scenarios are personalised, with every attribute expressed relative to each vendor's verified current conditions, which anchors the choices to circumstances the respondent already knows (Taylor and Cummings 1999); and vendors were told at the outset that their responses would inform PMC policy. Appendix A5.2 sets out these and the other design responses to hypothetical bias, inattention and fatigue. A randomised information treatment, described in Section 3.3, is embedded between choice rounds.

3.1 Stated Preferences Experiment

Cognitive Testing: The vignette was tested with approximately 50 street vendors in Patna before it was finalised. Given that many street vendors have low levels of literacy, the format had to be accessible to respondents regardless of their reading ability. We presented draft versions of the choice tasks and solicited feedback on the framing of the question, the language used to describe attribute levels, and the visual presentation of the two market alternatives. On that basis we iteratively refined the design, incorporating visual aids and images to represent attributes and to facilitate comparison between the status quo and the alternative. Appendix Figure A2 presents the final format.

Experiment Design: Each respondent completed six choice tasks. The first task served as a practice round (round 0) to familiarise participants with the choice format and has been excluded from the analysis. The remaining five rounds (rounds 1–5) constitute our analytical sample. In each choice task, respondents chose between two alternatives: (1) remain in their current location, or (2) relocate to a vending zone. Both alternatives were described by the same six attributes, whose values for the zone varied systematically across choice occasions, with the exception that the zone consistently offered legal rights to operate, the feature distinguishing formal from informal vending.

Vendors are told that the new marketplace is identical to their current one except in the attributes shown, and that storage capacity and transport costs are unchanged. This holds fixed the supply-side conditions that would otherwise confound valuations of the offered terms, so that differences across profiles operate through the attributes we vary. It also means the experiment holds each vendor's scale of operation fixed by construction, a limitation we return to in Section 4.8. Appendix A5.3 reproduces the enumerator script.

Attribute Selection and Levels: The six attributes were selected based on formative qualitative interviews and existing literature on informal sector formalisation. Table 2 presents the full set of attributes and their possible values in the vending zone alternative. Attribute levels were determined through pilot testing to reflect realistic variation observed in actual vending markets while ensuring sufficient spread to identify prefer-

ences. Appendix A5 details how the scenarios for the experiment were selected.

Table 2: Possible Values of Attributes in the New Marketplace

Attributes	Levels
Monthly Rent (% of monthly sales)	0, 1, 2, 4, 8
Number of spots	same, 20, 40, 80, 160
Expected number of vendors relocating (% of vendors)	[0, 100]
Expected number of competitors (absolute change)	-10, -5, 0, +5, +10
Expected number of daily customers (% change)	-50, -20, 0, 20, 50 if current ∈ [10-60]
	-30, -15, 0, 15, 30 if current ∈ [61-120]
	-25, -10, 0, 10, 25 if current ∈ [121-170]
Distance from residence (absolute change)	-1, 0, +1, +3, +5

A key design feature is that attribute values for the alternative zone were defined *relative to each vendor's self-reported status quo*, elicited in the survey module preceding the choice tasks. For example, if a vendor reported a current commuting distance of 2 km and the design specified +1 km, the alternative was presented as a 3 km commute. All vendors therefore face the same experimental variation in relative terms while the absolute values are grounded in their actual situation. All pairwise attribute correlations in the design matrix are below 0.2, so individual attribute effects can be estimated without multicollinearity concerns (Appendix A5).

Two features of the assignment bear on what the design can identify, and we state them here rather than alongside the results. Reductions in commuting distance were floored at each vendor's existing commute, so an offer of a one-kilometre shorter commute frequently delivered no change at all; the experiment therefore identifies the cost of a longer commute but not the value of a shorter one. The assigned peer share was constructed by shifting each vendor's own elicited belief by a randomised shock of -50, -20, 0, $+20$ or $+50$ percentage points, rounding to the nearest ten points because the card presents the share as a count out of ten vendors, and censoring at zero and 100; the range of exogenous variation actually delivered is therefore narrower than the range the design nominally spans. Section 4.7 reports both.

3.2 Discrete Choice Model

Vendor i chooses in round t between remaining at the current site and relocating to the offered zone, and relocates if the net utility gain is positive:

$$U_{it}^{\text{relocate}} - U_{it}^{\text{stay}} = V_{it} + \varepsilon_{it} > 0, \tag{1}$$

with ε_{it} logistic and independent across i and t. The deterministic component is

$$V_{it} = \bar{\alpha} + z_i'\gamma + a_i + x_{it}'\beta + d_i \cdot \text{km}_{it}, \tag{2}$$

so that the model is a binary panel logit with random coefficients. The vector x_{it} collects the attributes of the offer, which vary across occasions by design, together with round

dummies and the information treatment interacted with post-treatment rounds. The vector z_i collects 27 vendor characteristics measured before the choice tasks, and each vendor's elicited belief about the share of peers who would relocate.

Functional form is tested, not imposed. Every attribute takes a small number of discrete levels, so each can enter as a full set of level dummies, a specification that imposes no shape on any attribute and only additivity across them. Simpler forms are restrictions of that saturated model, tested by nested likelihood-ratio tests. Two restrictions survive and the rest do not. Commuting distance enters as free dummies at the offered levels,

$$\sum_{\ell \in \{-1,\,+1,\,+3,\,+5\}} \beta^D_\ell \cdot \mathbf{1}\{\Delta\text{Distance}_{it} = \ell\},$$

because no parametric restriction on it is accepted: the penalty at one kilometre is far larger than linearity permits. Rent, assigned as a percentage r_{it} of each vendor's own monthly sales, enters as a linear spline with knots at one and two per cent,

$$\beta^R_1 \min(r_{it}, 1) \;+\; \beta^R_2 \big(\min(r_{it}, 2) - 1\big)^+ \;+\; \beta^R_3 (r_{it} - 2)^+,$$

which is the free-dummy model subject to a single restriction the data do not reject. The remaining attributes, the number of spots, expected competitors, expected customers and the assigned peer share s, enter linearly, each having been tested against its own free dummies. Appendix A6 reports every test, including against the fully linear specification with which this project began, which the saturated model rejects. Additivity across attributes is maintained throughout and is not tested: the design was not constructed to identify interactions between distance and rent, so the quantities in (5) rest on it as an assumption.

Vendors differ. Two coefficients are random and estimated as independent normals. The vendor-specific intercept $a_i \sim \mathcal{N}(0, \sigma_a^2)$ is the private value of a zone, the part of it that observed characteristics do not explain. The vendor-specific slope $d_i \sim \mathcal{N}(0, \sigma_d^2)$ allows each additional kilometre of commute to cost vendors differently. The mean of d_i is zero by construction rather than by estimate: no continuous distance term enters the fixed part of (2), so the average response to distance is carried entirely by the free dummies β^D_ℓ, and σ_d measures dispersion around a mean response whose shape is unrestricted. Section 4.2 reports why both random coefficients are needed and Appendix A7 reports the alternatives, including a latent class model and non-parametric shapes for the distribution of a_i.

The likelihood integrates (2) over a_i and d_i on a product rule: a 241-node rectangle rule spanning six standard deviations for the intercept, crossed with twelve-point Gauss–Hermite for the distance slope. Gauss–Hermite is not used for the intercept because its dispersion is large relative to the logistic scale, $\hat{\sigma}_a \approx 4.9$, and a Hermite rule places too few nodes in the region that carries the mass. We confirm each reported log-likelihood on a much finer grid, where it moves by less than 0.01. Every model is started from three values of each standard deviation and all three reach the same optimum. Standard errors are clustered at the marketplace, the level at which the choice scenarios were assigned, and we report alongside them a 200-replication marketplace pairs bootstrap; the two imply the same inferences for every quantity in Section 4, with interval widths differing

by at most a third. The estimation sample is all five rounds for 2,910 vendors, 14,550 choice occasions across 203 marketplaces; the twelve vendors dropped from the surveyed sample are those missing a covariate in z_i, and they account for the three marketplaces that fall out with them.

Quantities. The implied choice probability for a vendor with given (a_i, d_i) is

$$P(\text{relocate}) = \frac{\exp(V_{it})}{1 + \exp(V_{it})}, \tag{3}$$

and predicted take-up for any zone design averages this over vendors and over the distribution of the random coefficients. For attributes other than rent we report trade-offs in kilometres of commute, using the per-kilometre rate between the +1 and +3 km levels,

$$\text{Trade-off}_k = \frac{\hat{\beta}_k}{\big(\hat{\beta}^D_{+3} - \hat{\beta}^D_{+1}\big)/2}, \tag{4}$$

which is a ratio of coefficients within one model and so is free of the scale normalisation. We report these alongside the same ratios from a fixed-effects conditional logit, which assumes nothing about how vendors differ.

Because rent is not linear below two per cent of sales, the cost of a move cannot be stated by dividing a distance coefficient by a rent slope. We report it instead as the maximum rent a zone can carry at each offered distance. Let $W_i(\ell) = \bar{\alpha} + z_i'\gamma + a_i + \beta^D_\ell + \tilde{x}_{it}'\tilde{\beta}$ denote utility at distance ℓ with every attribute other than rent at the vendor's status quo. The maximum acceptable rent $r^*_i(\ell)$ solves

$$\beta^R_1 \min(r, 1) + \beta^R_2 \big(\min(r, 2) - 1\big)^+ + \beta^R_3 (r - 2)^+ = -W_i(\ell), \tag{5}$$

which has a closed form within each segment of the spline. We report this for the average vendor and as a distribution across vendors, censored at eight per cent of sales because the design never offered more. The break-even distance, at which r^* reaches zero even at a free spot, is obtained by linear interpolation between adjacent offered levels and is reported as approximate for that reason.

Because the elicited peer belief enters z_i, the intercept is not itself the value of a zone and every valuation is evaluated at the sample mean of that belief, with the value at a belief of zero reported alongside. Section 4.4 explains why the belief is retained despite this: without it, the assigned peer share inherits the construction by which it was generated, and the coefficient used to bound coordination effects is no longer shock-identified.

3.3 Information Experiment

The awareness data reported in Section 2 motivate the information experiment. At the start of the survey, we randomly assigned vendors to treatment and control groups within each marketplace. After completing round 3 of the choice tasks, treatment group vendors received detailed information about the Street Vendors Act (2014) and concrete evidence of enforcement in Chandigarh, where approximately 3,000 non-compliant vendors were displaced in late 2024 (the full treatment text is reproduced in Appendix A5.4). The information was delivered orally by enumerators to accommodate varying literacy levels.

Control group vendors received no information and proceeded directly to round 4. All vendors then completed two further choice tasks (rounds 4–5) in the same format as previous rounds. Treatment and control groups are comparable on observable characteristics prior to information provision, with a marginal difference in weekly sales (Appendix Table A4).

To identify how enforcement information shifts the perceived value of the vending zone, we allow the treatment to shift the intercept of (2) alone, constraining attribute slopes to be common across treatment and control:

$$V_{it} = \bar{\alpha} + \underbrace{\lambda \cdot \text{treat}_i \cdot \text{post}_t}_{\text{treatment effect}} + \sum_{t=2}^{5} \rho_t \mathbf{1}\{\text{round} = t\} + z_i'\gamma + a_i + x_{it}'\beta + d_i \cdot \text{km}_{it}, \qquad (6)$$

where $\text{post}_t \in \{0, 1\}$ indicates rounds 4–5 and $\text{treat}_i \in \{0, 1\}$ indicates treatment assignment. The coefficient of interest is λ. The main effect of treat_i is not estimated separately: assignment was randomised within marketplace, so any time-invariant difference between the arms is absorbed by $z_i'\gamma$ and the random intercept a_i. The main effect of post_t is absorbed by the round dummies ρ_t. This specification embodies the conceptually correct null: if the treatment conveys information about whether the zone will function as advertised, it should raise the overall value of the zone without changing the relative weights vendors place on its design attributes. It is also a restriction we test rather than assume. In a conditional-logit specification of the same design, interacting the treatment with every design attribute does not reject the restriction that the treatment shifts the intercept alone ($\chi^2(8) = 5.82$, $p = 0.67$), and a further test allowing the treatment effect to differ across the range of estimated valuations does not reject a common effect ($p = 0.18$).

Because the treatment shifts only the intercept, all five rounds can be pooled in a single estimation, and (6) is the specification we report throughout Section 4 rather than a separate exercise. The pre-treatment rounds identify the design trade-offs and the post-treatment rounds identify λ, while the round dummies absorb learning, fatigue and any drift common to both arms. The treatment effect is converted to percentage points of monthly sales through (5) and to kilometres through (4).

4 Results

We present the results in the order in which the modelling choices were made. Section 4.1 shows what the choices look like with no functional form imposed on any attribute, and establishes that neither distance nor rent is linear. Section 4.2 shows that vendors differ far more in what a zone is worth to them than observed characteristics can account for. Section 4.3 reports the main estimation table, and the four subsections that follow report the quantities the paper turns on: what a zone is worth, what relocation costs, which vendors a zone reaches, and how much the three remaining explanations account for.

Willingness to pay is measured in percentage points of monthly sales, because the numéraire is the rent slope above two per cent of sales. The design assigned rent as a percentage of each vendor's own reported sales, so the recovered parameters are rates on own sales rather than rupee amounts for a representative vendor. Rupee figures throughout are conversions at median monthly sales in the sample, INR 60,667, so one percentage point

of sales is INR 607. Intervals are the 200-replication marketplace pairs bootstrap; the clustered delta-method intervals imply the same inferences throughout, with widths differing by at most a third.

Two distinct shares appear below and are not interchangeable. Predicted take-up averages each vendor's choice probability over the estimated distribution of the random coefficients, and is the quantity reported in Table 7 and used throughout the paper when a share of vendors is quoted. The share of vendors with a positive valuation counts the vendors for whom the zone is worth more than the status quo in the deterministic part of utility, with no logistic error, and is the quantity traced across rent levels in Figure 14. The second is the larger of the two wherever it exceeds a half, and we name which is which at each use.

4.1 Neither Distance Nor Rent Is Linear

Every attribute in the design takes a small number of discrete levels, assigned at random, so the relocation rate at each level can be read without a model at all. Figure 6 plots those raw shares for all six attributes, with marketplace-clustered confidence intervals. Figure 7 reports the commuting distance levels inside a fixed-effects logit with a free dummy at each offered level, which imposes neither a shape on distance nor a distribution on vendor intercepts. Both show the same two departures from linearity.

Commuting distance costs far more at the first kilometre than at later ones. The penalty attached to a one-kilometre longer commute is 0.43 of the penalty attached to a three-kilometre one (95% CI [0.31, 0.56]), where linearity would require one third and would place the ratio much closer to the interval's lower end only by coincidence. Rent is non-linear in the opposite direction: it is not free at low levels, as the earlier draft of this paper assumed, and it bites hardest in the middle of its range rather than at the top.

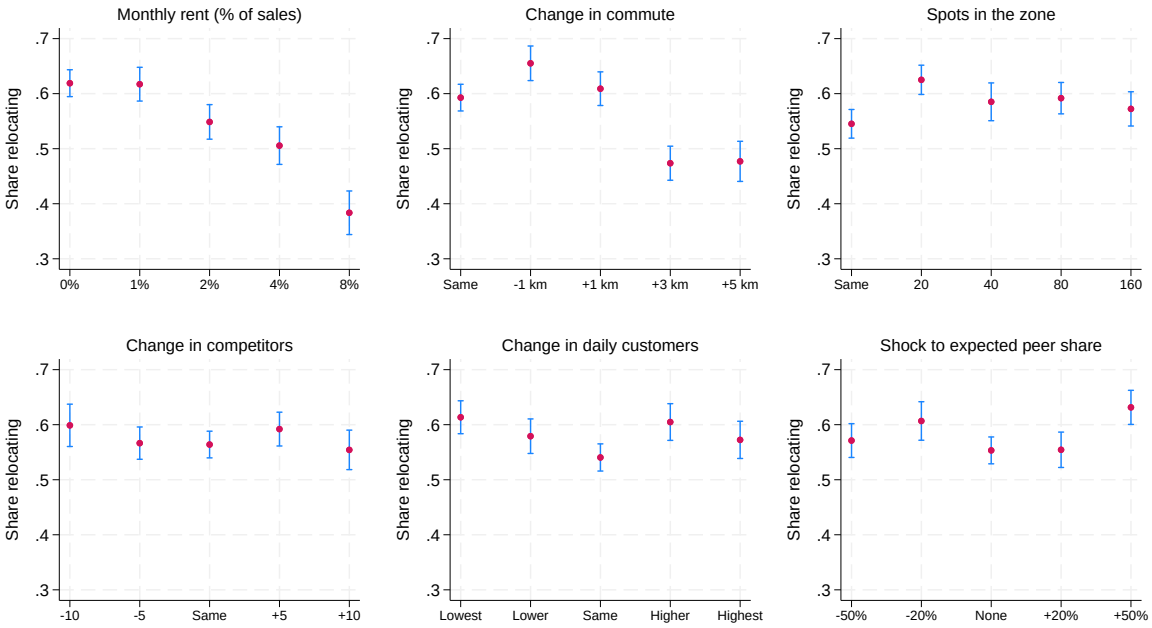
We therefore carry both attributes into the estimation in the least restrictive form the design supports. Distance enters as free dummies at each offered level. Rent enters as a linear spline with knots at one and two per cent of sales, which is the free-dummy model subject to a single restriction that the data do not come close to rejecting. Appendix A6 reports the restriction tests behind both choices, including the linear specification with which this project began and which the free-dummy model rejects. The rent ladder reported below assumes that distance and rent enter additively. The design was not constructed to identify their interaction, so we maintain additivity rather than test it.

4.2 Vendors Differ, and Observables Do Not Say How

A model in which all vendors share the same preferences cannot reproduce the data. Figure 8 plots the observed distribution of the number of rounds in which a vendor relocates, from zero to five, against the distribution a homogeneous conditional logit predicts. The data are concentrated at the two ends: 748 vendors never relocate and 1,021 always do, while the homogeneous model places far too many vendors in the middle.

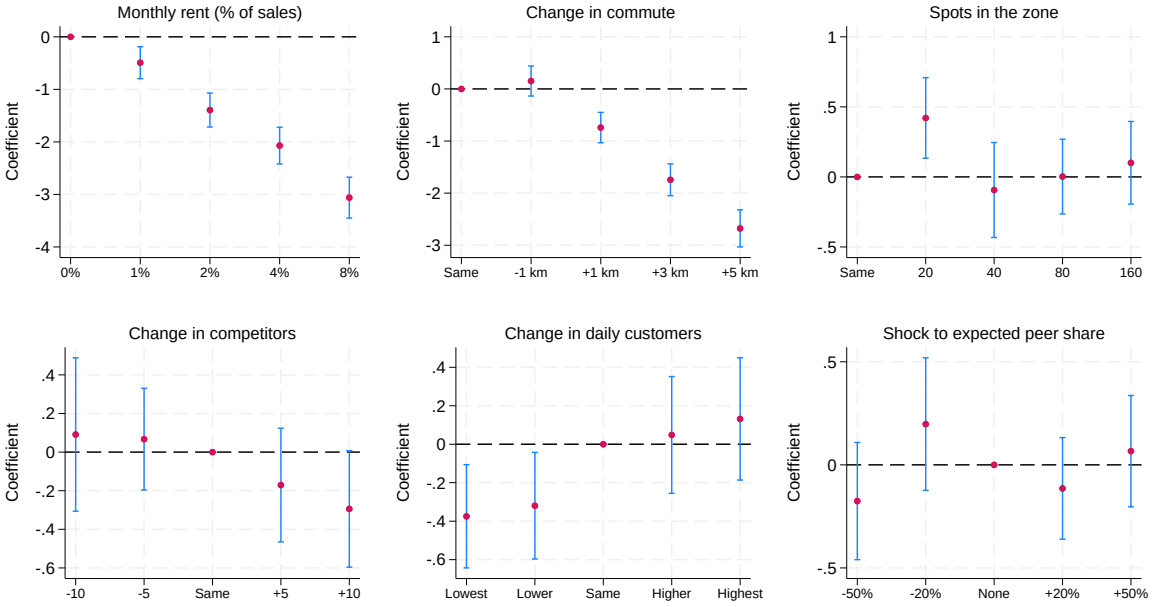
Allowing vendors to differ repairs this and is worth far more than the covariates. Table 3 reports held-out fit on the control group's final two rounds, with each vendor's posterior formed from their training rounds alone. Once the 27 vendor characteristics are already

Figure 6: Relocation Rate at Each Attribute Level



Raw shares over all five rounds; 95% CIs clustered by marketplace.

Figure 7: Commuting Distance, Free Level Dummies



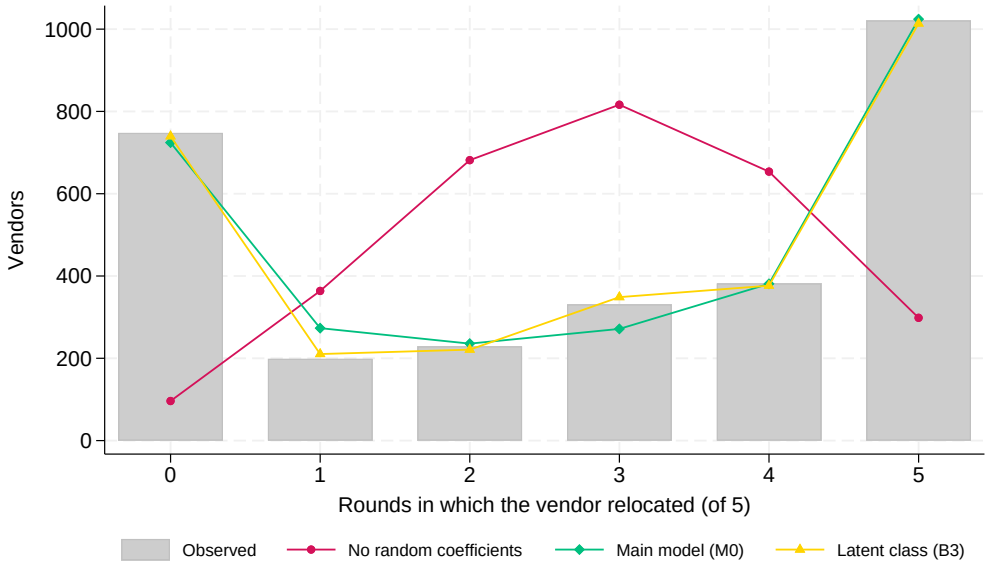
FE logit with vendor fixed effects and free level dummies;
95% CIs clustered by marketplace. Base: no change.

in the model, allowing vendors to differ is worth a further 760 held-out log points, against about 29 from the characteristics themselves.

The two variance figures in Table 4 measure different things and are reported together for that reason. The 27 characteristics generate a predicted spread in the value of a zone with a variance of 4.43, which is 15.5% of the total spread across vendors. Yet adding them reduces the *unexplained* variance by only 6.1%. Both are correct, and together they say that the covariates are largely orthogonal to the vendor-specific deviation they were meant to absorb: they shift where a vendor sits on average without explaining why that vendor differs from others like them.

We therefore estimate a mixed logit with two independent random coefficients, a vendor-specific valuation of the zone and a vendor-specific per-kilometre cost, specified in Section 3.2. Appendix A7 reports the alternatives we estimated against it, including a three-class latent class model and flexible non-parametric shapes for the distribution of valuations, none of which pays for its parameters.[1]

Figure 8: Observed and Predicted Distribution of Relocation Counts



4.3 Main Estimates

Table 5 reports four specifications. Column (1) is a conditional logit with every attribute entering linearly, the specification with which this project began. Column (2) replaces every attribute with free level dummies, and the likelihood-ratio test of (1) against (2) rejects linearity. Column (3) is a fixed-effects conditional logit in the adopted forms,

[1] One richer specification does. Adding a vendor-specific cost of changing site at all improves the likelihood by 11.3 log points for one parameter and is the only specification that also improves out of sample. It leaves the maximum rent at an unchanged commute, the break-even distance and the value of information unchanged to within one per cent, and moves the intermediate rungs of the ladder by at most eight per cent; its one substantive difference is the dispersion of the distance response at short moves, which the specification used here understates by a factor of about 2.5. Statements below about how widely vendors differ at short distances are therefore lower bounds. Appendix A7 reports it in full.

Table 3: Held-out Fit

	Log-likelihood	Brier	Hit rate
Conditional logit	-1827.3	0.224	0.638
Conditional logit + 27 covariates	-1798.8	0.219	0.647
Mixed logit + 27 covariates	-1037.9	0.110	0.854
Latent class, three types	-1063.0	0.114	0.847
Gain in held-out log-likelihood			
The 27 covariates alone	28.5		
Heterogeneity, on top of the covariates	760.8		
Both together	789.4		

Notes: Fit on the control group's final two rounds, which no model saw in estimation; each vendor's posterior is formed from their training rounds alone. Brier is the mean squared prediction error and the hit rate the share of choices called correctly, both on the same held-out rounds. The lower panel decomposes the gain over the bare conditional logit: the covariates are worth little, and almost all of the improvement comes from allowing vendors to differ once the covariates are already in.

Table 4: What the Observed Characteristics Explain

	Value
Fitted SD of the random intercept	4.91
SD of the covariate-predicted intercept	2.11
SD of the posterior intercept	4.20
Share of zone-value variation from covariates (%)	15.5
Fall in unexplained variance when covariates added (%)	6.1

Notes: All quantities are in the utility scale of the zone intercept. The two percentages measure different things and are reported together for that reason: the covariates generate a predicted spread that is 15.5 per cent of the total variation across vendors, yet adding them reduces the *unexplained* variance by only 6.1 per cent. They shift where a vendor sits on average without explaining why that vendor differs from others like them.

which imposes no assumption whatever on how vendors differ, at the cost of using only vendors whose choices vary across rounds. Column (4) is the main specification. The bottom panel reports the same trade-offs for every column in kilometres of commute, which is where the columns are comparable: coefficient levels in column (3) are on a different scale, because removing intercept heterogeneity inflates them, but ratios are not, and columns (3) and (4) agree on every one.[2] Covariate coefficients are in Appendix Table A3.

[2]Column (3) uses the 1,141 vendors whose choices vary across rounds, the only vendors that identify a fixed-effects logit; the other columns use all 2,910. Its log-likelihood is conditional and has no BIC.

Table 5: Main Estimates

	(1) Rent/dist linear	(2) Rent/dist free	(3) FE logit	(4) Main model (M0)
main				
Rent (% of sales)	-0.175***			
	(0.013)			
Rent, 0-1% (per point)				-0.581***
				(0.146)
Rent, 1-2% (per point)				-1.082***
				(0.146)
Rent, above 2% (per point)				-0.376***
				(0.034)
1%		-0.215***	-0.510***	
		(0.067)	(0.149)	
2%		-0.551***	-1.417***	
		(0.073)	(0.157)	
4%		-0.833***	-2.088***	
		(0.090)	(0.166)	
8%		-1.428***	-3.083***	
		(0.108)	(0.185)	
Change in commute (km)	-0.189***			
	(0.016)			
-1 km		0.053	0.265	0.506**
		(0.120)	(0.235)	(0.234)
+1 km		-0.326***	-0.760***	-0.900***
		(0.068)	(0.142)	(0.147)
+3 km		-0.773***	-1.766***	-2.074***
		(0.069)	(0.150)	(0.175)
+5 km		-0.993***	-2.705***	-3.393***
		(0.089)	(0.173)	(0.262)
Information arm	-0.009	-0.005		
	(0.069)	(0.069)		
Information x post	0.184***	0.185***	0.393**	0.506***
	(0.055)	(0.055)	(0.154)	(0.175)
SD: intercept				4.912***
				(0.201)
SD: per-km slope				0.725***
				(0.049)
Log likelihood	-8968.88	-8958.91	-1601.75	-6005.87
Parameters	59	65	29	53
BIC (on vendors)	18408.34	18436.26		12434.47
Vendors	2910	2910	1141	2910
Choice occasions	14550	14550	5705	14550
Marketplaces	203	203	182	203
LR: rent and distance linear vs free		19.94		
LR degrees of freedom		6		
LR p-value		0.0028		

Notes: Standard errors clustered at the marketplace in parentheses. Column (1) enters rent and distance linearly; column (2) replaces both with free level dummies; column (3) is a fixed-effects conditional logit in the adopted forms; column (4) is the main

4.4 The Value of Formality

Vendors value what a zone offers. A municipality designating an existing site as a vending zone, so that no vendor's commute changes, could charge the average vendor up to 4.48 per cent of monthly sales (95% CI [3.34, 5.68]), or INR 2,718 per month at median sales. At a free spot on those terms, predicted take-up is 67.9 per cent ([65.2, 70.7], Table 7). That figure averages each vendor's choice probability; the share of vendors whose valuation of the zone is positive at all is slightly higher, at 69.4 per cent ([66.6, 72.2]), and it is that share, not take-up, that Figure 14 traces across rent levels.

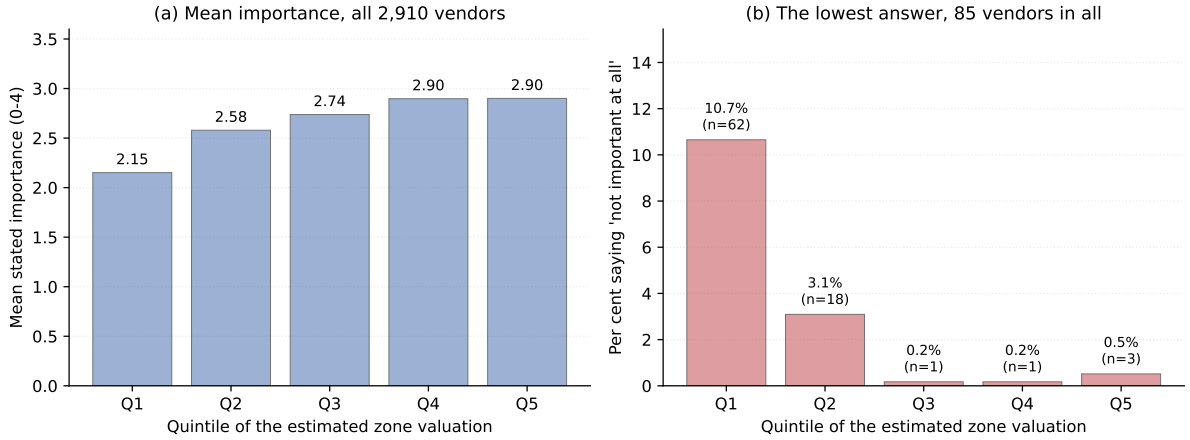
That valuation is contingent on what a vendor expects of their peers, and we report it both ways because the contingency is itself informative. The figure above is evaluated at the sample-mean belief, that 40 per cent of peers would also relocate. Evaluated instead at a belief that no peers would move, the value of the zone falls from 2.60 to 0.52 in utility, or to roughly 1.2 to 1.5 per cent of sales. We include the elicited belief in utility throughout, because without it the assigned peer share is badly confounded: the assigned share was constructed by shifting each vendor's own belief by a randomised shock, and omitting the belief loads that construction onto the coefficient we use to bound coordination effects. The cost is that the average vendor's own expectation then sits inside the intercept. Section 4.7 shows that exogenous shifts in the assigned share move nothing, which together with the contingency documented here suggests that the belief correlation reflects unobserved vendor optimism rather than genuine strategic complementarity.

What vendors say they are buying is consistent with those magnitudes. Among those choosing the zone, 85% cite protection from eviction or harassment and 60% the reduction of stress and uncertainty, while expected footfall and amenities are cited less often (Figure 13, panel (b)). The value of formality here is largely insurance value, which is what the formalisation literature would lead one to expect, since registration raises neither productivity nor earnings for most firms and the welfare case rests instead on risk reduction and, for own-account entrepreneurs, on inclusion in social insurance (de Mel, McKenzie, and Woodruff 2013; Benhassine et al. 2018; Farias and Hsu Rocha 2026; Ulyssea 2020).

A stated attitude the model never saw lines up with it. Legal-right importance was elicited after the choice tasks, and after treatment for half the sample, so it cannot enter the model. Comparing it against the model's output is nonetheless a check the estimation never saw, and a strong one for that reason. Mean stated importance rises monotonically across quintiles of the estimated valuation of a zone, from 2.15 in the lowest to 2.90 in the highest on a five-point scale, and the full cross-tabulation is far from independent ($\chi^2(16) = 323.8$); every vendor in the estimation sample answered the item (Figure 9, panel (a)). The same pattern appears at the bottom of the scale, where vendors in the lowest quintile are some twenty times more likely than those in the highest to say legal rights do not matter at all, 10.7 per cent against 0.5. We report that beside the means rather than leading with it, because only 85 vendors of 2,910 give that answer and three of the five quintiles contain fewer than five of them (panel (b)). Because the item is post-choice, part of this alignment may be rationalisation by vendors justifying choices already made, and we state that rather than leave it out; the reading is harder to apply to vendors who refused the zone in every round, whose answers are not justifying

a choice they made.

Figure 9: Stated Legal-Right Importance by Estimated Valuation



Legal-right importance was elicited after the choice tasks and, for the treated arm, after the information was given, so it is not a covariate in any model: this is a comparison the estimation never saw. Quintiles are of the posterior zone valuation under the main model, 582 vendors each. Panel (b) counts are on the bars because three of its five bars rest on one to three vendors; the contrast the text draws is between the first and the last. Chi-squared(16) = 323.7 on the full five-by-five table.

4.5 The Cost of Moving

Relocation consumes the value of the legal right, and it does so unevenly. Table 6 and Figure 10 report the maximum rent the average vendor would accept at each distance the design offered. It falls from 4.48 per cent of monthly sales at an unchanged commute to 2.09 per cent ([1.61, 3.34]) when the commute is one kilometre longer, and to 0.90 per cent ([0.18, 1.46]) when it is three kilometres longer. Between the two offered levels either side, the valuation reaches zero at about 3.8 kilometres of additional commute ([2.94, 4.75]). At the five-kilometre move the design offered, there is no rent, not even zero, at which the average vendor accepts; this is the same fact as the break-even distance rather than a second one, and we report it as such.[3]

More than half of what the legal right is worth is consumed by the first kilometre alone, and the two kilometres after it together cost less than that first one. The free dummies in Section 4.1 establish this without a functional form, and the step ratio of 0.43 ([0.31, 0.56]) states it as a single quantity. We do not decompose the distance response into a fixed cost of moving and a constant per-kilometre rate, because the design identifies that split only weakly: offered distances cluster at the levels drawn, and the two terms move in opposite directions across samples while every quantity reported here stays put. What the data support is the cost of named moves and the ladder above.[4]

[3] No distance between three and five kilometres was ever offered, so the break-even figure interpolates linearly between those two levels. We report it to one decimal place for that reason.

[4] The experiment also offered commutes one kilometre *shorter* than the status quo, and we find no measurable effect of such an offer. This does not establish that a shorter commute is not valued. Reductions were floored at the vendor's existing commute, so of the 1,540 occasions on which a shorter commute was offered, 986 (64.0%) delivered no change at all, and the mean realised reduction was 0.36 kilometres rather than one. The design does not identify the value of a shorter commute.

That the penalty attaches to leaving the site rather than to changing anything is not status quo bias: no comparable penalty attaches to a change in the rent, the expected number of competitors, the expected number of customers or the peer share (Appendix A9). The number of spots is the one exception, and its step is smaller than distance's. It is what the vendor gives up by leaving. Across the sample, the average vendor expects only 14.9% of their regular customers to follow to a new site and the median vendor just 5%, and 79% of vendors expect lower customer volume there (Figure 12), all elicited before any choice profile was shown. For a firm trading on the street, location is closer to the production technology than to an overhead, because the customer base does not move with the firm.

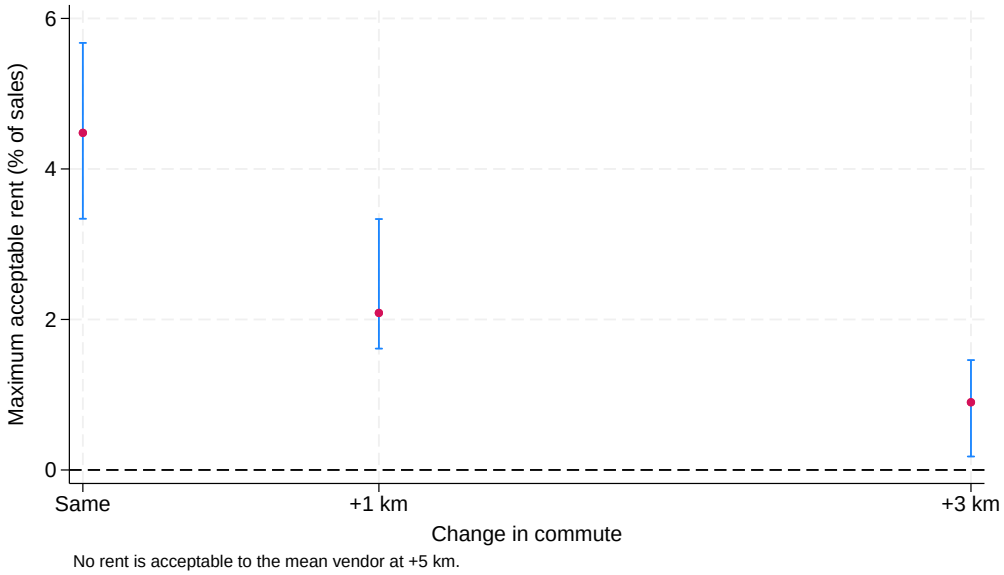
The ladder above is the average vendor's, and vendors differ widely around it. At an unchanged commute the median vendor's maximum rent is 3.15 per cent of sales, with a quartile range of 1.67 to 5.19, and 44.4% of vendors sit above eight per cent of sales, the highest rent the design ever offered, so their ladder entries are censored there. The same quartiles after a one-kilometre move are 1.38, 2.65 and 4.78.

Table 6: The Value of a Zone and the Cost of Distance

	Estimate	95% CI
Panel A. Maximum acceptable rent, per cent of monthly sales		
At an unchanged commute	4.48	[3.34, 5.68]
after a +1 km move	2.09	[1.61, 3.33]
after a +3 km move	0.90	[0.18, 1.46]
after a +5 km move	none	
At an unchanged commute, with information	5.83	[4.46, 7.66]
Break-even distance (km)	3.79	[2.94, 4.75]
Panel B. Trade-offs, in kilometres of commute		
Rent, per point below 1 per cent	0.99	[0.43, 1.69]
Rent, per point between 1 and 2 per cent	1.84	[1.40, 2.59]
Rent, per point above 2 per cent	0.64	[0.47, 0.90]
Information treatment	0.86	[0.39, 1.56]
Peer share, per 10 percentage points	0.00	[-0.08, 0.08]
Panel C. The zone in utility		
Value at the mean peer belief	2.60	[2.12, 3.07]
at a belief that no peers move	0.52	[0.13, 0.90]
with information	3.10	[2.54, 3.78]
Step ratio, $\beta(+1)/\beta(+3)$	0.43	[0.31, 0.56]

Notes: Mixed logit with a random intercept and a random per-kilometre slope, 2,910 vendors and 14,550 choice occasions. Intervals are a 200-replication marketplace pairs bootstrap. Rent is a percentage of each vendor's own monthly sales; one point is INR 607 at median sales. At a +5 km move there is no rent, not even zero, at which the average vendor accepts, which is the same fact as the break-even distance. The break-even distance interpolates between the +3 and +5 km levels, the only two the design offered either side of it.

Figure 10: Maximum Acceptable Rent by Distance



4.6 Which Vendors a Zone Reaches

Because the value of a zone is so dispersed, which vendors a zone draws depends on how good its terms are, and a municipality cannot identify those vendors in advance. Table 7 and Figure 11 report predicted take-up at each of the seven zone designs the model can evaluate without extrapolation, with every other attribute held at each vendor's own status quo. Take-up falls from 67.9% at a free zone with no change in commute to 61.9% one kilometre away, 53.6% three kilometres away, and 45.5% five kilometres away; charging two per cent of sales at an unchanged commute costs eleven points.

Grouping vendors by what the 27 observed characteristics predict about them spreads take-up at the best design from 49.9% to 86.0%, a range of 36 points. Grouping them instead by their estimated valuation, which uses their own choices, spreads it from 7.5% to 99.9%, a range of 92 points. The gap between the two is the part of the variation a survey cannot see. Quoting either figure alone misleads in opposite directions: heterogeneity dominates, part of it is predictable, that part is large in take-up terms, and it is still much the smaller part of the total.

The share of vendors whose choices respond to the terms on offer depends on how the question is posed, and we report it two ways rather than choose. Across the seven policy-relevant designs in Table 7, 34.9% of vendors ([32.0, 37.3]) change their predicted decision by more than 20 percentage points between the best and worst. Across the full grid of levels the design drew, which includes combinations no municipality would offer, the figure is 57.8% ([49.6, 64.8]). We take the first as the headline, because it is the one a policy could act on, and note that the spread across definitions is 23 points.

Inattention does not account for the vendors whose choices never vary. Only 4.2% of respondents failed the embedded attention check, failure rates do not differ across the range of estimated valuations ($p = 0.13$), and the vendors who refuse the zone in every round score no lower on cognitive ability than those who accept in every round. If

yea-saying drove the unconditional acceptances, the failure rate should stand out among those vendors, and it does not.

The stated reasons for vendors' choices make the same point in their own words. Figure 13 reports the reasons elicited after the final choice task, for staying in panel (a) and for relocating in panel (b). Among vendors choosing to stay, high rent and distance from home are the two most common reasons, followed by features of the incumbent spot itself: an established customer base, the time required to rebuild it, and attachment to the current area. Coordination concerns are uncommon.[5] For firms operating without premises, the price of formality is inherently spatial: it requires exchanging an established location for a new one.

Table 7: Predicted Take-up across the Design Ladder

Design	Δ km	Rent (%)	Take-up (%)	95% CI	With information	Difference
D0	0	0	67.9	[65.2, 70.7]	71.1	3.2
D1	0	1	64.2	[61.8, 66.9]	67.5	3.3
D2	0	2	56.7	[54.0, 59.4]	60.3	3.5
D3	1	0	61.9	[58.7, 64.8]	65.3	3.4
D4	1	1	58.0	[54.4, 60.8]	61.4	3.5
D5	3	0	53.6	[50.7, 57.1]	56.9	3.3
D6	5	0	45.5	[42.0, 48.9]	48.5	3.0

Notes: Predicted share of vendors relocating at each zone design, averaging over the estimated distribution of vendor valuations, with every attribute other than distance and rent held at each vendor's own status quo and the peer share at each vendor's own elicited belief. Rent is a percentage of own monthly sales. The last column is the effect of the information treatment at that design. These are the seven designs the model can evaluate without extrapolating beyond the offered levels.

4.7 What Does Not Explain Empty Zones

Three of the four explanations commonly advanced for unoccupied vending zones account for much less than the spatial cost, and one further design response can be ruled out alongside them. It is their joint weakness, rather than the size of the distance estimate on its own, that makes siting the binding constraint.

Rent: real, and bounded. Rent bites from the first percentage point and most steeply between one and two per cent of sales. A percentage point of rent is worth about 1.0 kilometre of commute below one per cent ([0.43, 1.69]), 1.8 kilometres between one and two per cent ([1.40, 2.59]), and 0.6 kilometres above two per cent ([0.47, 0.90]). The same result read off the valuation distribution rather than off predicted take-up is the share of vendors for whom the zone is worth more than the status quo at each rent the experiment actually offered, plotted in Figure 14. At an unchanged commute, 69.4% of

[5]Reported ties within the marketplace do not predict the valuation of a zone either. Vendors were asked how many others in their market they drink tea with, could ask to mind their stall, share business information with, could borrow from, would lend to, and are related to. These counts cover relations with other vendors in the current marketplace, so they do not speak to attachment to customers, to the location itself, or to ties elsewhere in the city. Appendix A8 reports the estimates.

Figure 11: Predicted Take-up across the Design Ladder

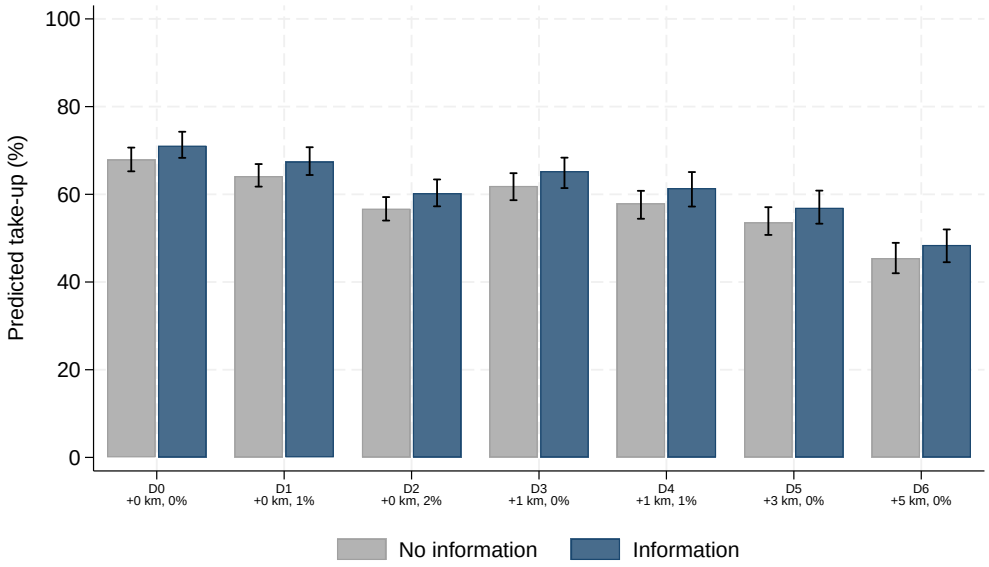
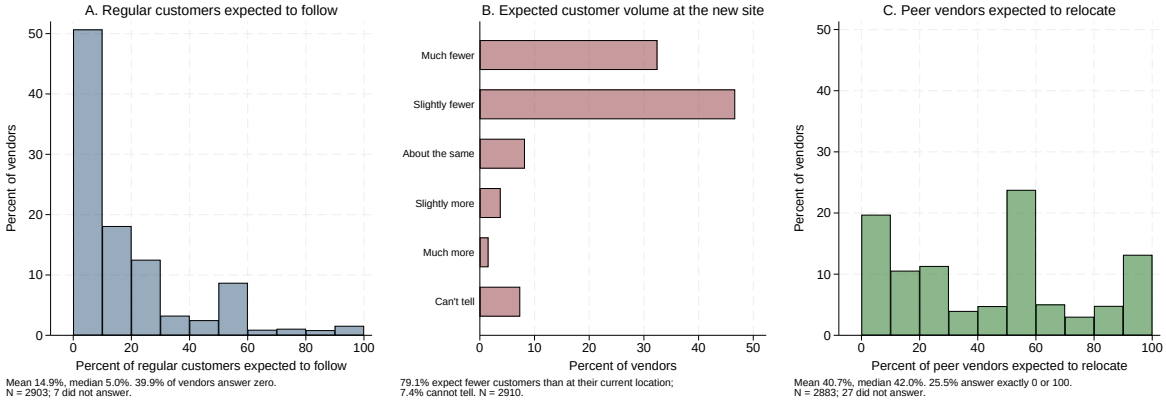
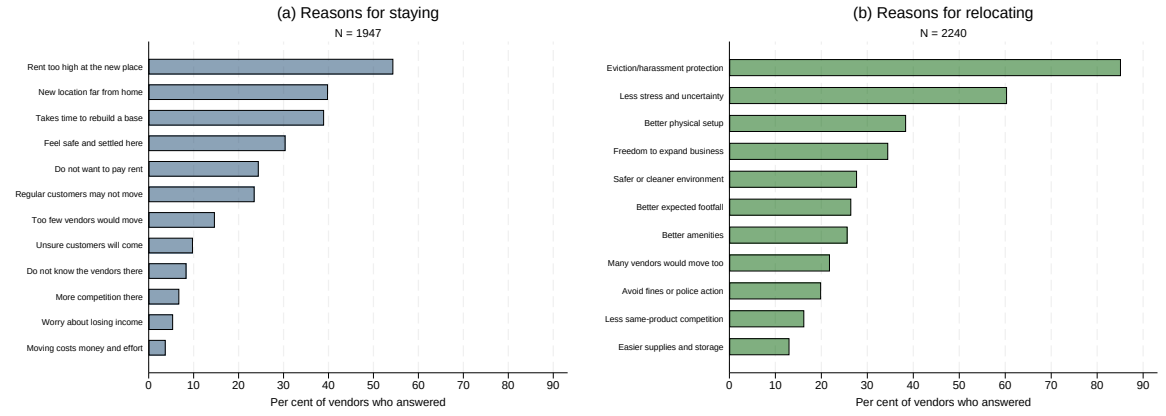


Figure 12: Pre-Choice Belief Elicitations



All three questions were asked before the first choice profile was shown:
they are items 1-3 of the choice module and the cards are item 4.
Estimation sample, one observation per vendor.

Figure 13: Stated Reasons for Staying and for Relocating



Reasons were elicited after the final choice task and, for the treated arm, after the information was given. They are descriptive only: no quantity estimated in this paper uses them. Each question was asked of vendors who made that choice at least once, so the denominators differ and a vendor who did both appears in both panels. Respondents could give more than one reason, so shares do not sum to 100. Other is omitted: 1 respondent for staying and 9 for relocating.

vendors have a positive valuation at no rent, 66.8% at one per cent of sales, 60.7% at two, 54.3% at four and 44.4% at eight. In predicted take-up, the measure used throughout the paper, raising the rent from nothing to two per cent of sales costs a municipality eleven points, while moving the zone one kilometre costs six (Table 7). Figure 14 also plots the schedule after a one-kilometre move, which lies below it at every rent. Rent matters because it is a new obligation rather than a substitute for one already borne: vendors pay nothing for their current pitch and only 9% make any informal payment. Over the short moves a municipality is most likely to be choosing between, rent and distance are therefore comparable instruments; it is against the three- and five-kilometre moves, which cost fourteen and twenty-two points, that distance dominates and rent cannot buy the siting decision back.

Coordination: no detectable effect. We find no effect of the expected share of relocating peers on the choice to relocate. The estimate is -0.0005 kilometres of commute per ten percentage points, zero to three decimal places, with a 95% confidence interval of $[-0.077, 0.077]$; in money it is -0.0007 per cent of sales $[-0.122, 0.110]$. The interval, not the insignificance, is what rules out a large effect, and it rules out anything larger than about 0.08 kilometres per ten percentage points. The range over which this is established is narrower than the design nominally spans, because the assigned share was censored at zero and 100 and 35.5% of occasions sit at one of those bounds, so a vendor already expecting ten per cent of peers to move received at most a ten-point reduction. Over the range the shock actually delivered, the bound is about 0.15 kilometres. In Table 7, setting the expected peer share to 100 per cent, the whole marketplace relocating together, moves predicted take-up at no design by more than 0.3 points.

This is the opposite of what the belief data alone would suggest, and the contrast is why the experimental variation is necessary. Expected peer relocation is strongly associated with willingness to relocate in the cross-section, and read that way would place coordination at the centre of the problem. But pessimism about peers travels with a lower valuation of the zone itself, so the observational association confounds the effect of peer behaviour with whatever else separates optimistic vendors from pessimistic ones, which is the simultaneity problem of Section 3 appearing in the data. When the share is assigned rather than believed, nothing happens. Zones are not empty because vendors are locked in a low relocation equilibrium.

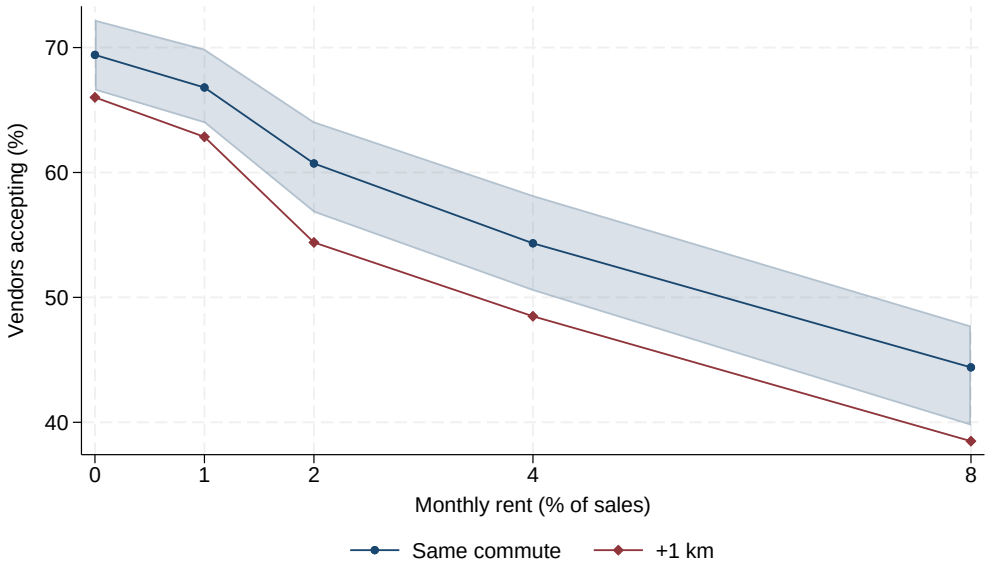
Zone size. The number of spots offered in the zone does not attract vendors, and the estimate cannot be distinguished from zero across the full design range of 20 to 160 spots (Table 5). One reading is that, conditional on expected competitors and the expected share of peers relocating, additional spots represent room for entrants who are not yet present, which a vendor with an established position has no reason to value. The estimate is small enough that we do not lean on that interpretation. What it does rule out is the simplest policy response to an empty zone, which is to build a larger one. This attribute is the number of spots, not the floor area of each one, which the choice tasks hold fixed; Section 4.8 treats spot dimensions as a feature of the Bihar rules rather than as an estimated quantity.

Enforcement information. The information treatment raises the value of a zone and changes none of the trade-offs. The Treat $\times$ Post coefficient on the zone intercept is re-

ported in Table 5; in a conditional-logit specification of the same design, interacting the treatment with every attribute does not reject the restriction that it shifts the intercept alone ($\chi^2(8) = 5.82$, $p = 0.67$). In willingness-to-pay terms the treatment is worth 1.35 per cent of monthly sales ([0.62, 2.25]), or 0.86 kilometres of commute ([0.39, 1.56]). It raises the maximum rent the average vendor would accept at an unchanged commute from 4.48 to 5.83 per cent of sales (Table 6), and lifts predicted take-up by about 3.3 points at every one of the seven designs in Table 7.

Given that only 1.3% of vendors could name the Act at baseline, and only 7.7% were aware of any such policy (Figure 5), this is consistent with an information channel: most vendors did not know the policy existed or was being enforced elsewhere, and learning this raises the value they place on a zone. Because the treatment combines general policy information with a specific example of enforcement, we cannot separately identify whether the effect operates through basic awareness or through the credibility a concrete example confers, and the extremely low baseline awareness makes that distinction difficult to isolate. That same low baseline means the estimate bounds what information provision can achieve rather than describing the return to a marginal campaign. Information can raise the value of a vending zone, but it cannot substitute for addressing where the zone is.

Figure 14: Share of Vendors with a Positive Valuation, by Rent Level, at an Unchanged Commute and after One Kilometre



4.8 Space Constraint and the Formalisation Gap

The one dimension of zone design our experiment cannot price is the space allocated to each vendor, because the choice tasks hold operating scale fixed by construction. Under the Bihar implementation of the Act each spot measures 6 by 4 or 5 by 3 feet. In our sample 85% of vendors operate from a stall or cart rather than a floor display (Table 1), and many sell perishable produce or cooked food requiring storage, preparation area, and customer-facing display. We do not measure current footprints, so the comparison rests on the stated spot dimensions alongside the observed distribution of stall types rather

than on a direct estimate.

This offers an explanation for the divergence between the willingness to relocate we estimate and the near-zero take-up observed. Our choice tasks tell vendors the zone differs from their current spot only on the attributes shown, and that storage capacity is unchanged. The vendors who accept a zone at an unchanged commute are therefore vendors who would relocate if a zone could accommodate them at their current scale, which the zones as built cannot. Even vendors willing to move on every other margin face a structural impediment that rent subsidies, information campaigns, or coordination mechanisms cannot resolve without first addressing the space allocated per vendor. This is a second spatial margin alongside the siting of the zone.

5 Conclusion

Place-based formalisation fails on the terms of the offer rather than on the value of what is offered. The legal right itself is not what vendors are rejecting: a municipality designating an existing site as a zone, so that no commute changes, could charge the average vendor up to 4.5 per cent of monthly sales for it, 68 per cent of vendors accept such a zone when the spot is free, and 85 per cent of those choosing a zone cite protection from eviction or harassment as the reason. What consumes that value is relocation, and it does so unevenly. More than half of what the legal right is worth is lost to the first kilometre of additional commute alone, and the two kilometres after it together cost less than that first one. The maximum rent the average vendor would accept falls from 4.5 per cent of sales at an unchanged commute to 2.1 per cent one kilometre away and 0.9 per cent three kilometres away, and reaches zero at about 3.8 kilometres, beyond which a zone is refused even when the spot is free. The three other explanations for empty zones account for much less. Rent matters but is bounded: charging two per cent of sales costs eleven points of take-up, against fourteen points for moving the zone three kilometres and twenty-two for moving it five. Building larger zones is not a remedy, since the number of spots has no detectable effect. And randomised variation in the share of peers a vendor expects to relocate, the explanation most often advanced for unoccupied formal space, has no detectable effect at all. What has to change is whether the vendor has to move, how far if they do, and the floor area of the spot itself.

5.1 Implications for design

Our estimates price three of those margins and bracket the fourth. The spot dimensions fixed by the Bihar rules are smaller than the space most vendors appear to need to run their operations, so relocation would require a reduction in operating scale that our choice tasks hold constant by construction (Section 4.8). Expanding the space allocated per vendor is therefore a precondition for translating willingness into take-up, including among the 68 per cent of vendors who would accept a zone at an unchanged commute.

The first lever is not an exchange rate at all. Because the cost of distance is front-loaded, siting a zone somewhat closer recovers the later kilometres and not the first. Where an existing vending site can itself be designated as a zone, the legal right is delivered without triggering that cost, and this is close to the only option our estimates support rather than merely the cheapest. It is also the option the municipality has the least room to choose.

The sites vendors occupy are the sites whose congestion the policy exists to relieve, and 90 per cent of vendors in our sample trade along main roads or streets with the remaining 10 per cent on pedestrian walkways, so in-situ designation will be feasible in some marketplaces and not others. Our estimates price what displacement costs the vendor; they do not price what leaving vendors in place costs the city, and the comparison that would settle the choice at any particular site requires both.

Where displacement is unavoidable, the scope for buying the distance back is narrow. Take-up at a free zone falls from 68% at an unchanged commute to 62% one kilometre away, 54% three kilometres away and 46% five kilometres away. The average vendor would accept at most 2.1 per cent of sales after a one-kilometre move, against 4.5 per cent for the same zone designated in place, so displacement costs the municipality something like 2.4 points of monthly rent per vendor in every case where the vendor accepts at all. Rent and distance are comparable instruments over this range rather than one dominating the other: waiving a two per cent rent buys back eleven points of take-up, about what moving the zone two kilometres closer achieves. A municipality weighing a site it can designate in place but must charge for against a free site further out should prefer the closer one only when it is closer by more than that. A free zone one kilometre away draws more vendors, 62 per cent, than a zone in place charging two per cent of sales, 57 per cent.

Rent should nonetheless not be set casually, because the earlier view that a small rent is free does not survive. Rent bites from the first percentage point of sales and most steeply between one and two per cent, a percentage point being worth about 1.0, 1.8 and 0.6 kilometres of commute in the three segments. There is no band inside which rent is costless and no threshold to avoid crossing. What there is instead is an instrument that works, but one whose every percentage point is charged against a budget that relocation has already mostly consumed.

Two further levers do less than is often supposed. Information provision raises what a zone is worth by about 1.3 per cent of monthly sales, or 0.9 kilometres of commute, and lifts take-up by a little over three points at every design we evaluate, without changing any trade-off: in a conditional-logit specification of the same design, interacting the treatment with every attribute does not reject the restriction that it shifts the intercept alone. Only 1.3 per cent of vendors could name the Street Vendors Act at baseline, and only 7.7 per cent were aware of any such policy, so that figure bounds what information can achieve from a standing start rather than describing the return to a marginal campaign. Coordination does less still. The standard prescription, that vendors must be moved together so that none is first, is not supported: assigning a vendor the expectation that the entire marketplace will relocate moves predicted take-up by about one hundredth of a point at every design. Mechanisms designed to solve a coordination problem will not fill these zones, because there is no coordination problem to solve.

Enforcement is the instrument this literature identifies as most effective at inducing formalisation (Andrade, Bruhn, and McKenzie 2016; Ulyssea 2020), and it carries a corollary here. Raising the cost of informality while zone spots remain too small to house existing operations would not move these vendors into the formal economy; it would make their informality more expensive without opening a viable alternative. Singapore, the one case we found in which the empty zone equilibrium was decisively broken, did exactly this,

withdrawing street hawker licences rather than improving the zones (Appendix A1), and the cost of any mismatch between what vendors needed and what the zones offered fell on the vendors.

A further implication follows from how little of this a municipality could anticipate. Vendors differ enormously in what a zone is worth to them, and the 27 characteristics we observe, including earnings, education, tenure and exposure to eviction, generate 15.5 per cent of that variation while reducing the unexplained part by only 6.1 per cent. Targeting therefore captures only a minority of the variation, and a programme cannot be aimed with any precision at the vendors most likely to take it up. What a municipality can vary instead is the terms of the offer, which is the case for treating siting as the policy variable rather than targeting.

Taken together, the estimates locate the failure of these programmes in a constraint rather than in a mistake. Zones sit empty not because the state has failed to offer formality or because vendors do not want it, but because the only instrument available for conferring legal status on a firm without premises requires moving it, and moving it consumes most of what the legal status is worth. The state is not absent here, and the firms are not myopic. What binds is that the site a vendor is most willing to be formalised in is the site the city is least willing to leave them in.

5.2 Transition costs and the long run

Our estimates price the decision to relocate at a point in time, and that decision bundles two things a policymaker would want to separate: the cost of moving and the cost of being somewhere else. The front-loading we document separates them in shape, since a cost that falls almost entirely on the first kilometre is not a cost of distance, but it does not separate them by duration, and part of it is plausibly transitional. Among vendors choosing to stay, 39 per cent cite the time needed to rebuild a customer base, and smaller shares cite the money and effort of setting up again (Figure 13). Such costs are by construction temporary, so part of what we attribute to leaving the site may be the disruption of moving rather than the standing disadvantage of the destination. If so, the short-run cost of relocation exceeds its long-run cost, and a design that looks unattractive on our estimates could still leave vendors better off once adjustment is complete. Several channels would push that way: secure tenure lengthens the horizon over which it is worth investing in a spot (Harari and Wong 2026), recognition carries eligibility for the collateral-free lending under PM SVANidhi, destinations need not stay as they begin (Franklin 2025; Agness and Getahun 2024), and a lawful spot removes a recurring source of loss rather than a one-off one. The evidence also cuts the other way, since relocated households in Mumbai were no better off than those who stayed (Barnhardt, Field, and Pande 2017) and formalisation raises neither productivity nor earnings for most firms (de Mel, McKenzie, and Woodruff 2013; Farias and Hsu Rocha 2026). We observe the choice and not what follows from it.

What does follow from our estimates, whichever way that evidence eventually settles, is an asymmetry worth taking seriously in design. The transition cost is borne by the vendor with certainty and immediately; the long-run gain is uncertain and deferred, and for a firm with the profit levels reported in Table 1 that asymmetry is not a matter of

patience but of whether the intervening months are survivable. This is an argument for compensating the move rather than assuming it will pay for itself, whether through a rent holiday over the period a customer base is rebuilt or through the working capital that PM SVANidhi already provides. A rent holiday is also the one use of rent our estimates clearly support, since rent is a weak lever against distance but is well suited to covering a cost that is itself temporary.

5.3 Limitations

These estimates are recovered from stated choices, which is what the setting permits: 99 per cent of vendors have never changed marketplace and the zones that exist are empty, so there is no revealed variation in the terms of relocation to estimate from, and a revealed-preference design here would recover the preferences of whoever had already moved. The experiment buys variation the field does not supply, at the cost that the choices are hypothetical. The design personalises every attribute to each vendor's verified circumstances and the choices were introduced as informing PMC policy, both of which narrow the gap; what they cannot guarantee is the level. Comparisons across designs are identified by randomisation and are the quantities we report and interpret.

Because the choice tasks hold operating scale constant, the space allocated per vendor is bracketed rather than priced. The Bihar spot dimensions are smaller than the operations most vendors appear to run, and that is consistent with the observed pattern of take-up, but we cannot say what a vendor would pay to avoid shrinking or how that cost ranks against distance and rent. The account we give of the gap between willingness and observed take-up is accordingly an interpretation the evidence supports rather than a quantity we estimate, and it rests on the stated spot dimensions alongside the observed distribution of stall types rather than on measured footprints. Six attributes also do not exhaust the design of a vending zone: tenure duration, enforcement inside the zone, infrastructure, and eligibility for schemes such as PM SVANidhi are held constant, so the estimates price what we vary conditional on the rest of the offer being as vendors imagine it.

Three features of the design set the range over which particular results hold. The experiment identifies the cost of a longer commute and not the value of a shorter one, because reductions were floored at each vendor's existing commute and most such offers delivered no change. The null on coordination is established over the range of peer shares the randomised shock actually delivered, which is narrower than the range the design nominally spans, because the assigned share was censored at zero and 100. And it is identified from an anonymous share of relocating peers: reported ties within a marketplace are few relative to its size (Appendix A8), and the design does not separate a vendor's response to that share from their response to knowing which particular vendors were moving.

One restriction in our specification the data reject. Allowing the cost of changing site at all to vary across vendors, as well as the cost of each kilometre, fits better on every criterion we examine. It leaves the maximum rent at an unchanged commute, the break-even distance and the value of information unchanged to within one per cent, and moves the intermediate rungs of the ladder by at most eight per cent; what it changes materially is the dispersion of the distance response over short moves, which it more than

doubles (Appendix A7). Statements about how widely vendors differ at short distances are therefore lower bounds, and we report the more conservative model.

What we price is the relocation decision, not its consequences. Whether vendors who relocate earn more, are evicted less, or are better off is a separate question that preferences cannot settle, and we do not measure the congestion costs that motivate relocation, so these estimates bound what displacement costs the vendor without establishing when it is nonetheless warranted. The setting is a single city; the instrument is common across the developing world (Appendix A1), while the spot dimensions, rent levels and commuter-led demand structure are Patna's. The stated reasons in Figure 13 were elicited after the final choice round and are post-treatment for the information arm; the beliefs in Figure 12 were elicited before any choice profile was shown.

References

Agness, Daniel and Tigabu Getahun (2024). "The Long-Term Human Capital Effects of Subsidized Housing: Evidence from Ethiopia". Working Paper.

Andrade, Gustavo Henrique de, Miriam Bruhn, and David McKenzie (2016). "A Helping Hand or the Long Arm of the Law? Experimental Evidence on What Governments Can Do to Formalize Firms". In: *The World Bank Economic Review* 30.1, pp. 24–54.

Barnhardt, Sharon, Erica Field, and Rohini Pande (Jan. 2017). "Moving to Opportunity or Isolation? Network Effects of a Randomized Housing Lottery in Urban India". In: *American Economic Journal: Applied Economics* 9.1, pp. 1–32. DOI: 10.1257/app.20150397. URL: https://www.aeaweb.org/articles?id=10.1257/app.20150397.

Benhassine, Najy et al. (2018). "Does inducing informal firms to formalize make sense? Experimental evidence from Benin". In: *Journal of Public Economics* 157, pp. 1–14.

Bénit-Gbaffou, Claire (2016). "Do Street Traders Have the 'Right to the City'? The Politics of Street Trader Organisations in Inner City Johannesburg, Post-Operation Clean Sweep". In: *Third World Quarterly* 37.6, pp. 1102–1129. DOI: 10.1080/01436597.2016.1141660.

Boonjubun, Chaitawat (2017). "Conflicts over Streets: The Eviction of Bangkok Street Vendors". In: *Cities* 70, pp. 22–31.

Brock, William A. and Steven N. Durlauf (Apr. 2001). "Discrete Choice with Social Interactions". In: *The Review of Economic Studies* 68.2, pp. 235–260. ISSN: 0034-6527. DOI: 10.1111/1467-937X.00168. eprint: https://academic.oup.com/restud/article-pdf/68/2/235/4369341/68-2-235.pdf. URL: https://doi.org/10.1111/1467-937X.00168.

Bromley, Ray (2000). "Street Vending and Public Policy: A Global Review". In: *International Journal of Sociology and Social Policy* 20.1/2, pp. 1–28. DOI: 10.1108/01443330010789052.

Bromley, Rosemary D. F. (1998). "Informal Commerce: Expansion and Exclusion in the Historic Centre of the Latin American City". In: *International Journal of Urban and Regional Research* 22.2, pp. 245–263.

Bromley, Rosemary D. F. and Peter K. Mackie (2009). "Displacement and the New Spaces for Informal Trade in the Latin American City Centre". In: *Urban Studies* 46.7, pp. 1485–1506. DOI: 10.1177/0042098009104577.

Bruhn, Miriam and David McKenzie (2014). "Entry Regulation and the Formalization of Microenterprises in Developing Countries". In: *The World Bank Research Observer* 29.2, pp. 186–201. DOI: 10.1093/wbro/lku002.

Campos, Francisco, Markus Goldstein, and David McKenzie (2023). "How should the government bring small firms into the formal system? Experimental evidence from Malawi". In: *Journal of Development Economics* 161, p. 103045.

Coletto, Diego, Lema Jaber, and Linus Vanhellemont (2021). "Street Vendors in Lima in the Time of COVID-19: Guilty or Oppressed?" In: *The Canadian Geographer / Le Géographe canadien* 65.4, pp. 435–447. DOI: 10.1111/cag.12712.

Cooper, Russell and Andrew John (Aug. 1988). "Coordinating Coordination Failures in Keynesian Models*". In: *The Quarterly Journal of Economics* 103.3, pp. 441–463. ISSN: 0033-5533. DOI: 10.2307/1885539. eprint: https://academic.oup.com/qje/article-pdf/103/3/441/5376673/103-3-441.pdf. URL: https://doi.org/10.2307/1885539.

Crentsil, Aba Obrumah and George Owusu (2018). "Accra's Decongestion Policy: Another Face of Urban Clearance or Bulldozing Approach?" In: *African Cities and the Development Conundrum*. Brill Nijhoff, pp. 213–228.

Crossa, Veronica (2009). "Resisting the Entrepreneurial City: Street Vendors' Struggle in Mexico City's Historic Center". In: *International Journal of Urban and Regional Research* 33.1, pp. 43–63. DOI: 10.1111/j.1468-2427.2008.00823.x.

de Mel, Suresh, David McKenzie, and Christopher Woodruff (Apr. 2013). "The Demand for, and Consequences of, Formalization among Informal Firms in Sri Lanka". In: *American Economic Journal: Applied Economics* 5.2, pp. 122–50. DOI: 10.1257/app.5.2.122. URL: https://www.aeaweb.org/articles?id=10.1257/app.5.2.122.

Donovan, Michael G. (2008). "Informal Cities and the Contestation of Public Space: The Case of Bogotá's Street Vendors, 1988–2003". In: *Urban Studies* 45.1, pp. 29–51. DOI: 10.1177/0042098007085100.

Eidse, Noelani, Sarah Turner, and Natalie Oswin (2016). "Contesting Street Spaces in a Socialist City: Itinerant Vending-Scapes and the Everyday Politics of Mobility in Hanoi, Vietnam". In: *Annals of the American Association of Geographers* 106.2, pp. 340–349. DOI: 10.1080/00045608.2015.1117936.

Farias, Alison and Roberto Hsu Rocha (July 2026). "Formalizing microentrepreneurs through targeted tax systems: Evidence from Brazil". en. In: *Journal of Public Economics* 259, p. 105677. ISSN: 00472727. DOI: 10.1016/j.jpubeco.2026.105677. URL: https://linkinghub.elsevier.com/retrieve/pii/S0047272726001131 (visited on 07/29/2026).

Franklin, Simon (Jan. 2025). *The Demand for and Impacts of Government Housing: Evidence from Ethiopian Lotteries*. Working Papers 986. Queen Mary University of London, School of Economics and Finance. URL: https://ideas.repec.org/p/qmw/qmwecw/986.html.

Gateway Research Centre (2022). *Street Vendors in Kampala City, Uganda: A Question of Urban Economic Policy Reforms*. Policy Brief. Gateway Research Centre. URL: https://www.gatewayresearchcentre.org/policy-briefs/street-vendors-in-kampala-city-uganda-a-question-of-urban-economic-policy-reforms.

Gisanabagabo, Sabin and Emmanuel Jules Masengesho (2024). "Effects of Locating Street Vendors to Formal Markets on Their Socio-Economic Living Conditions in Kigali City". In: *Journal of Research Innovation and Implications in Education* 8.3, pp. 503–512.

Harari, Mariaflavia and Maisy Wong (July 2026). "Slum Upgrading and Long-Run Urban Development: Evidence from Indonesia". In: *The Review of Economic Studies* 93.4, pp. 2646–2679. ISSN: 0034-6527. DOI: 10.1093/restud/rdaf090. URL: https://doi.org/10.1093/restud/rdaf090 (visited on 08/05/2026).

ILO (2018). "Women and Men in the Informal Economy: A Statistical Picture". In: *Geneva: International Labour Organisation*.

India, The Gazette of (2014). *The Street Vendors (Protection of Livelihood and Regulation of Street Vending) Act, 2014*. https://prsindia.org/files/bills_acts/acts_parliament/2014/the-street-vendors-act,-2014.pdf.

Lall, Somik V., Mattias K. A. Lundberg, and Zmarak Shalizi (Jan. 2008). "Implications of alternate policies on welfare of slum dwellers: Evidence from Pune, India". In: *Journal of Urban Economics* 63.1, pp. 56–73.

Mas, Alexandre and Amanda Pallais (Dec. 2017). "Valuing Alternative Work Arrangements". In: *American Economic Review* 107.12, pp. 3722–3759. URL: https://ideas.repec.org/a/aea/aecrev/v107y2017i12p3722-59.html.

McFadden, Daniel (1974). "Conditional Logit Analysis of Qualitative Choice Behavior". In: *Frontiers in Econometrics*. Ed. by Paul Zarembka. New York: Academic Press, pp. 105–142.

Middleton, Alan (2003). "Informal Traders and Planners in the Regeneration of Historic City Centres: The Case of Quito, Ecuador". In: *Progress in Planning* 59.2, pp. 71–123. DOI: 10.1016/S0305-9006(02)00061-2.

Mohammed Darwish, Hoda Al-Amir (2023). "A Planning Approach to Re-Envisioning Street Vendor Spaces in Cairo". In: *Journal of Urban Regeneration and Renewal* 16.3, pp. 303–317. DOI: 10.69554/UZQG5750.

Msoka, Colman Titus and Hezron Makundi (2023). "Formalisation of Street Vending in Dar es Salaam: Implementation and Enforcement of the Wamachinga Identity Card Initiative". In: *Forum for Development Studies* 50.2. DOI: 10.1080/08039410.2022.2112277.

Musoni, Francis (2010). "Operation Murambatsvina and the Politics of Street Vendors in Zimbabwe". In: *Journal of Southern African Studies* 36.2, pp. 301–317. DOI: 10.1080/03057070.2010.485786.

National Archives of Singapore (2026). *Moving Hawkers From the Streets to "Suites"*. Policy History, Archives Online. URL: https://www.nas.gov.sg/archivesonline/policy_history/card?id=23 (visited on 08/08/2026).

Permatasari, Anggraeni et al. (2014). "Strategic Analysis Relocating Street Vendor through 3D Negotiation. Case Study: Street Vendor Surakarta, Indonesia". In: *Procedia - Social and Behavioral Sciences* 115, pp. 436–443.

Rocha, Rudi, Gabriel Ulyssea, and Laísa Rachter (2018). "Do lower taxes reduce informality? Evidence from Brazil". In: *Journal of Development Economics* 134, pp. 28–49. DOI: 10.1016/j.jdeveco.2018.04.003. URL: https://doi.org/10.1016/j.jdeveco.2018.04.003.

Roever, Sally and Caroline Skinner (2016). "Street Vendors and Cities". In: *Environment and Urbanization* 28.2, pp. 359–374. DOI: 10.1177/0956247816653898.

Song, Shangcong (2020). "Street Stall Economy in China in the Post-COVID-19 Era: Dilemmas and Regulatory Suggestions". In: *Research in Globalization* 2, p. 100030.

Steel, William F., Tony Dogbe Ujoranyi, and George Owusu (2014). "Why Evictions Do Not Deter Street Traders: Case Study in Accra, Ghana". In: *Ghana Social Science Journal* 11.2, pp. 52–76.

Takeuchi, Akie, Maureen Cropper, and Antonio Bento (July 2008). "Measuring the welfare effects of slum improvement programs: The case of Mumbai". In: *Journal of Urban Economics* 64.1, pp. 65–84. URL: https://ideas.repec.org/a/eee/juecon/v64y2008i1p65-84.html.

Taylor, Laura O. and Ronald G. Cummings (June 1999). "Unbiased Value Estimates for Environmental Goods: A Cheap Talk Design for the Contingent Valuation Method". In: *American Economic Review* 89.3, pp. 649–665. URL: https://ideas.repec.org/a/aea/aecrev/v89y1999i3p649-665.html.

Ulyssea, Gabriel (2020). "Informality: Causes and Consequences for Development". In: *Annual Review of Economics* 12, pp. 525–546.

Appendix

A1 Relocation Policies Around the World

Figure A1: News Articles from around the world



The formalisation model we study in Patna, in which vendors are offered legally protected spots in a designated zone in exchange for relocating, is not an idiosyncratic feature of Indian urban policy. Table A1 documents 54 policy instruments in 27 countries across six regions, all but one dated between 2001 and 2026; the exception is Singapore's hawker resettlement programme of the 1970s and 1980s, the historical template from which much of the contemporary model derives. Just over 70% of the entries involve designating or constructing a physical site to which vendors are expected to move, which is the margin our experiment evaluates. The approach spans low-income settings with very large informal sectors (Tanzania, Uganda, Bangladesh), middle-income settings undergoing rapid redevelopment (Indonesia, Vietnam, Peru, Egypt), and high-income ones where vending is a small but politically salient share of employment (Singapore, the United States), under governments of every stripe. It is also durable: R. Bromley (2000) reviewed the same family of instruments a quarter of a century ago, and the design has changed little since.

Table A1: Street vendor formalisation, relocation and vending zone policies

Country and city	Policy or programme	Type
Panel A: South Asia		
India	National Policy on Urban Street Vendors (2004, rev. 2009)	L
India	Street Vendors (Protection of Livelihood and Regulation of Street Vending) Act (2014)	L, Z, R
India	PM SVANidhi micro-credit scheme (2020)	R
India, Bhubaneswar	Participatory vending zone regularisation (from 2007)	Z, R

Continued on next page

Table A1, continued

Country and city	Policy or programme	Type
Bangladesh, Dhaka	City Hawker Management Policy (2026)	L, Z, R
Bangladesh, Dhaka	DNCC temporary markets on six open fields (2026)	Z
Nepal, Kathmandu	KMC space allocation proposals and clearance (2010s)	E
Panel B: Southeast and East Asia		
Thailand, Bangkok	BMA vending ordinance (2002) and Designated Areas Directive (2005)	L, Z, R
Thailand, Bangkok	Return the Footpath to Pedestrians (from 2014) (Boonjubun 2017)	E, Z
Thailand, Bangkok	Revised BMA rules and hawker centre plan (2020s)	Z, R
Indonesia, Surakarta	Notoharjo and Panggung Rejo relocation (2006); Regulation No. 3 (2008) (Permatasari et al. 2014)	Z
Indonesia, Yogyakarta	Malioboro to Teras Malioboro relocation (2022)	Z
Indonesia, Jakarta	Relocation to designated lots (from 2015)	Z, E
Vietnam, Hanoi	Vending ban, 62 streets and 48 public spaces (2008) (Eidse, Turner, and Oswin 2016)	E
Vietnam, Ho Chi Minh City	Civilized Streets (2013 to 2015); sidewalk reclamation (2017)	E
Vietnam, Ho Chi Minh City	Nguyen Van Chiem and Bach Tung Diep food zones (2017)	Z, R
China	Street stall economy, State Council measures (2020) (Song 2020)	L, Z
China, Shenzhen	Designated vending areas (2023)	Z
Singapore	Hawker resettlement programme (1970s to 1986)	Z, R
Singapore	Hawker Centres Upgrading Programme (2001); SEHC model (2015)	Z
Philippines, Quezon City	Revised Market Code, Ordinance SP-2459 (2015)	L, Z, R
Philippines, Manila	Ordinance No. 8349 on designated vending areas	L, Z
Philippines	House Bill 10724, proposed national framework (2024)	L, Z, R
Panel C: Sub-Saharan Africa		
South Africa, Durban	Informal Economy Policy and Warwick Junction upgrade	L, Z
South Africa, Johannesburg	Operation Clean Sweep (2013) (Bénit-Gbaffou 2016)	E, Z
Tanzania	Vendor recognition (2016); Wamachinga ID cards (2018) (Msoka and Makundi 2023)	R
Tanzania, Dar es Salaam	Machinga Complex and Kariakoo pitch allocation (from 2017)	Z
Rwanda, Kigali	Vending ban and mini-market programme (2017, 2022) (Gisanabagabo and Masengesho 2024)	E, Z

Continued on next page

Table A1, continued

Country and city	Policy or programme	Type
Uganda, Kampala	KCCA gazetted markets directive (2022)	E, Z
Kenya, Nairobi	Pop-Up Markets and Street Vending Act (2019)	L, Z, R
Kenya, Nairobi	Mwariro and Karandini market relocation (2017 to 2019)	Z
Ghana, Accra	AMA decongestion by-laws (2011) and clearance exercises (Crentsil and Owusu 2018; Steel, Ujoranyi, and Owusu 2014)	E, Z
Zimbabwe	Operation Murambatsvina (2005) (Musoni 2010)	E, Z
Zimbabwe, Harare	Designated vending sites directive (2015)	E, Z
Liberia, Monrovia	Memorandum of Understanding with the City Corporation (2018)	L, R
Senegal, Dakar	Petersen Market eviction and PASI negotiations (2023)	E
Panel D: Middle East and North Africa		
Egypt, Cairo	Relocation of downtown vendors to Torgoman (2014)	E, Z
Egypt, Cairo	Six alternative public markets, incl. El Tunisi, El Turjuman and Ahmed Helmy (Mohammed Darwish 2023)	Z
Egypt, Cairo	Al-Tounsi Market relocation, UNESCO management plan (2010 to 2014)	Z
Panel E: Latin America		
Mexico, Mexico City	Programa Integral de Rehabilitacion del Centro Historico (2007) (Crossa 2009)	Z, E
Peru, Lima	La Parada closure and Santa Anita relocation (2012) (Coletto, Jaber, and Vanhellemont 2021)	Z, E
Colombia, Bogota	IPES REDEP kiosk network (Donovan 2008)	Z, R
Colombia, Bogota	Politica Publica de Vendedores Informales 2023 to 2035; RIVI register	L, R
Colombia, Bogota	Public space decree, confianza legitima protocol (2026)	L, Z
Ecuador, Quito	Relocation to eleven Centros Comerciales de Ahorro (2003) (Middleton 2003; R. D. F. Bromley 1998)	Z
Brazil	Microempreendedor Individual, Complementary Law 128 (2008)	L, R
Brazil, Sao Paulo	To Legal, Termo de Permissao de Uso for ambulantes	R
Chile, Santiago	Ordenanza No. 127 on comercio estacionado (rev. 2026)	L, R
Paraguay, Ciudad del Este	Municipal formalisation project	Z, R
Panel F: North America		
United States, California	Safe Sidewalk Vending Act, SB 946 (2018)	L, R
United States, Los Angeles	City and County Sidewalk Vending Ordinances (from 2018)	L, R

Continued on next page

Table A1, continued

Country and city	Policy or programme	Type
United States, California	SB 972, retail food code reform (2022)	L, R
United States, New York City	Local Law 18, permit expansion and enforcement office (2021)	L, R
United States, New York City	Advisory Board recommendations on designated spaces (2022)	Z

Notes: Non-exhaustive. Street vending is regulated mainly at municipal level, so the full universe runs to thousands of by-laws, most undocumented in academic or English-language sources. This is a curated sample weighted towards cases with an identifiable relocation or zone-creation component and a citable source. Academic treatments are cited where they exist; otherwise the policy name links to its primary source. Types: L, legislation or policy framework; Z, designation or construction of vending zones or market facilities; R, registration or licensing of vendors; E, clearance paired with an offer of an alternative site. Entries may carry more than one code. Spanish and Portuguese titles are given without diacritics.

A subset of these cases matches the Patna setting closely, in that a purpose-built facility was constructed and then under-occupied for a documented reason. Mohammed Darwish (2023) evaluates six alternative public markets in Cairo and identify unprofitable location as the crucial factor behind vendor rejection, alongside a failure to accommodate vendors' operating requirements in the design of the stalls. In Surakarta, the Notoharjo market was initially well attended but drew persistent complaints of insufficient footfall as its facilities deteriorated (Permatasari et al. 2014). In Lima, larger wholesalers moved to the new Santa Anita facility while smaller traders, for whom the rents were unaffordable, stayed put (Coletto, Jaber, and Vanhellemont 2021). In Kampala, a survey conducted after the 2022 directive to trade only from gazetted markets found average daily earnings of UGX 50,000 among street traders against UGX 7,000 among those trading from the zoned markets, with 85% of relocated traders reporting lower income (Gateway Research Centre 2022). In Nairobi, the Mwariro market was built with roughly 350 stalls against some 7,000 stalls whose traders were to be moved out of the central business district (PesaCheck, 2018), and stalls there remained unoccupied years after completion (*Daily Nation*, 2024). These map onto the three constraints our estimates identify: distance and foregone footfall, rent, and allocated space. The qualitative literature has recognised the pattern, with R. D. F. Bromley and Mackie (2009) showing how stall allocation rules and pricing in new Latin American municipal markets excluded the poorest traders, and Roever and Skinner (2016) reporting vendor complaints about the viability of relocation sites in Lima and Mexico City. What it does not do is quantify how much of the rejection is attributable to distance, to rent, or to the spot itself, since these move together across cases and cannot be separated with observational data. That decomposition is our contribution.

Singapore is the instructive contrast, being one of the few documented cases in which the empty zone equilibrium was broken. In 1975, one in six stalls in the new hawker centres sat empty while some 15,000 hawkers still traded on the street, and the reasons recorded for their reluctance were that street spots carried no rent, that hawkers had established clienteles they feared losing, and that a street stall could be enlarged simply by adding tables and chairs (National Archives of Singapore 2026). Those are the three margins our estimates identify, half a century earlier and on another continent. The equilibrium was

reversed not by improving the zones but by removing the alternative: licensed hawkers had to occupy their allocated stalls within two weeks or face a month's suspension, and new street hawker licences ceased to be issued in 1978. Street hawker numbers fell to around 4,000 within five years and to roughly 1,000 by 1983, and the last street hawker was relocated in 1986 (National Archives of Singapore 2026).

A2 Representativeness

To assess the representativeness of our sample, we compare characteristics of survey respondents to vendors who were recruited but did not complete the survey. Table A2 presents this comparison. Survey respondents differ from non-respondents on two observable characteristics. Participants are more likely to sell produce (59.2% vs 49.5%), and less likely to sell cooked food and drinks. Participants are also slightly less likely to be male (88.3% vs 91%). Age and the possession of stalls or carts are balanced across the two groups. These patterns are consistent with two features of the data collection process. First, street food vendors require more continuous attention to preparation and service than produce vendors, leaving them with less downtime to participate in a 45-minute survey, plausibly explaining their lower participation rate. Second, in husband-and-wife operations, recruited husbands frequently delegated the main survey to their wives, which likely accounts for the lower share of male respondents among participants relative to recruits. Taken together, the differences in sample composition reflect operational and household constraints on survey participation rather than systematic selection on the economic or locational characteristics that are central to this paper's analysis.

Table A2: Balance Table for Participants and Non-participants

	(1)	(2)	(3)
Variable	Non-Participants	Participants	Difference
Age	37.687	38.636	0.949
	(12.298)	(12.756)	(0.584)
Male	0.910	0.883	-0.027*
	(0.287)	(0.322)	(0.015)
Fresh Produce	0.495	0.592	0.097***
	(0.500)	(0.492)	(0.023)
Cooked Food/Drinks	0.505	0.408	-0.097***
	(0.500)	(0.492)	(0.023)
On Main Roads or streets	0.887	0.899	0.012
	(0.317)	(0.302)	(0.014)
On Pedestrian Walkways	0.113	0.101	-0.012
	(0.317)	(0.302)	(0.014)
Have Stalls/Carts	0.856	0.851	-0.005
	(0.351)	(0.356)	(0.016)
On the Floor	0.142	0.149	0.007
	(0.349)	(0.356)	(0.016)
Observations	564	2,922	3,486

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

A3 Covariate Coefficients

Table A3 reports the coefficients that Section 4.3 omits from the estimation table: the four held design attributes, the elicited peer belief, and the 27 vendor characteristics. The characteristics shift the mean of the vendor-specific valuation of a zone and are the quantities behind the targeting results in Section 4.6. The education missingness indicator is separated in 87 of the 200 bootstrap replications and is held at the estimation bound there; it is excluded from every representative-vendor quantity in the paper, so it affects no reported number, but it is not a well-identified coefficient and should not be interpreted.

Table A3: Vendor Characteristics

	(1) Linear	(2) Level dummies	(3) Main model (M0)
main			
20	-0.031	-0.109	
	(0.056)	(0.066)	
40	-0.127*	-0.161**	
	(0.074)	(0.079)	
80	-0.041	-0.132*	
	(0.062)	(0.068)	
160	0.028	-0.038	
	(0.063)	(0.072)	
-10	-0.008	0.010	
	(0.085)	(0.084)	
-5	0.089	0.030	
	(0.057)	(0.062)	
+5	0.013	-0.045	
	(0.057)	(0.064)	
+10	-0.042	-0.130*	
	(0.072)	(0.077)	
Lowest	-0.098	-0.145*	
	(0.075)	(0.076)	
Lower	-0.015	-0.112*	
	(0.068)	(0.066)	
Higher	0.000	-0.059	
	(0.078)	(0.077)	
Highest	0.044	-0.005	
	(0.069)	(0.070)	
-50%	0.088	0.028	
	(0.068)	(0.073)	
-20%	0.058	-0.016	

continued on the next page

Table A3: Vendor Characteristics *(continued)*

	(1) Linear	(2) Level dummies	(3) Main model (M0)
	(0.068)	(0.071)	
+20%	-0.031	-0.115	
	(0.068)	(0.071)	
+50%	0.120*	0.041	
	(0.063)	(0.068)	
Belief: peers relocating (%)	0.013***	0.013***	0.040***
	(0.001)	(0.001)	(0.004)
belief_peer missing	0.800**	0.805**	2.876**
	(0.404)	(0.405)	(1.388)
Spots (10s)			0.010
			(0.007)
Competitors (10s)			-0.160***
			(0.061)
Daily customers (10s)			0.113***
			(0.028)
Peers relocating (%)			-0.000
			(0.002)
Expected regulars following (%)	0.010***	0.010***	0.030***
	(0.002)	(0.002)	(0.006)
belief_cust==Much fewer	-0.196	-0.194	-0.607
	(0.148)	(0.148)	(0.492)
belief_cust==Slightly fewer	0.104	0.107	0.322
	(0.143)	(0.143)	(0.483)
belief_cust==About the same	0.171	0.169	0.531
	(0.183)	(0.182)	(0.594)
belief_cust==Slightly more	0.392*	0.399*	1.134
	(0.234)	(0.235)	(0.751)
belief_cust==Much more	0.518	0.519	1.434
	(0.315)	(0.317)	(0.916)
belief_regulars missing	1.559***	1.547***	4.058***
	(0.390)	(0.390)	(1.471)
Marketplace eviction frequency, leave-one-out	-0.034	-0.033	-0.118
	(0.056)	(0.056)	(0.190)
harassment = yes	0.171*	0.172*	0.481*
	(0.089)	(0.088)	(0.283)
eviction = yes	-0.003	-0.003	-0.001

continued on the next page

Table A3: Vendor Characteristics *(continued)*

	(1) Linear	(2) Level dummies	(3) Main model (M0)
	(0.084)	(0.084)	(0.267)
license = yes	0.300***	0.297***	1.013***
	(0.092)	(0.091)	(0.292)
log(weekly profit + 1)	-0.060	-0.060	-0.259
	(0.057)	(0.057)	(0.183)
Male	0.014	0.012	0.046
	(0.119)	(0.120)	(0.394)
Married	0.235**	0.235**	0.936***
	(0.098)	(0.098)	(0.326)
Number of children	-0.013	-0.014	-0.071
	(0.030)	(0.030)	(0.100)
cognitive_ability	-0.034	-0.035	-0.119
	(0.072)	(0.072)	(0.232)
Social ties (tea, information, family)	0.049	0.049	0.131
	(0.045)	(0.045)	(0.154)
Economic ties (minding, borrowing, lending)	-0.096**	-0.096**	-0.335**
	(0.046)	(0.046)	(0.151)
log(market size)	-0.110**	-0.112**	-0.310**
	(0.046)	(0.046)	(0.156)
Age (years)	0.002	0.002	0.006
	(0.003)	(0.003)	(0.011)
Education (years)	0.012	0.012	0.027
	(0.008)	(0.008)	(0.026)
Education missing	2.021**	2.048**	7.668***
	(0.921)	(0.928)	(2.875)
On a pedestrian walkway	-0.335**	-0.337**	-1.015**
	(0.134)	(0.135)	(0.431)
Away from main roads	0.253	0.250	0.981
	(0.250)	(0.249)	(0.890)
Pays an informal fee at the current spot	-0.132	-0.137	-0.454
	(0.149)	(0.150)	(0.465)
log(status quo daily customers)	0.103	0.103	0.386
	(0.083)	(0.083)	(0.288)
log(1 + status quo commute, km)	-0.030	-0.017	-0.146
	(0.076)	(0.079)	(0.248)

continued on the next page

Table A3: Vendor Characteristics *(continued)*

	(1) Linear	(2) Level dummies	(3) Main model (M0)
Observations	14550	14550	

Notes: The four held design attributes, the elicited peer belief, and the 27 vendor characteristics of the main model. The attributes vary across choice occasions by design; the characteristics were all measured before the choice tasks. Standard errors clustered at the marketplace in parentheses. Columns match Table 5; the fixed-effects column is omitted because vendor characteristics are absorbed by the fixed effects. bc2-bc6 are the categories of expected customer volume, with the don't-know answer as the base; the miss_ variables flag item non-response on the belief questions, which are zero-filled so that no vendor is dropped. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A4 Balance for the Information Experiment

Table A4: Balance Table For Information Experiment

Variable	(1) Control Group	(2) Treatment Group	(3) Difference
Age	38.39	38.88	0.49
	(12.63)	(12.87)	(0.47)
Education	6.08	6.08	-0.00
	(4.77)	(4.75)	(0.18)
Male	0.88	0.88	0.00
	(0.32)	(0.32)	(0.01)
Married	0.82	0.82	-0.00
	(0.38)	(0.38)	(0.01)
Single	0.16	0.16	0.00
	(0.36)	(0.36)	(0.01)
Widowed/Divorced	0.02	0.02	0.00
	(0.14)	(0.15)	(0.01)
Backward Caste	0.88	0.86	-0.02
	(0.33)	(0.35)	(0.01)
Scheduled Caste/Tribe	0.09	0.10	0.01
	(0.29)	(0.31)	(0.01)
General	0.03	0.04	0.00
	(0.18)	(0.20)	(0.01)
Fresh Produce	0.60	0.59	-0.01
	(0.49)	(0.49)	(0.02)
Cooked Food/Drinks	0.40	0.41	0.01
	(0.49)	(0.49)	(0.02)
On Main Roads or streets	0.90	0.90	-0.00
	(0.30)	(0.30)	(0.01)
On Pedestrian Walkways	0.10	0.10	0.00
	(0.30)	(0.30)	(0.01)
Have Stalls/Carts	0.85	0.86	0.01
	(0.36)	(0.35)	(0.01)
On the Floor	0.15	0.14	-0.01

	(0.36)	(0.35)	(0.01)
Work 7 days a week	0.95	0.94	-0.00
	(0.22)	(0.23)	(0.01)
Hours of Operation	9.19	9.25	0.06
	(3.63)	(3.70)	(0.14)
Have employees	0.30	0.33	0.02
	(0.46)	(0.47)	(0.02)
Yesterday's Sales (INR)	2,927.72	3,039.17	111.45
	(2,477.57)	(3,211.32)	(106.40)
Yesterday's Profits (INR)	612.02	624.78	12.75
	(586.92)	(662.10)	(23.18)
Weekly Sales (INR)	18,493.02	20,014.04	1,521.03*
	(15,425.35)	(28,443.03)	(851.26)
Weekly Profits (INR)	4,270.38	4,249.35	-21.03
	(3,992.45)	(4,118.06)	(150.13)
Own Other Businesses	0.15	0.14	-0.02
	(0.36)	(0.34)	(0.01)
Never Changed Marketplaces	0.99	0.99	-0.00
	(0.08)	(0.08)	(0.00)
Upto 1 year	0.06	0.07	0.01
	(0.23)	(0.25)	(0.01)
1 to 5 years	0.21	0.21	-0.01
	(0.41)	(0.40)	(0.02)
5 to 10 years	0.21	0.19	-0.02
	(0.41)	(0.39)	(0.01)
More than 10 years	0.52	0.53	0.01
	(0.50)	(0.50)	(0.02)
Policy Awareness	0.08	0.07	-0.01
	(0.27)	(0.26)	(0.01)
Risk Aversion	3.57	3.62	0.05
	(1.72)	(1.73)	(0.06)
Observations	1,432	1,490	2,922

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

A5 Experimental Design Details

A5.1 Scenario Construction

The choice scenarios for this stated preferences experiment were constructed through a systematic filtering process applied to a full factorial design. We began by generating all possible combinations of six attributes, each with five levels, creating a complete factorial design of 15,625 scenarios. To ensure realistic and cognitively manageable choice tasks, we applied two key filtering criteria. First, we identified and removed all weakly dominating scenarios—alternatives that were unambiguously better than the status quo across all attributes with at least one attribute strictly better—as such scenarios provide limited information about vendor trade-offs. We retained only "mixed" scenarios where some attributes were improvements while others represented trade-offs, along with the status quo scenario and scenarios where only rent varied from current conditions. Second, to

maintain respondent comprehension and reduce cognitive burden, we restricted scenarios to those where a maximum of three attributes differed from the status quo, ensuring participants would face manageable comparisons rather than overwhelming changes across all six dimensions simultaneously.

From the pool of valid scenarios that satisfied both filtering criteria, we randomly selected 150 unique choice scenarios using a fixed random seed to ensure reproducibility. These 150 scenarios were then systematically organised into 25 blocks, with each block containing six choice tasks. Each respondent was assigned to one block, meaning they faced six binary choices comparing a new market alternative (described by varying attribute levels) against their current market (the status quo). This blocking strategy balanced the need for sufficient statistical power across the full sample with the practical constraint of minimising respondent fatigue. The random selection and assignment process ensured that the chosen scenarios maintained variation in attribute levels while avoiding systematic biases.

Table A5 reports pairwise correlations between attribute levels in the design matrix. All correlations are below 0.2, confirming that the design supports independent estimation of attribute effects without multicollinearity concerns.

Table A5: Correlation Matrix of Experimental Design Attributes

	Market Rent	Number of Spots	Commuting Distance	Relocating Vendors	Competitors	Customers
Market Rent	1.00					
Number of Spots	-0.12	1.00				
Commuting Distance	-0.19	-0.08	1.00			
Relocating Vendors	0.08	-0.11	-0.15	1.00		
Competitors	0.11	-0.02	-0.11	-0.12	1.00	
Customers	-0.04	0.04	0.19	-0.12	0.05	1.00

Note: Correlations calculated using coded attribute levels (1-5) from the experimental design matrix.

A5.2 Bias Mitigation

Hypothetical bias is mitigated by two design choices: personalising the status quo to each vendor's verified current conditions (Taylor and Cummings 1999), and informing vendors at the outset that their responses would directly inform actual PMC policy decisions, increasing incentive compatibility.

Inattention bias is addressed by including a rent attention check in the instrument. Approximately 4.2% of respondents failed this check, suggesting inattention is not a problem in the sample.

Survey fatigue is addressed by (a) the ordering of the information treatment, so that the round dummies absorb any drift in response patterns from fatigue, and (b) a risk aversion task implemented as an incentivised game between choice modules to maintain engagement.

A5.3 Sample Stated Preferences Vignette

Figure A2 reproduces a sample choice task as presented to respondents. Given the low literacy levels among vendors, each attribute was represented using visual aids alongside verbal descriptions read aloud by the enumerator. The status quo column reflected each vendor's self-reported current conditions, elicited in the survey module immediately preceding the choice tasks.

Enumerator Script for Choice Tasks

The following script was read aloud by enumerators before each choice task was presented. The script was standardised across all enumerators and all respondents.

> *For the rest of this exercise, we want you to imagine that you are being offered a stall in a newly established marketplace and you need to choose whether you want to stay in the current marketplace or you want to move to this new market …*
>
> *…The new market offered does not differ from your current market except for the characteristics listed here. Imagine that the amount of the products that you can store, and transportation costs for products remain the same in the new marketplace …*
>
> *…In this exercise, we want you to imagine that each card represents a new scenario of the world. The only options available to you are the ones shown on the card. We want you to tell us which option you prefer out of the two. There are no "right" or "wrong" answers- you should only choose the option that you like better (even if you may not like any of them).*

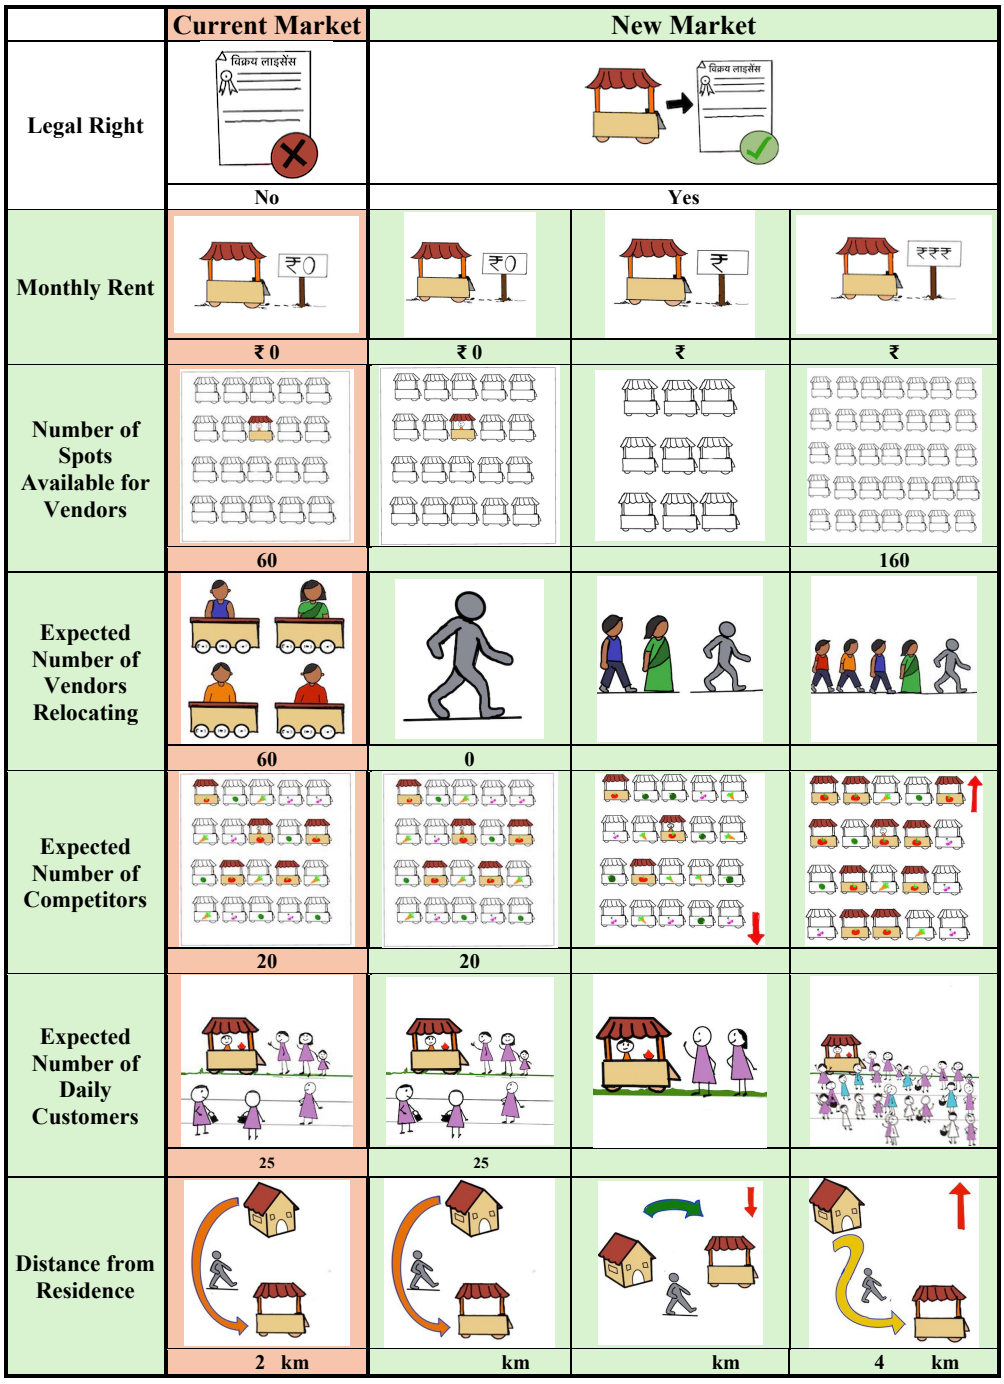
The six attributes presented in each choice task, and their possible values, are described in Table 2.

A5.4 Information Treatment Text

Control group respondents received no information and proceeded directly to round 4. The following was read aloud by enumerators to treatment group vendors immediately after round 3:

> *Let me tell you about a policy that concerns businesses like yours. It is called the Street Vendors Act of 2014. Under this act businesses like yours are offered spaces in specified vending areas with formal stalls to sell products. This is a national level policy and after its complete implementation, street vendors will not be allowed to work in unauthorized vending areas and will face fines if found to do so. For example, in Chandigarh, the municipality has situated several vendors in new marketplaces. News in December 2024 reported that about 3,000 vendors who didn't abide by the policy are not being allowed to vend in their old location in Chandigarh. The vendors who choose to move to the new marketplace will not have the option to return to their old location once the policy has been fully implemented.*

Figure A2: Sample Stated Preferences Choice Task



Notes: The choice task presents two alternatives: the vendor's current informal location (left column) and a hypothetical formal vending zone (right column with 3 sub-columns). Attribute values in the vending zone column vary across choice tasks according to the experimental design described in Section 3.1. The three sub-columns are used to depict three scenario possibilities: values remain at status-quo levels, decrease or increase. Visual aids replace text-based comparisons to accommodate low-literacy respondents. The practice round, which precedes the five analytical rounds, used a similar format and is excluded from all estimation.

A6 Functional Form

Each attribute in the design takes five discrete levels, assigned at random, so each can enter utility as a full set of level dummies. That specification imposes no shape on any attribute and only additivity across them. Every simpler form reported in the paper is a restriction of it, tested by a nested likelihood-ratio test. Table A6 reports those tests. The order matters for how the results should be read: there is one maintained model and a sequence of testable restrictions on it, rather than a search across specifications.

Distance. No parametric restriction on distance is accepted. The penalty at a one-kilometre longer commute is 0.43 of the penalty at three kilometres (95% CI [0.31, 0.56]), where linearity requires one third, and linearity above zero is rejected on the full sample ($p = 0.041$). We therefore retain the dummies. We do not report a decomposition into a fixed cost of moving and a per-kilometre rate: the design offers few distinct distances and the two terms move in opposite directions across samples while every reported quantity stays put, so the split is weakly identified.

Rent. The adopted form is a linear spline with knots at one and two per cent of sales. Against free rent dummies it imposes one restriction, which is not rejected ($p = 0.93$ to 0.95 across the two samples), and the three segment slopes are 0.99, 1.84 and 0.64 kilometres of commute per percentage point. Earlier drafts of this paper reported that rent carries no disutility below one per cent of sales, on a smaller and outcome-selected sample; the saturated model on the full sample rejects that ($p = 0.039$), and rent entering linearly from zero is rejected as well ($p = 0.008$).

Additivity. The maximum-rent ladder of Section 4.5 reads the cost of distance and the cost of rent off the same utility index, so it assumes the two are additive. This is the one maintained assumption the saturated model does not itself test: 150 scenarios cannot support a fully interacted design, and the experiment was not constructed to identify the interaction. We maintain additivity rather than test it.

A7 Choice of Heterogeneity Model

Section 4.2 establishes that vendors differ. How to model that difference is a separate decision from functional form, and a harder one, because the candidate models are not all nested and no single test settles it. Attribute forms are held at the adopted specification throughout, so these models differ only in their treatment of heterogeneity.

Discrete types. A latent class logit with three classes estimates the share of each type, their utilities, and class membership from one likelihood. It fits the distribution of choice counts very well, far better than any continuous specification, and its classes line up with

Table A6: Restriction Tests on the Saturated Model

Restriction	Attribute	χ^2	df	p	Verdict
Full sample, all five rounds					
(i) no step at zero	distance	0.01	1	0.9147	not rejected
(ii) linear above zero – D2	distance	4.18	1	0.0409	rejected
(iii) decrease effect is zero	distance	1.89	1	0.1695	not rejected
(iv) symmetry – D1	distance	0.39	1	0.5307	not rejected
(v) 1pct is free – R1 assumes	rent	4.25	1	0.0392	rejected
(vi) linear above 2pct – R3	rent	0.00	1	0.9488	not rejected
(vii) linear from zero	rent	11.81	3	0.0081	rejected
Marginal vendors					
(i) no step at zero	distance	7.31	1	0.0068	rejected
(ii) linear above zero – D2	distance	0.17	1	0.6793	not rejected
(iii) decrease effect is zero	distance	0.02	1	0.8864	not rejected
(iv) symmetry – D1	distance	8.11	1	0.0044	rejected
(v) 1pct is free – R1 assumes	rent	2.48	1	0.1152	not rejected
(vi) linear above 2pct – R3	rent	0.01	1	0.9297	not rejected
(vii) linear from zero	rent	29.11	3	0.0000	rejected

Notes: Each row is a nested likelihood-ratio test of one restriction against the free-dummy specification, which imposes no shape on any attribute. The adopted forms are those the data do not reject: free dummies for distance, and a linear spline in rent with knots at one and two per cent, whose single restriction against free rent dummies is row (vi). Rent being free below one per cent, assumed in an earlier draft of this paper, is rejected on the full sample, as is rent entering linearly from zero.

a variable the model never saw.[6] Four and five classes are rejected against three.

Continuous heterogeneity. A mixed logit with a normal random intercept beats the three-class model on log-likelihood, on BIC and on held-out prediction, and a Vuong test distinguishes them decisively ($V = 7.31$, $p < 10^{-12}$). Two flexible alternatives were estimated to check whether the normal is the wrong shape: a non-parametric distribution estimated on a grid of support points with the slopes held common, and mixtures of two and of three normal components. Neither pays for its parameters. BIC rises mechanically with the number of covariates, so figures from models with different covariate sets are not alternatives to one another and cannot be quoted side by side.

What the chosen model does not fit. The model we report predicts 273 vendors relocating in exactly one of the five rounds against 199 observed, a χ^2 of 34.7 against the latent class model's 2.0. The flexible shapes above repair this and still lose on BIC. We report the model that predicts choices better out of sample rather than the one that reproduces this marginal distribution best, and state the trade-off rather than leaving it to be found.

The form of the random distance response. The reported model lets the per-kilometre cost vary across vendors. An alternative lets the cost of changing site at all vary instead, and a third lets both. Table A7 reports all five variants. Of the two single choices, the per-kilometre form is clearly the better, by 53 BIC points. But one random distance coefficient is not enough: adding a vendor-specific cost of any change in commute improves the likelihood by 11.3 log points for one parameter (LR 22.5, boundary $p = 1.0 \times 10^{-6}$), improves BIC by 14.5 points, and is the only specification that also improves held-out prediction, on all three measures at once.

That richer model leaves the headline quantities unchanged: the break-even distance moves from 3.792 to 3.795 kilometres, the maximum rent at an unchanged commute from 4.48 to 4.44 per cent of sales, and aggregate predicted take-up agrees to within 0.3 points at every design in Table 7. The intermediate rungs of the ladder move by at most eight per cent, the largest being the maximum rent after a one-kilometre move, 2.09 against 2.25 per cent. Where the two models differ materially is the dispersion of the distance response at short moves, which the richer one more than doubles: they agree to within about two per cent about who responds to a five-kilometre move and disagree by a factor of 2.5 about a one-kilometre one. The quantities in Section 4.6 that depend on that dispersion, the responsive share and the take-up spread across vendors, are therefore lower bounds as reported here. We report the simpler model because its standard errors rest on the bootstrap used throughout the paper, and we state the restriction it imposes rather than omitting it. One assumption inside the comparison is untested: the two random distance coefficients are estimated as independent, and a correlated version would fit at

[6] Legal-right importance, elicited after the choice tasks. The class that refuses the zone in every scenario is many times more likely to say legal rights do not matter at all. Because the item is post-choice, part of this may be rationalisation.

least as well, so the statistic above is a lower bound on what the second coefficient is worth.

Table A7: Forms of the Random Distance Response

Random part	Log lik.	k	BIC	ΔBIC	Held-out	vs. M0
Per-kilometre slope only *(reported)*	-6005.87	53	12434.47	0.00	-1037.92	0.00
Any-move step only	-6032.40	53	12487.52	53.05	-1038.01	-0.08
Any-increase step only	-6024.42	53	12471.57	37.10	-1038.53	-0.60
Per-kilometre slope and any-move step	-5994.62	54	12419.93	-14.54	-1033.43	4.49
Per-kilometre slope and any-increase step	-5999.90	54	12430.49	-3.98	-1035.64	2.28

Notes: All five models share the adopted attribute forms and the same 27 covariates, and differ only in how the distance response varies across vendors. BIC is computed on vendors. Held-out is the log-likelihood on the control group's final two rounds, which no model saw. Of the two single random coefficients the per-kilometre slope is the better by 53 BIC points. Adding an any-move step to it improves the likelihood by 11.3 points for one parameter (LR 22.5, boundary $p = 1.0 \times 10^{-6}$), improves BIC by 14.5 and is the only specification that also improves held-out prediction. We report the simpler model and treat statements about dispersion at short distances as lower bounds.

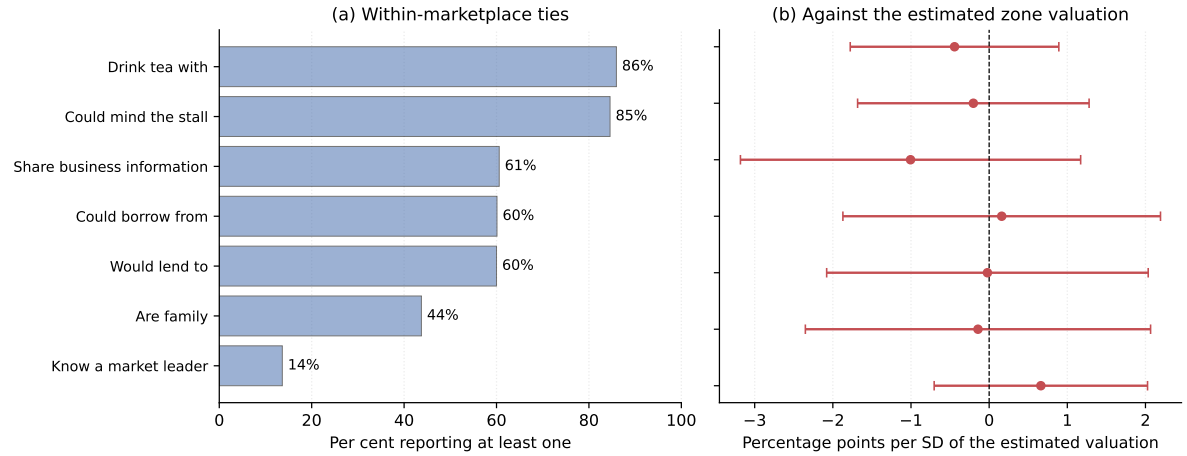
A8 Within-Marketplace Ties

Vendors were asked how many others in their marketplace they drink tea with, could ask to mind their stall, share business information with, could borrow from, would lend to, and are related to, alongside whether they know of anyone occupying an informal leadership role. Panel (a) of Figure A3 reports the share holding at least one tie of each kind. Social contact is widespread: 86 per cent name at least one vendor they drink tea with and 85 per cent at least one who could mind their stall. Economic relations are less common and are reported in similar numbers in each direction, 60 per cent naming at least one vendor they could borrow from and 60 per cent at least one they would lend to. Against a median enumerated market size of 50 vendors, the counts are small relative to the market.

These record how many vendors a respondent holds each relation with and not which ones, so they are within-marketplace degree counts and no network statistic beyond degree is recoverable from them. They are also bounded by the current marketplace, so ties elsewhere in the city and relations with customers or suppliers are not recorded.

Panel (b) relates each tie to the estimated valuation of a zone. Each estimate is a separate regression of the indicator on the vendor's posterior valuation, with marketplace fixed effects absorbed and standard errors clustered at the marketplace. None of the seven intervals excludes zero, and the point estimates run from -1.0 to $+0.7$ percentage points per standard deviation of the valuation against base rates between 14 and 86 per cent, so this is a tight null rather than an uninformative one. Ties within the marketplace do not predict how much a vendor values a zone. That is consistent with the null on the assigned peer share in Section 4.7 but is a separate and weaker statement, since these counts speak only to relations with other vendors in the current market.

Figure A3: Within-Marketplace Ties and the Estimated Valuation



Panel (a): share of the 2,910 estimation-sample vendors reporting at least one tie of each kind within their current marketplace, against a median enumerated market size of 50. Panel (b): one regression per relation of the indicator on the posterior zone valuation, marketplace fixed effects absorbed, standard errors clustered at the marketplace, 95 per cent intervals. The counts record how many vendors a respondent holds each relation with and not which ones, so no network statistic beyond degree is recoverable.

A9 Robustness

Is the distance cost status-quo bias? The paper attributes the front-loaded distance cost to leaving the site rather than to a general reluctance to depart from the status quo. Table A8 tests this directly, adding to the adopted specification a step for *any* change from the vendor's status quo on each attribute, so that each step is identified off a discontinuity at zero having already allowed the level to matter.

Table A8: Discontinuity at Zero, by Attribute

Attribute	Step	s.e.	p
Commuting distance	−0.134	0.060	0.026
Rent	−0.070	0.059	0.237
Number of spots	−0.093	0.045	0.039
Expected competitors	−0.026	0.050	0.605
Expected customers	−0.063	0.053	0.233
Assigned peer share	−0.005	0.051	0.918

Notes: Each row is the coefficient on an indicator for any change from the vendor's status quo on that attribute, entered alongside the attribute's own level. Standard errors clustered at the marketplace. Distance, rent and spots are estimated on a linear-in-level specification so that the step is separately identified: against the rent spline, "any rent change" is collinear with the first spline segment, and against the free distance dummies the step is collinear with the dummies themselves. The remaining rows are from the adopted specification with the steps added. The linear-in-level form understates the distance discontinuity relative to the free dummies, where the +1 km coefficient is −0.90 against a per-kilometre rate of −0.59; it is used here because a placebo requires the same functional form on both sides of the comparison.

Rent carries no discontinuity, and neither do expected competitors, expected customers or the assigned peer share. Rent is the informative case: it is the price-like attribute a status-quo reading would most expect to bite, and a vendor who simply disliked change would show a step there. The number of spots is the exception, carrying a step of −0.09, about two-thirds the size of the distance step in the same specification. We report it rather than set it aside. Note what it is not: the spots *level* remains indistinguishable

from zero, so vendors are indifferent to how many spots a zone offers and mildly averse to the number differing from their current market. That is a status-quo effect on one attribute, not a demand for capacity, and it does not account for a distance cost that consumes the entire value of the legal right within four kilometres.