

Do Local Markets Make Cities More Unequal? Evidence from Grocery Retail

Mateus Morais*

London School of Economics

October 2026

[Click here for the most recent version](#)

Abstract

This paper shows that local grocery markets make cities more unequal. I combine the price of every barcode at every formal grocery store in Maceió, Brazil, with a quantitative spatial model in which a neighborhood's grocery stores, their assortments and prices, its rent and the incomes of its residents are determined jointly. Richer neighborhoods pay more for the same barcode, yet their grocery cost of living is lower because they can reach more stores. Giving every household access to every store in the city would be worth 12.3 percent of total consumption to the poorest quartile of households and 6.8 percent to the richest, reducing the ratio of the richest quartile's real consumption to the poorest's by 5.1 percent.

JEL Classifications: R13, R23, L81, D12, D63, C43, O18

Keywords: spatial inequality, grocery retail, local cost of living, price indices, consumption amenities

*I am grateful to Daniel Sturm, Catherine Thomas and Veronica Rappoport for guidance and encouragement. I thank participants at the 2026 Urban Economics Association Summer School for suggestions. I also thank Economiza Alagoas for making its price data available. Any remaining errors are my own. Email: m.p.morais@lse.ac.uk.

1 Introduction

Spatial inequality is usually measured by households' earnings and what they pay to live where they do. Wages, housing costs and amenities differ across locations, and households of different incomes sort across them accordingly (Moretti, 2013; Diamond, 2016). Yet where a household lives also determines what it can buy, and hence its real income. Across cities, the cost of living depends on which products are available locally and at what prices (Handbury, 2021; Diamond and Moretti, 2021). Within a city the same logic applies at a finer scale, since each neighborhood has its own grocery stores, with their own shelves and prices, and two households with the same income can face different costs of living depending on where they live. This margin of spatial inequality has received far less attention, although it weighs most on poorer households, who spend a larger share of their budget on groceries.¹

Whether local markets widen or narrow inequality within a city is not obvious. Retailers choose where to open, what to stock and what to charge in response to the households around them. If stores in richer neighborhoods charge more for the same products, poorer households pay less and local markets compress inequality. If richer neighborhoods instead attract more stores with wider assortments of higher-quality products, local markets amplify it. Answering the question requires observing the stores each household can reach and the price of every product each of them sells, so that access, assortments and prices are measured on the same goods. At the same time, retailers' choices can affect households' decisions about where to live.

This paper asks whether local grocery markets make cities more unequal, and finds that they do. I combine administrative records of the price of every barcode sold by every formal grocery store in Maceió, Brazil, with a quantitative spatial model of the city. In the model, the stores that serve each neighborhood, what they stock and what they charge, the neighborhood's rent and the incomes of its residents are determined jointly. Estimating the model allows me to quantify a grocery cost of living by location within the city and by household income. Richer neighborhoods pay more for the same barcode, yet their grocery cost of living is lower. Their residents can reach more stores, and since households differ in which store suits them, a wider choice of stores outweighs the higher prices they face. Giving every household access to every store in the city without travel costs would reduce the ratio of the richest quartile's real consumption to the poorest's by 5.1 percent.

The data come from Economiza Alagoas, a price-transparency program in which the state treasury publishes the price, the barcode, and the seller recorded on every electronic invoice issued by retailers above a revenue threshold (SEFAZ-AL, 2026c). For six months up to 23 November 2025, I collect 14,687,078 transaction prices for 124,995 grocery barcodes at 2,923 establishments in the Maceió metropolitan area, and I place each establishment in its census tract, the unit

¹In the 2017–2018 Brazilian household consumption survey, groceries are 11.0 percent of the budget of the poorest quartile of households and 6.9 percent of the budget of the richest (Instituto Brasileiro de Geografia e Estatística, 2018).

at which the 2022 census reports household income. The records differ from the usual sources of local prices in three ways. They cover every barcode a store sells, where price surveys follow a fixed basket of items in a sample of outlets; they record the price at which each sale took place, and they cover every formal store. Independent stores make up 80 percent of the establishments.

Observing the same barcode at many stores, and many barcodes at each store, lets me measure the quality of every barcode from the prices at which it sells. I separate the log price of each barcode into a barcode effect, common to every store that sells it, and a store effect, common to every barcode the store sells, in the way wages are separated into worker and firm effects (Abowd et al., 1999). The barcode effect, measured relative to the other barcodes in its HS-8 category, is a store-invariant measure of product quality. It accounts for most of the within-category variance in log prices, 0.52 against 0.02 for the store effect.

Three facts describe how the grocery market of a neighborhood varies with the income of its residents, and they affect household utility in opposite directions. First, richer neighborhoods have more grocery stores within reach. The income elasticity of the number of sellers within 1.0 km is 0.507. Second, richer neighborhoods sell more distinct barcodes across stores, with an income elasticity of 0.310, and the additional barcodes sit at the top of each category's quality distribution. Third, the same barcode costs more in a richer neighborhood, with an income elasticity of 0.056. The first two facts favor richer neighborhoods and the third favors poorer ones. Which dominates depends on how households value an additional store or a better product against a lower price, which requires a model of where households shop and what they buy.

The model builds the grocery price index of each neighborhood from the decisions of households and stores, down to the individual barcode. Households choose among the stores within reach, trading travel costs against idiosyncratic tastes for individual stores (Miyauchi et al., 2025), and their valuation of quality rises with income (Handbury, 2021; Faber and Fally, 2022). Stores are multiproduct retailers (Hottman et al., 2016) that decide whether to operate, which barcodes to stock and what to charge. A store that lowers one price gains shoppers who also buy its other barcodes. Rents and residential choices follow the quantitative spatial framework of Ahlfeldt et al. (2015) and Redding and Rossi-Hansberg (2017). As well as in rent, locations in the model differ in grocery prices and assortments, which shapes where households choose to live. The model accounts for the first two facts. Richer households spend more on groceries, and so locations near them can support more stores. Stores near richer households stock premium barcodes because their customers value quality more. I estimate the model on the city and take its observed stores and prices as an equilibrium, so the model measures what the existing grocery market is worth to each neighborhood.

The first result is that access to stores, not prices, dominates variation in the grocery cost of living across neighborhoods (Section 6.2). A neighborhood with one percent higher income has a grocery price index 0.128 percent lower. The effective number of stores a household chooses among, which rises with the number of stores within reach and with how evenly the household's

purchases spread across them, lowers the index by 0.168 percent, more than the total difference. Prices raise it by 0.065 percent, and higher product quality offsets part of that rise.

The second result is that unequal access to stores costs the poorest households in the city the most (Section 7). I measure that cost against a benchmark in which every household can shop at every store in the city without travel costs, with prices and assortments at their estimated values while households move and rents adjust. Equal access would be worth 12.3 percent of total consumption to the poorest quartile of households and 6.8 percent to the richest. With prices and assortments held at their estimated values, the gains come from reaching more stores rather than from lower prices or better products. Removing travel costs while each neighborhood keeps the stores it already reaches is worth 2.6 percent of consumption to the poorest quartile. Lower budget shares for groceries account for 76 percent of the difference between the gains of the poorest and the richest quartiles, and the worse store access for poorer neighborhoods accounts for the rest.²

The third result is that competition for shoppers plays a large role in grocery pricing at the store level (Section 6.3). A store that raises one price loses customers for everything else it sells, and this margin offsets an amount equal to 58.5 percent of the mean markup. Markups are close to flat across the quality distribution, from 0.2334 in the bottom decile to 0.2291 in the top, so high-quality barcodes cost more because they cost more to supply, not because stores earn more on them.

The last result concerns how grocery retail responds when a neighborhood becomes richer (Section 8). Moving higher-income households into a middle-income neighborhood brings new store entry and higher-quality products to its shelves, and the price index falls by 0.164 percent for each percent increase in mean income. The rent increases that accompany the change in residents' incomes leave incumbents of every income group worse off. However, better stores raise the welfare of incumbents in every group.

This paper relates to three literatures. The first measures how the cost of living varies across space and across income groups. I allow grocery prices and access to stores and products to be determined jointly, whereas most studies quantify the welfare effects of grocery prices or of store access separately. Price indices that account for new products and for changes in quality and tastes (Feenstra, 1994; Broda and Weinstein, 2010; Redding and Weinstein, 2020) show that larger cities offer greater product variety and that its value depends on household income (Handbury and Weinstein, 2015; Handbury, 2021), new products directed at richer households lower their inflation relative to that of poorer ones (Jaravel, 2019), and Diamond and Moretti (2021) show how local prices shape the geography of consumption across U.S. cities. Within cities, poorer neighborhoods can face higher grocery prices (Eizenberg et al., 2021), local demand affects retail prices and variety (Stroebe and Vavra, 2019; Borraz et al., 2024), and supermarket entry lowers

²The size of the gains depends on tastes being independent across stores, and would be smaller if households that liked one store also liked similar stores nearby.

the cost of living of the households it serves (Atkin et al., 2018). Studies of food deserts find that differences in supply explain little of the nutritional gap between rich and poor households (Handbury et al., 2015; Allcott et al., 2019).

A second literature studies consumption amenities in quantitative urban models, where amenities respond to the income of local residents and reinforce residential sorting (Diamond, 2016; Couture and Handbury, 2020; Couture et al., 2024). Hoelzlein (2023) and Almagro and Domínguez-lino (2025) model the joint sorting of residents and amenities within a city and measure consumption amenities by the number of establishments in a location. Miyauchi et al. (2025) show that consumption trips shape the value of residential locations, and businesses anticipate and follow neighborhood change (Glaeser et al., 2018, 2023; Behrens et al., 2024; Hoelzlein and Miller, 2024). I construct the location-specific consumption amenity from what stores sell and what they charge, allowing it to respond to neighborhood income through assortments and prices as well as through the number of establishments.

A third literature, in industrial organization, studies how retailers compete in location, assortment, and price. Models of entry infer the intensity of competition from the number of firms a market supports (Bresnahan and Reiss, 1991; Berry, 1992; Mazzeo, 2002; Seim, 2006), and models of spatial competition show how distance shapes where stores locate and what they charge (Davis, 2006; Houde, 2012; Ellickson et al., 2020). A multiproduct firm internalizes substitution across its own products (Nocke and Schutz, 2018), and a retailer whose customers buy a basket on each trip can lower one price to draw shoppers who buy the rest (Lal and Matutes, 1994; Smith, 2004; Chen and Rey, 2012; Rhodes, 2015; Thomassen et al., 2017). Assortments are chosen jointly with prices (Draganska et al., 2009; Eizenberg, 2014; Fan and Yang, 2020), adapt to local tastes (George and Waldfogel, 2003; Quan and Williams, 2018), and rise in quality with the size of the market (Ellickson, 2007; Berry and Waldfogel, 2010). These models take the households in each market as given. In this paper, the income composition of the households a store serves is an equilibrium outcome of residential sorting, and it is the channel through which sorting reaches the shelf.

Section 2 describes the data and the quality measure, and Section 3 documents how the number of stores within reach, the barcodes they sell and the prices they charge vary with neighborhood income. Section 4 presents the model and Section 5 its estimation. Section 6 reports the model estimates, how stores set prices, and the decomposition of the grocery price index, Section 7 measures the cost of unequal access to stores, and Section 8 studies how grocery retail and the composition of neighborhoods respond to each other. Section 9 concludes.

2 Data

A household’s local retail environment is the set of stores within its reach and the goods those stores sell, so describing it means measuring both. The data assembled for this paper covers

both for an entire city. I observe the prices at which the formal grocery retailers of the Maceió metropolitan area sold each barcode over a period of six months, yielding 14,687,078 transaction records for 124,995 distinct barcodes across 2,923 establishments, and I geocode each establishment to the census tract in which it sits.

These records come from Economiza Alagoas, a price-transparency program operated by the Ministry of Finance of Alagoas (SEFAZ-AL). The program disseminates information from electronic consumer invoices (Nota Fiscal de Consumidor eletrônica, NFC-e) so that households can compare prices across establishments, and exposes the same records through a public API (SEFAZ-AL, 2026c,a). The API returns the commercial content of each invoice, the price charged, the product sold and the establishment that sold it; the buyer’s identifiers and the tax amounts assessed on the invoice are withheld (SEFAZ-AL, 2026b). Coverage extends to the whole of Alagoas, and I collect the price data for the 14 municipalities of the Maceió metropolitan area.

The dataset is comprehensive because issuing the invoice is compulsory for sellers above a revenue threshold. Alagoas phased the requirement in by firm size, and it was fully in force well before the sample window opens (Estado de Alagoas, 2015; Secretaria de Estado da Fazenda de Alagoas, 2017). The exemption defines the sample’s lower boundary. Individual micro-entrepreneurs (*microempreendedores individuais*) may issue the invoice but are not obliged to, and that status is capped at annual gross revenue of about US\$14,500 (Receita Federal do Brasil, 2025b), so any shop the boundary leaves out is small. Such shops add to counts of sellers but carry little weight in the model’s access measure, which weights each store by how much households spend there. Set against the number of grocery retailers Google Places reports as operating in the city, the sample holds 120% as many (Google, 2026), so the invoice data record more sellers than an independent directory finds.³

The unit of observation is a barcode–establishment pair, which exists when the barcode is recorded as transacted at that establishment during the window. A variety is a distinct barcode (the Global Trade Item Number, GTIN) and a product category is an eight-digit Mercosur Common Nomenclature code, which coincides with the eight-digit level of the Harmonized System and is referred to as HS-8; a single HS-8 category groups many barcodes. Each record reports a barcode, its HS-8 category, a product description, a unit of measure, a listed price and a sale price, the date of the last sale, and establishment identifiers with geographic coordinates. I use the sale price throughout. Geocoding places all but 3.2% of establishments in a census tract (Appendix A), and the median tract covers 0.06 square kilometres.

The sample window runs from 23 May 2025 to 23 November 2025, over which I query the API repeatedly. Each query returns, for a barcode at an establishment, the last transaction recorded there within the preceding 10 days, so repeated sweeps yield several observations for a given

³Appendix A documents that comparison and sets the price data against the barcode datasets this literature ordinarily uses.

pair, which I collapse to a single representative price, the median across sweeps. The median is taken over collection sweeps and not over sales, so it is not a sales-weighted price. The analysis therefore carries one price per pair for the whole window, and no result below depends on how prices move within it. Among pairs observed more than once, 52 percent never change price (Appendix A).

I restrict attention to food and grocery categories (HS-2 chapters 01–24), whose consumption is local. These span the grocery basket, from fresh meat, dairy and produce through staples such as rice, flour and sugar to packaged foods, non-alcoholic drinks and beer. The invoice universe also contains establishments that are not grocery retailers, among them gas stations, restaurants, and kiosks that sell a few food items alongside their main business. Two restrictions remove them. Records carrying no barcode after the harmonization documented in Appendix A, 0.3% of grocery records, are dropped, since a good that cannot be identified cannot be compared across stores, and this alone removes 1,412 establishments whose grocery activity is entirely unbarcoded. The seller counts of Section 3 therefore cover sellers of barcoded groceries, and a shop that sells only unbarcoded goods lies outside them. Among those that remain I require at least two distinct grocery barcodes, which drops a further 707 single-item sellers while discarding 0.02% of barcoded grocery transactions. Table 1 reports the resulting sample alongside the full HS-8 universe, and Appendix A, Table 18, traces it through to the estimation sample.

Table 1: Price sample: grocery categories and the full product universe

	Grocery	Full
Price observations	14,687,078	39,825,912
Barcode–establishment pairs	1,655,811	6,685,255
Barcodes	124,995	1,231,974
Establishments	2,923	8,125
Product groups (HS-8)	875	4,885

Notes: The grocery sample restricts the price data to food and grocery categories (HS-2 chapters 01–24) and to establishments carrying at least two distinct grocery barcodes, between 23 May 2025 and 23 November 2025. The full column is the entire HS-8 universe of formal establishments in the Maceió metropolitan area before these restrictions.

Two limitations follow from the invoice design. The first is that the price data records prices and not quantities, so every quantity the model needs is a modelled object rather than a measured one. The expenditure shares the demand system requires come from the consumption survey described below, not from the price data, and they are national rather than local.

The second is that sellers below the invoice threshold are absent. The mandate binds establishments, not workers, so the question is how much shop-based grocery retail is run by unregistered businesses. Among food-retail employers in Maceió, meaning those running a business with employees, 92.4% of the 36 respondents report that the business holds a CNPJ (Instituto Brasileiro de Geografia e Estatística, 2019). The unregistered remainder of the sector is concentrated among own-account sellers, of whom only 15.3% report one, namely street ven-

dors, open-air market traders (*feirantes*) and door-to-door sellers, who operate no fixed shop and would not enter a store-level sample however it were built.

Scale varies widely across the sample. The median establishment carries 107 distinct grocery barcodes, against 12 at the twenty-fifth percentile and 1,484 at the ninetieth. Variety is concentrated in the larger stores. The 77% of establishments carrying at least 10 barcodes account for 99% of all grocery variety-listings, so the small sellers add stores without adding much variety. The sample spans all 14 municipalities of the metropolitan area and the roughly 2,500 firms operating in them.

The price records concentrate in the largest stores. Table 2 classifies establishments into formats from their assortment. 67% of grocery price records come from 204 hypermarkets, and the remainder is spread across supermarkets, convenience stores, and specialty shops. A record is the last sale of a barcode at a store within a sweep’s look-back window, so these shares measure how many barcodes a format sells and how often they are recorded, not its share of sales.

Table 2: Grocery establishments by store format

	Establishments	Transactions (%)	Median varieties	Median HS-8 categories
Hypermarket	204	67	3,867	260
Supermarket	947	31	512	52
Convenience/minimarket	774	2	80	16
Specialty food store	324	1	35	5
Incidental seller	674	0	3	2
All establishments	2,923	100	–	–

Notes: The grocery sample is establishments carrying at least two distinct grocery barcodes (HS-2 chapters 01–24); single-barcode sellers are excluded. Transactions are price records, as in Table 1. Format is assigned from the assortment. Incidental sellers carry 2–9 distinct grocery barcodes; specialty stores concentrate more than 70% of grocery transactions in a single HS-2 chapter and carry at most 4 chapters (bakeries, butchers, beverage shops); the remaining broad-line stores are split by barcode count (convenience/minimarket under 200, supermarket 200–1,999, hypermarket 2,000 or more). The median columns are within-format.

Stores belonging to no chain are 80% of the grocery establishments in the sample, and chain coverage is spatially sparse, as Figure 1 shows. In the city of Maceió, chain stores are present in only 16% of census tracts, against 46% for independents, so of the 639 tracts containing any grocery retailer, 433 of them, or 68%, are served by independents alone.

In Brazil, independents are not the marginal segment they are in the United States. Grocery stores whose owner operates fewer than four outlets are 43 percent of American stores but only 11 percent of sales (Cho and Volpe, 2017), while the comparable Brazilian channel is 33 percent of outlets and 30 percent of grocery billing (NielsenIQ, 2025).

Figure 2 shows the variation that the decomposition (1) uses. It plots the price of Coca Cola 350ml, a physically identical item with an unambiguous unit, at each of the 925 establishments that carry it. The cheapest store charges US\$0.41 and the most expensive US\$2.15, a ratio of 5.2

for the same can.

Product quality. I measure product quality from the price data itself, which prices the same barcode at many establishments and many barcodes at each establishment. I decompose the log price of barcode i in HS-8 category n at establishment f into

$$\log p_{ifn} = \delta_i + \theta_f + \beta_n + \varepsilon_{ifn}, \quad (1)$$

the cross-sectional analog of the two-way fixed-effects wage decomposition of Abowd et al. (1999). The unit of observation is the barcode–establishment cell, priced at its median transaction over the sample window and weighted by the number of records behind that median, since a median over many records measures the cell’s price more precisely. As a worker effect is identified within the set of firms that workers’ moves connect, δ_i is identified within a connected set of the bipartite barcode–establishment graph, which barcodes listed in more than one establishment link together. A barcode sold in a single store is identified through that store’s other barcodes, but its effect has no second store to be checked against. Of all barcodes, 54.8% appear in more than one establishment, and the largest connected set retains 99.8% of them. Every barcode belongs to one category, so δ_i and β_n are separately identified only up to a constant within it. I normalize δ_i to mean zero within category, weighting by the number of underlying price observations, so that it measures what a product costs relative to the average barcode in its own group wherever it is stocked. I call δ_i quality. It carries most of the within-category variance of log prices, at 0.52 against the 0.02 the establishment effect contributes. It also lines up with a product attribute that played no part in estimating it, since barcodes carrying an organic certification score 0.30 log points above their category peers. Appendix A.7 reports the variance decomposition, the external validation, and the checks on additive separability.

Product attributes. Open Food Facts is a collaborative database of label and composition attributes for packaged food (Open Food Facts, 2026). I download a fixed revision of its product dump and keep the records whose barcode appears in the price data, retaining organic and premium labels and brand. 15,143 barcodes match, 12.1% of the grocery universe. None of these attributes enters the construction of δ_i , so they serve as an external check on it. Coverage skews toward branded packaged goods, so the matched set is a validation slice and not a representative sample.

Household demand. The 2017–2018 Brazilian household consumption survey (Pesquisa de Orçamentos Familiares) is IBGE’s national household budget survey. It records a household’s expenditures item by item, over reference periods running from a seven-day food diary to twelve-month recall for durables, together with its income and the state it is in (Instituto Brasileiro de

Geografia e Estatística, 2018). It supplies the grocery and housing budget shares of each income quartile, and its food diary, which records quantities alongside expenditure, gives the unit values from which Section 5 estimates the two elasticities. The Alagoas subsample is too thin to measure those shares by quartile, so they use the full national sample. Appendix A documents the mapping from survey product codes to HS-8.

Neighborhoods. Population, income, and housing measures come from the IBGE censuses and spatial layers, and two census years are used. The 2022 census is the current picture of the city and the basis for the analysis; it is enumerated at the census tract, the geography at which prices, establishments, and neighborhoods are matched throughout, and it supplies the income distribution the model’s four income groups are cut from. The model and every tract-level exhibit use the 1,725 tracts with at least 100 occupied households and a recorded income, since tract income is not meaningfully measured in smaller tracts, and the exhibits about stores use those of them that hold a grocery seller. The tract file reports the mean and variance of the income of household heads but not its distribution. The 2022 sample microdata, released in 2026, report household income for every weighting area in the country, and are what the income-composition map of Appendix A.6 is estimated on. Rents come from the 2010 census. The 2022 questionnaire did not collect the rent a household pays, so 2010 remains the most recent census that measures it. Section 5 states which census supplies each model input. The 2010 sample is a national file, and is used here beyond Maceió. It reports household income and rent at the level of weighting areas for every Brazilian city, which for the 125 cities with at least ten weighting areas gives the 2,774 neighborhoods on which Section 5 estimates the sorting elasticity.

Table 3 describes the metropolitan area on 2022 data.

Retail margins. The Annual Survey of Commerce (Pesquisa Anual de Comércio) reports, for each retail activity, the margin firms earn over the cost of the goods they sell (Instituto Brasileiro de Geografia e Estatística, 2024). In 2024 the rate ranged from 26.5% in supermarkets and hypermarkets to 52.3% in specialised food stores. I blend the grocery activities into a single rate of 33.4%, weighting each by its share of the barcode-establishment pairs of the sample’s stores whose activity code maps to one of the survey’s classes, so that the blend reflects where the sample’s own activity sits rather than the composition of Brazilian retail at large. Stores whose activity code maps to no class, such as wholesale-retail stores, do not enter the weights. This is a markup over cost, not a margin over sales. Section 5 uses it to set the level of marginal costs, which the price data does not identify separately from markups.

Establishment registry. The Receita Federal firm registry supplies each establishment’s legal name, economic activity, opening date, and firm-size indicator, joined to the invoice records on the firm identifier after the addresses have been standardized. It is what distinguishes a chain

from an independent, since it reports how many establishments a firm operates nationally. The registry is also published in bulk for the whole country. Each establishment record carries its economic activity code, postcode, opening date, and, for those that have closed, a closing date, so the register recovers the grocery establishments in operation at earlier dates. The register keeps each establishment’s current postcode, so an establishment that moved while keeping its registration is placed at its later address, and the reconstructed stock locates such establishments where they traded last. Section 5 uses the 1995 stock as an instrument for the sorting elasticity.

Terrain. Elevation comes from the Copernicus GLO-90 digital elevation model, a global raster at ninety-metre resolution (European Space Agency and Airbus, 2022). From it I compute the mean elevation of each weighting area and the share of its land that is steeply sloped, which enter the same estimation as controls.

Appendix A, Table 14, inventories every source the paper uses.

Table 3: Maceió metropolitan area, 2022 Census

	Value
Census tracts	2,221
Population	1,276,151
Households	439,799
Mean household size	2.90
Mean monthly income, household head (2025 US\$)	585

Notes: All rows are 2022 Census values for the 14 municipalities of the Maceió metropolitan area, aggregated over all 2,221 of its census tracts, of which the analysis uses the 1,725 with at least 100 occupied households and a recorded income. Households are occupied private households (*domicílios particulares ocupados*), which excludes vacant and occasional-use dwellings; mean household size is IBGE’s published mean number of residents per occupied private household. Mean household income is the tract-level mean monthly nominal income of household heads with income, averaged across tracts weighted by occupied households. Income is monthly, converted to 2025 US dollars. Nominal 2022 reais are deflated to 2025 by the IPCA and divided by the average 2025 market exchange rate of 5.59 reais per dollar. All monetary quantities in the paper use this conversion.

3 Motivating Facts

This section presents the facts about a neighborhood’s grocery environment that motivate the model’s key features. Across the census tracts of the Maceió metropolitan area,⁴ I document how three margins of that environment vary with neighborhood income. Those margins are how many sellers a household can buy from, what those sellers stock, and what they charge for an identical good. Sellers choose where to operate and what to put on their shelves in response to the households around them, so the three are determined jointly with a neighborhood’s income.

⁴The census tract is the unit throughout, and the level at which neighborhood income is measured.




Figure 1: Chain and independent grocery retailers

Notes: Each point is an establishment in the grocery sample. A retailer is a chain if its firm root operates more than one retail establishment in Brazil, and independent otherwise. The map covers the city of Maceió, which holds 77% of the sample's establishments.

Each of the three margins moves the real cost of a household's groceries, and they do not move it in the same direction, so which dominates is a quantitative question. Section 4 puts them in a common unit and separates decisions the data show only in combination, and Section 6.2 quantifies the role of each margin.




Figure 2: Price of a single barcode across establishments

Notes: Each point is one establishment carrying Coca Cola 350ml between 23 May 2025 and 23 November 2025, ordered by price; the dashed line is the median, US\$0.76. 3 establishments posting more than three times the median are omitted, on the ground that a price that far above the median for a single can is more likely a unit-of-measure error in the record than a price a shopper paid. Omitting them understates the dispersion the figure reports.

Grocery sellers locate where incomes are higher. Grocery shopping is local, so the revenue a seller can capture depends on the households close to it. Groceries are a necessity, and their share of household spending falls as income rises, from 11.0% in the bottom quartile of the national consumption survey to 6.9% in the top (Section 2). Grocery spending still rises with income, and households in the top quartile spend 2.9 times as much on groceries as those in the bottom. A neighborhood of richer households is therefore a deeper grocery market, and more sellers there can cover the fixed costs of operating. Allcott et al. (2019) find this pattern in the United States, where low-income neighborhoods have fewer supermarkets.

For each census tract I count the grocery sellers within 1.0 km of it and estimate

$$\mathbb{E}[S_\ell \mid y_\ell] = \exp(a + \gamma^s \log y_\ell), \quad (2)$$

where S_ℓ is that count and y_ℓ is the mean income of household heads in the tract. Estimation is by Poisson pseudo-maximum likelihood. The average tract in the lowest income quartile has 34 grocery sellers within 1.0 km, against 84 in the highest, an income elasticity of 0.507 across 1,725 tracts. Figure 3 plots the count against log tract income.

Richer neighborhoods have more sellers per household and not only more sellers. A neighborhood with more households supports more sellers whatever its income, so the raw count

could be reading population rather than income. Controlling for the number of households within the same radius leaves an income elasticity of 0.265, against 0.507 unadjusted, and sellers per thousand households run from 5.4 in the lowest income quartile to 9.1 in the highest. The count rises with income at other radii too, with the ratio between the top and bottom quartile at 2.4 within 0.5 km and 2.2 within 2.0 km (Appendix B.1).




Figure 3: Grocery sellers within 1.0 km, by neighborhood income

Notes: Each point is the mean count of grocery sellers within 1.0 km of a census tract, across the tracts in an equal-count bin of log tract income; bars are 95% confidence intervals for the bin mean and the line is a loess fit. Distance runs from the tract's representative point to each establishment, so the disc crosses tract boundaries. The count rises with income across the whole range. Appendix B.1 repeats the exercise at 0.5 and 2.0 km.

Richer neighborhoods offer a wider assortment, and a different one. A store stocks a barcode only if the sales it expects cover the cost of carrying it. Which barcodes do so depends on the households a store sells to, so a neighborhood's assortment reflects the composition of local demand and not only its level. Handbury and Weinstein (2015) document how the variety of goods available differs across US cities, and Handbury (2021) how the cost of living depends on the income of the household doing the buying. The facts below are the within-city counterpart. The number of distinct barcodes sold in a tract rises with its income, with an elasticity of 0.310 across the 857 tracts that hold at least one seller (Appendix B.2).

That total conceals two responses that separate by quality. I estimate the income elasticity of

variety for the barcodes at each quality percentile q ,

$$\mathbb{E}[B_{\ell q} | y_{\ell}] = \exp(a_q + \gamma_q^v \log y_{\ell}), \quad (3)$$

where $B_{\ell q}$ counts the distinct barcodes of quality percentile q sold anywhere in tract ℓ , y_{ℓ} is the mean income of household heads in the tract, and q ranks a barcode against the others in its own HS-8 category by the quality index of Section 2. Estimation is by Poisson pseudo-maximum likelihood over every tract–percentile pair, with standard errors clustered by tract.⁵ Run on establishments instead of tracts, the same regression gives the second profile in Figure 4.




Figure 4: Income elasticity of variety, by quality percentile and level

Notes: Each series is $\hat{\gamma}_q^v$ from equation (3), estimated by Poisson pseudo-maximum likelihood at each quality percentile and clustered by census tract. The neighborhood series counts the barcodes sold anywhere in a tract, the establishment series those sold by one store. Bands are 95% confidence intervals and the lines are loess fits. Percentiles rank barcodes within HS-8 category, so a high percentile is the premium end of a category and not of the basket. Both series rise with quality, and the establishment series is negative at the bottom of the distribution.

The neighborhood elasticity rises gently across the lower half of the quality distribution and steeply above it, from -0.032 at the fifth percentile and 0.110 at the tenth to 0.818 at the ninetieth and 1.029 at the ninety-ninth.

Stores in richer neighborhoods carry slightly more barcodes, and different ones. The establishment elasticity is -0.328 at the fifth percentile and 0.929 at the ninety-ninth, and pooled over

⁵A pair holding no barcode at that quality is a count of zero and not a missing observation. Regressing the log count would drop those pairs, and the tracts with no low-quality barcode are disproportionately the rich ones.

the quality percentiles it is 0.083, against 0.310 for tracts. A store in a richer neighborhood drops the cheapest barcodes and adds premium ones, which leaves its total close to unchanged.

The distance between the two profiles reflects both how many sellers a tract holds and how much their assortments overlap, and it is not constant. At the bottom of the quality distribution a neighborhood's total holds up while each store drops cheap barcodes, so the barcodes its stores together carry offset what each one sheds. At the top the two profiles nearly meet, so the premium response is visible in what individual stores stock. Appendix B.3 treats the measurement error in q .

The same barcode costs more in a richer neighborhood. The same product sells at markedly different prices across stores. The median grocery barcode carried in at least ten establishments has a cross-store standard deviation of log price of 0.16. A store one standard deviation above the mean charges roughly 17% more than a store at the mean, for a physically identical good, before any adjustment for store or neighborhood. What follows shows how much of that dispersion is organized by neighborhood income.

Stroebel and Vavra (2019) find that retail prices respond to local housing wealth, while DellaVigna and Gentzkow (2019) show that US chains set nearly uniform prices across their stores. Within a city, Eizenberg et al. (2021) find the opposite pattern in Jerusalem, where residents of poor peripheral neighborhoods pay some of the highest grocery prices and rarely shop outside them. In Montevideo, Borraz et al. (2024) find that new housing lowers local grocery prices and raises the variety available, as incumbents respond to the added demand. Most sellers in this setting are single-establishment independents (Section 2), which set their prices store by store.

I compare a barcode with itself across neighborhoods, estimating

$$\log p_{i\ell} = \gamma^p \log y_\ell + \omega_i + u_{i\ell}, \quad (4)$$

where $p_{i\ell}$ is the price of barcode i in tract ℓ and ω_i is a barcode fixed effect.⁶ Standard errors are clustered by census tract, the level at which the regressor varies. Over the 857 tracts that hold at least one seller, the income elasticity of price is 0.056. Figure 5 plots the residualized price against log tract income, which rises by 5.1 log points from the lowest income bin to the highest.

In the full sample the elasticity does not vary detectably with quality, unlike the elasticity of variety. It survives a control for the number of local competitors, keeps its level among single-establishment independents, and is unchanged on barcodes observed at both ends of the income distribution. Among independents, however, it is steeper for lower-quality barcodes (Appendix B.4 and Appendix B.5).

⁶A category fixed effect is neither included nor needed, since each barcode belongs to one HS-8 category and category effects are nested within ω_i . The dependent variable is the median of the barcode's price records in the tract, pooled across the stores that carry it, so a cell is a barcode in a neighborhood.




Figure 5: Price of the same barcode, by neighborhood income

Notes: Log price residualized on barcode fixed effects, so the comparison is of a barcode with itself across neighborhoods. Residuals are averaged within tract, and each point is the mean across the tracts in an equal-count bin of log tract income, weighted by the price records each tract contributes. Bars are cluster-robust 95% confidence intervals for the bin mean, clustering on tracts, and the line is a loess fit. The profile dips across the lowest-income bins, where the bin means are least precise, and rises from there to the top of the income range.

What these facts cannot separate is cost from markup. A store in a richer neighborhood pays more for floorspace, and rents are not observed at the barcode level, so a cost difference and a demand difference produce the same price elasticity here. Section 4 assumes marginal cost carries no neighborhood index, which is the assumption that makes the separation.

These are equilibrium quantities, not effects of income. The number of sellers within 1.0 km of a tract, the assortment they carry and the prices they post are chosen together by firms that observe the neighborhood. More sellers change what each one stocks and charges, and what they stock and charge changes how many can operate. A cross-section of census tracts separates none of this.

Section 4 supplies the separation. It writes down the household's choice of store and the firm's choice of prices and assortment, so that each fact becomes an outcome of primitives the model names. The assortment fact is the one the estimation uses directly: how the shelf's quality composition moves with the income of a store's customers identifies how the taste for quality rises with income (Section 5).

4 Model

This section presents a model of a city in which the grocery environment of a neighborhood, the prices and assortments of the stores that serve it, its rent, and its residential composition are determined jointly. The residential side shares its structure with the quantitative spatial framework of Ahlfeldt et al. (2015) and Redding and Rossi-Hansberg (2017), in which a location delivers an income, a housing cost and an amenity. The supply side is where the model departs from it. The consumption environment is produced by multiproduct retailers who choose where to operate, what to stock and what to charge, and who make those choices against the income mix of the households who shop with them, so groceries carry a local price index instead of a numeraire good priced the same everywhere.

Section 4.1 sets out demand, from residential choice through the store, category and barcode decisions that produce the index. Section 4.2 gives the retailers whose entry, assortment and pricing decisions determine it, Section 4.3 the floorspace market that prices the space households and stores compete for, and Section 4.4 how the pieces jointly determine the equilibrium.

4.1 Demand

A household chooses a neighborhood, then a store, then what goods to buy in the store. The first decision is standard, and it is set out here to fix notation and to show where the grocery price index enters utility. The store, category and barcode decisions build that index, and the subsection closes by decomposing it into the margins the rest of the paper measures.

Location choice. Let $\mathcal{L}$ denote the set of neighborhoods. Each household belongs to an income group $g \in \mathcal{G}$ and receives an exogenous income y_g . Conditional on neighborhood ℓ , a group- g household allocates that income across groceries, housing and a freely traded numeraire with Cobb–Douglas shares, so its indirect utility there is

$$u_{\ell g} = \frac{A_{\ell g} y_g \varepsilon_\ell}{\mathcal{P}_{\ell g}^{\beta_g} r_\ell^{\rho_g} p_0^{1-\beta_g-\rho_g}}, \quad (5)$$

where $A_{\ell g}$ is the amenity group g derives from neighborhood ℓ , ε_ℓ an idiosyncratic taste draw independent across households and neighborhoods, $\mathcal{P}_{\ell g}$ the price index of the grocery environment reachable from ℓ , r_ℓ the local rent, p_0 the price of the numeraire, and β_g and ρ_g the grocery and housing expenditure shares of group g . The numeraire is freely traded, so p_0 is common across neighborhoods and normalized to one, and the residual share $1 - \beta_g - \rho_g$ falls on it.

The draws ε_ℓ are Fréchet with shape κ_ε , common across groups. Writing $\bar{u}_{\ell g}$ for the deter-

ministic part of (5), the share of group g residing in ℓ is

$$s_{\ell g} = \frac{\bar{u}_{\ell g}^{\kappa_\varepsilon}}{\sum_{j \in \mathcal{L}} \bar{u}_{jg}^{\kappa_\varepsilon}}, \quad (6)$$

and the population it implies is $N_{\ell g} = s_{\ell g} N_g$, with N_g the total population of group g . Groups differ in where they sort through $\mathcal{P}_{\ell g}$, ρ_g and $A_{\ell g}$, each of which carries a group index.

Store choice. Having chosen a neighborhood, a group- g household chooses a store and a basket to solve

$$\max_{f \in \mathcal{F}_\ell, \{q_{ifn}\}} \frac{\nu_f U_g(\{q_{ifn}\})}{e^{\tau d_{f\ell}}} \quad \text{subject to} \quad \sum_{n \in \mathcal{N}} \sum_{i \in \mathcal{I}_{fn}} p_{ifn} q_{ifn} \leq E_g, \quad (7)$$

where $\mathcal{F}_\ell$ is the set of stores a household in ℓ can reach, $d_{f\ell}$ the distance from ℓ to store f and τ the trip cost per unit of distance, q_{ifn} the quantity of barcode i bought at store f and p_{ifn} its price there, $\mathcal{N}$ the set of categories, $\mathcal{I}_{fn}$ the barcodes on that store's shelf in category n , ν_f a Fréchet store-taste draw with shape $1/\lambda$, independent across households and stores, U_g a homothetic aggregator over the basket, and $E_g = \beta_g y_g$ the grocery budget of (5), which Cobb–Douglas shares hold fixed whatever groceries cost.

Solving the inner problem gives the unit expenditure function P_{fg} , the minimum cost of a unit of U_g at store f 's prices, so shopping at f delivers $\nu_f E_g / (e^{\tau d_{f\ell}} P_{fg})$. Maximizing over stores gives the unit cost of grocery consumption to a group- g household living in ℓ ,

$$\mathcal{P}_{\ell g} = \left[\sum_{f \in \mathcal{F}_\ell} \left(e^{\tau d_{f\ell}} P_{fg} \right)^{-1/\lambda} \right]^{-\lambda}, \quad (8)$$

with store-choice probabilities

$$s_{f\ell g} = \frac{(e^{\tau d_{f\ell}} P_{fg})^{-1/\lambda}}{\sum_{f' \in \mathcal{F}_\ell} (e^{\tau d_{f'\ell}} P_{f'g})^{-1/\lambda}}. \quad (9)$$

A tract that reaches no store has its index $\mathcal{P}_{\ell g}$ held at one, which acts as an outside option and keeps the index defined whichever stores operate.

The dispersion λ of those draws governs how much demand responds to a store's basket price. As $\lambda \rightarrow 0$ households concentrate on the cheapest basket, while larger λ spreads demand across stores. It also disciplines the supply side, since it sets how much traffic a store loses when it raises any one of its prices.

Consumption basket. The household buys the basket as a two-tier CES over the store's shelf, with elasticity of substitution σ_c across categories and $\sigma_b > \sigma_c > 1$ across the barcodes within a category, so the store index P_{fg} and the index $P_{fn,g}$ of category n at that store are

$$P_{fg} = \left(\sum_{n \in \mathcal{N}_f} P_{fn,g}^{1-\sigma_c} \right)^{\frac{1}{1-\sigma_c}}, \quad P_{fn,g} = \left(\sum_{i \in \mathcal{I}_{fn}} \tilde{\zeta}_i^{\alpha \log y_g} p_{ifn}^{1-\sigma_b} \right)^{\frac{1}{1-\sigma_b}}, \quad (10)$$

with within-category expenditure shares

$$s_{i|fn,g} = \frac{\tilde{\zeta}_i^{\alpha \log y_g} p_{ifn}^{1-\sigma_b}}{\sum_{j \in \mathcal{I}_{fn}} \tilde{\zeta}_j^{\alpha \log y_g} p_{jfn}^{1-\sigma_b}}, \quad (11)$$

where $\mathcal{N}_f$ is the set of categories store f stocks, over which every sum over n below runs, $\tilde{\zeta}_i$ is the quality of barcode i , and α sets how strongly the taste for quality varies with income. Income y_g is measured in 2025 reals per month, and the exponent is taken in those units, so $\alpha \log y_g$ gives both the weight a group places on quality and, through α , how that weight rises with income; a household with one real a month would choose among barcodes by price alone. Categories are HS-8 product categories. Quality is the only barcode-level shifter in demand, so a persistent barcode-specific premium is part of $\tilde{\zeta}_i$ and not a separate markup. Section 2 measures it as $\tilde{\zeta}_i = \exp(\tilde{\delta}_i)$, the exponential of the barcode's effect $\tilde{\delta}_i$ in a hedonic decomposition of prices, centred within its category (Appendix D).

The taste for quality enters the inner tier and not the outer one. Across categories, where a household is choosing how much rice to buy relative to coffee, substitution is constant at σ_c and the same for every group. Within a category it is choosing among barcodes that differ in quality, and how much that quality is worth to it varies with its income through the exponent $\alpha \log y_g$ on $\tilde{\zeta}_i$. The same shelf is therefore composed differently for a rich and a poor shopper at identical prices, and that is the channel through which the income composition of a neighborhood changes what a store finds worth selling.

Decomposition. These preferences let the price index be decomposed into six terms that add to it exactly, in the two groups the household's problem is made of. Getting to a store is proximity and the variety of stores it can choose among; what it meets once inside is the variety of

categories and of barcodes on the shelf, the quality of what sits there, and what that quality costs,

$$\begin{aligned}
\log \mathcal{P}_{\ell g} = & \underbrace{\sum_f s_{f\ell g} \tau d_{f\ell}}_{\text{proximity}} \\
& + \underbrace{\lambda \sum_f s_{f\ell g} \log s_{f\ell g}}_{\text{store variety}} + \underbrace{\frac{1}{\sigma_c - 1} \sum_{f,n} s_{f\ell g} s_{n|fg} \log s_{n|fg}}_{\text{category variety}} + \underbrace{\frac{1}{1 - \sigma_b} \sum_{f,n} s_{f\ell g} s_{n|fg} \log N_{fn}}_{\text{barcode variety}} \\
& + \underbrace{\frac{1}{1 - \sigma_b} \sum_{f,n} s_{f\ell g} s_{n|fg} \log \left(\frac{1}{N_{fn}} \sum_{i \in \mathcal{I}_{fn}} \tilde{\zeta}_i^{\alpha \log y_g} \right)}_{\text{quality}} \\
& + \underbrace{\frac{1}{1 - \sigma_b} \sum_{f,n} s_{f\ell g} s_{n|fg} \log \left(\frac{\sum_i \tilde{\zeta}_i^{\alpha \log y_g} p_{ifn}^{1 - \sigma_b}}{\sum_i \tilde{\zeta}_i^{\alpha \log y_g}} \right)}_{\text{price}}
\end{aligned} \tag{12}$$

where $s_{f\ell g}$ are the store-choice shares (9), $s_{n|fg} = P_{fn,g}^{1 - \sigma_c} / \sum_m P_{fm,g}^{1 - \sigma_c}$ the category shares within a store, and N_{fn} the number of barcodes on that store's shelf in category n . The price term is the quality-weighted mean price of the barcodes on the shelf, so it moves with the prices of the barcodes a store already carries and with the composition of its shelf. Splitting the count of barcodes and their average quality off from the price of what is on the shelf follows the decomposition of Hottman et al. (2016) and Redding and Weinstein (2020), which extends the variety correction of Feenstra (1994), here applied within a category and carrying the taste weights of (11). Appendix C.1 derives the decomposition.

The first two terms are the between-store block and the last four the within-store block, so the split by decision is a grouping of (12) and not a second decomposition. The three variety terms count stores, categories and barcodes, so a neighborhood has good variety access either because its stores are broad or because it reaches many of them. The quality term is an unweighted mean across the barcodes a store carries, so it measures what a household can reach and not what it buys, and the two differ whenever expenditure concentrates on a narrow part of the shelf. The proximity term is the trip cost a household pays at the shares it chooses, which is not the cost of setting τ to zero.

4.2 Supply

A store is an establishment at a site ℓ , where it occupies floorspace h_f . It chooses which barcodes to stock and what to charge for each of them, to maximize profit

$$\Pi_f = \max_{\{\mathcal{I}_{fn}\}, \{p_{ifn}\}} \sum_n \sum_{i \in \mathcal{I}_{fn}} (p_{ifn} - c_{in}) q_{ifn}(p_f; p_{-f}) - \sum_n F_n |\mathcal{I}_{fn}| - r_\ell h_f, \quad (13)$$

where $\mathcal{I}_{fn}$ is the set of barcodes carried in category n , p_f the store's own prices and p_{-f} its rivals', c_{in} the marginal cost of barcode i , F_n a stocking cost per barcode and r_ℓ the rent at the site.

Demand and costs. Aggregating the household side over groups, the quantity store f sells of barcode i is

$$q_{ifn} = \sum_g X_{fg} \frac{s_{n|fg} s_{i|fn,g}}{p_{ifn}}, \quad X_{fg} = \left(\sum_\ell N_{\ell g} s_{f\ell g} \right) E_g, \quad s_{n|fg} = \left(\frac{P_{fn,g}}{P_{fg}} \right)^{1-\sigma_c}, \quad (14)$$

where X_{fg} is what group g spends at the store, its visitors times their grocery budget, split by the category share $s_{n|fg}$ and the barcode share $s_{i|fn,g}$ of (11). The store's site enters through its visitors, whose numbers depend on the distances $d_{f\ell}$ in (9).

A price rise costs the store on three margins. It loses sales to shelfmates in the same category and to its own other categories, both through the shares in (14), and it loses shoppers altogether, since a higher P_{fg} lowers its store-choice share in (9).

Residual demand is made up of the households that can reach the store, so the income composition of the surrounding neighborhoods enters it through the group- g visitor weight

$$w_{fg} = \frac{\sum_\ell N_{\ell g} s_{f\ell g}}{\sum_{g'} \sum_\ell N_{\ell g'} s_{f\ell g'}}. \quad (15)$$

Quality valuation rises with income through (11), so $\{w_{fg}\}$ governs how far a store's demand tilts toward premium barcodes. It is the object through which residential sorting reaches the shelf, closing the loop back to Section 4.1.

Marginal cost varies with the category and with barcode quality, and is the same at every store,

$$c_{in} = c_0 \bar{c}_n \tilde{\zeta}_i^\phi, \quad (16)$$

with $\bar{c}_n$ a category cost shifter, ϕ the elasticity of marginal cost to quality and c_0 a common cost level. Rent does not enter it, since a store pays rent on its footprint whatever it sells, so floorspace costs enter the decision to operate and not the cost of an additional unit.

Prices. Writing the markup as $L_{ifn} = (p_{ifn} - c_{in})/p_{ifn}$, the condition for barcode i is

$$L_{ifn} = \frac{1 + C_{ifn}^{\text{cat}} + C_{ifn}^{\text{store}} - \mathcal{T}_{ifn}}{\Delta_{ifn}}, \quad (17)$$

where Δ_{ifn} is the elasticity of the barcode's own quantity to its own price. A store that ignored the sales its other barcodes recapture would charge the inverse of it. Taking the recapture into account, it charges more, because the sales a price rise drives away do not all leave the store. C_{ifn}^{cat} is the cannibalization it recovers on shelfmates in the same category and C_{ifn}^{store} the cannibalization it recovers in its other categories. $\mathcal{T}_{ifn}$ is the margin it gives up on those other categories when the price rise sends shoppers to a different store. Under the sign conditions (45), which hold at the estimates, cannibalization widens the margin and lost traffic narrows it. Appendix C.2 derives the condition and gives them in closed form.

Assortment and entry. A store stocks barcode i when doing so adds more to profit than it costs to stock, counting what the barcode earns inside its category, the spending it pulls into the category from the store's other categories, and the shoppers it draws to the store. Inside the category the addition falls short of the barcode's own margin revenue, since a new barcode takes sales from the shelfmates already there. A longer shelf also lowers the category's price index, which moves the store's own customers' spending toward the category, a gain when the category's margin exceeds the store's average. And it lowers the store's price index and so raises its traffic, which earns the margin on everything else the store sells. Writing $\mathcal{S}$ for the set already carried, which excludes i when the store carries it and is the current shelf when it does not, and $X_{nfg} = X_{fg} s_{n|fg}$ for what group g spends in category n at the store,

$$\Delta\pi_f(i) = \sum_g X_{nfg} s_{i,g}^{\text{post}} \left[(L_{ifn} - \bar{L}_{fn,g}) + L_{fn,g}^{\text{cat}} + L_{fg}^{\text{traffic}} \right] \geq F_n, \quad (18)$$

where $s_{i,g}^{\text{post}}$ is the share barcode i would hold once added, at the price (17) assigns it if the store does not yet carry it, $\bar{L}_{fn,g} = \sum_{k \in \mathcal{S}} L_{kfn} s_{k|fn,g}$ the share-weighted mean markup of what the store already carries in the category, and

$$L_{fn,g}^{\text{cat}} = \frac{\sigma_c - 1}{\sigma_b - 1} (\bar{L}_{fn,g} - \bar{L}_{f,g}), \quad L_{fg}^{\text{traffic}} = \frac{\bar{L}_{f,g} (1 - \bar{s}_{f,g})}{\lambda (\sigma_b - 1)}, \quad (19)$$

with $\bar{L}_{f,g} = \sum_{n'} s_{n'|fg} \bar{L}_{fn',g}$ the store's spending-weighted mean markup for group g , and $\bar{s}_{f,g}$ the store's average share of group g 's shoppers across the neighborhoods it draws from, by (9). Spending drawn from the store's other categories is worth the category's margin less the margin

it leaves behind, which $L_{fn,g}^{\text{cat}}$ records, and it vanishes when every category carries the store's average markup. The shoppers a longer shelf draws spend across the whole store rather than on the barcode that drew them, in their own group's shares, so what they are worth is $\bar{L}_{f,g}$, and L_{fg}^{traffic} states that in the same units as the markup difference beside it. A store adds a barcode when its markup beats the average markup of its shelf by enough to cover the stocking cost, scaled by the share it would capture, with the spending it pulls into the category and the traffic it brings in credited against the shortfall. Appendix C.3 derives the three terms. The condition governs the barcodes a store carries within the categories it stocks; which categories it stocks is taken from the observed shelf.

The post-entry share carries the taste weight of (11), so a premium barcode earns its place on a shelf where the households the store draws are rich enough to buy it. The set that clears (18) shifts up the quality distribution as the surrounding neighborhoods get richer, which is the mechanism behind the variety elasticity of Section 3 and the only place in the model where local income composition changes what is on the shelf rather than what it costs.

A site is a place a store can occupy, and a tract can hold several. A potential entrant is attached to each site, and it operates when $\Pi_f \geq 0$ in (13) and stays out otherwise. What it would earn depends on the stores already in the city, through p_{-f} and the residential allocation, so which sites are taken is determined jointly with the prices and shelves on them.

4.3 Floorspace market clearing

Households and stores compete for the same floorspace, which ties the two sides of the model together outside the grocery market. Households spend a share ρ_g of income on housing and a store occupies h_f at its site, while developers supply floorspace at a constant elasticity $\eta > 0$. Floorspace is owned by absentee landlords, so every household rents and the rents it pays do not return to the city's households as income, as in Ahlfeldt et al. (2015). Land also has a use outside the city's housing and retail, worth $\underline{r} > 0$, which no owner will supply below. The rent in ℓ is therefore

$$\bar{H}_\ell \left(\frac{\tilde{r}_\ell}{r_\ell^0} \right)^\eta = \sum_g \frac{\rho_g y_g N_{\ell g}}{\tilde{r}_\ell} + H_\ell^{\text{firm}}, \quad r_\ell = \max\{\underline{r}, \tilde{r}_\ell\}, \quad (20)$$

with $\bar{H}_\ell$ the floorspace stock at a reference rent r_ℓ^0 and H_ℓ^{firm} the floorspace the stores operating in ℓ occupy. The floor binds only where housing and retail together bid less than the outside use. Rent reaches retail only as the overhead of holding a site in (13), so retail meets the floorspace market through entry alone. Appendix C.4 shows the rent is unique.

4.4 Equilibrium

Given the parameters $(\sigma_b, \sigma_c, \lambda, \alpha, \kappa_\varepsilon, \beta_g, \rho_g, \tau, \eta, \underline{r})$, the barcode qualities $\{\xi_i\}$ and the cost shifters of (16), the location fundamentals $\{A_{\ell g}, \bar{H}_\ell, d_{f\ell}\}$, the overheads $\{F_n, h_f\}$ and group incomes and populations $\{y_g, N_g\}$, an equilibrium is a set of operating stores, assortments $\{\mathcal{I}_{fn}\}$, prices $\{p_{ifn}\}$, store-choice shares $\{s_{f\ell g}\}$, rents $\{r_\ell\}$ and location shares $\{s_{\ell g}\}$ in which (i) households solve the store-choice and grocery problem and the housing allocation; (ii) stores' prices satisfy the first-order conditions (17) at marginal costs (16); (iii) every shelf satisfies (18), so that no single addition or deletion improves it, every operating store earns $\Pi_f \geq 0$ in (13), and no potential entrant at an inactive site would earn $\Pi_f > 0$ by operating; (iv) location shares satisfy the Fréchet sorting (6); and (v) (20) holds in every neighborhood. Appendix C.4 establishes that an equilibrium exists conditional on the assortment and the set of operating stores, and that without conditioning on them one exists in mixed strategies, with condition (iii) holding in expectation, under sign restrictions on the pricing elasticities that hold at the point estimates (Remark 1). With the retail side held fixed, rents and the allocation of households are unique for every parameter value, so any multiplicity runs through the response of prices, shelves and entry.

5 Estimation

The model meets the data through the two conditions it imposes on the objects the retail data record. Prices must be consistent with optimal multiproduct pricing (17), and shelves with optimal stocking against the income mix of the customers each store draws (18). The retail data record no quantities, so the demand side is estimated on the household consumption survey.

Two parameters are recovered from the city, and each has its own equation. How strongly the taste for quality rises with income, α , comes from the assortment condition, where it governs which barcodes a store finds worth carrying given who shops there. The store-taste dispersion λ comes from the level of markups, anchored to an external measurement of retail margins.

The rest of the demand system is settled elsewhere. The two elasticities are estimated on a household consumption survey before the objective is formed and held fixed, the sorting dispersion is re-estimated on a national cross-section at each candidate λ , and the cost-quality exponent follows from splitting the price-quality slope into a cost part and a markup part. Each subsection below gives the reason for one piece of that division, and Table 4 collects them.

5.1 Elasticities of substitution

I estimate both elasticities on the food diary of the household consumption survey, which observes the quantities the retail data do not. The diary records the weight purchased alongside the expenditure, so every product has an observed unit value in each region. The two elasticities

differ only in the tier they are estimated at. For the upper tier, varieties are product categories and the estimand is substitution across them. For the lower tier, varieties are the product types within a category, the types of rice or of milk the survey distinguishes, and the estimand is substitution among them. Both use the double difference of Feenstra (1994) as refined by Soderbery (2015), which identifies the demand and supply elasticities jointly from the second moments of differenced residuals. Appendix D gives the estimator and the sample restrictions.

5.2 Taste for quality

How strongly the taste for quality rises with income is identified from the assortment condition, where α has a signature the pricing residual does not contain. A store carries the barcodes whose incremental contribution to profit covers the stocking cost, and by (18) that contribution runs through the share a barcode would hold once added, which carries the taste weight. Raising α makes premium barcodes worth carrying where a store's customers are rich and not worth carrying where they are poor. The shelf-composition moment is the score of the probability that the observed shelf is one no single addition or deletion improves (Proposition 1), so for each store and category it compares the barcodes actually carried against the model's probability of carrying them, weighted by the derivative of the log incremental profit with respect to α . It uses the store–category cells with at least 5 carried barcodes, since a shorter shelf has little composition to explain, and the barcodes a store could carry are those in the sample in the categories it already stocks. The parameter fits the covariance, across store–category cells, between the quality composition of a shelf and the income composition of the store's customers.

The moment does not depend on the level of the stocking cost or on the sorting block. The stocking cost is common within a store's category up to a store- and barcode-specific draw (40), and its common part is solved in each cell so that the expected number of barcodes carried equals the observed one. The count pins the cost, and the composition of the shelf, given that threshold, identifies α . The customers' income mix is built from the observed residential allocation, so α does not route through the sorting block.

The identifying assumption is that stores have no unobserved taste for quality that varies across the quality distribution within a category. A neighborhood amenity that makes a store's whole shelf better raises everything on it and cancels from a comparison of barcodes within one store and category. A store-level preference that tilted specifically toward premium barcodes would not cancel.

5.3 Store-taste dispersion

The store-taste dispersion is identified by the level of markups. A store that raises a price loses traffic at a rate λ sets, so the wider the dispersion of store tastes the less a price rise costs and the higher the margin a store holds. By Lemma 3, if a store's traffic elasticity were the same

across its shelf, every margin there would equal $1/(1 + \bar{\Omega}_f)$, set by λ and the store's share of the shoppers it draws. That share, like the weights on each income group, comes from the model's store-choice probabilities at observed prices and the observed residential allocation, since store visits are not observed. Differences in who buys each barcode spread margins around that level, and at the estimates they spread little, from 0.224 at the tenth percentile to 0.247 at the ninetieth, so the level of markups is governed by λ .

The level is anchored outside the sample. The hedonic decomposition recovers marginal cost only up to a normalization, so the common cost level is the single constant that sets the mean log markup over the sample's barcode-establishment cells to the commercial margin IBGE reports for grocery retail (33.4 percent over cost of goods sold, Section 2). The markup-level moment then requires the model's mean margin, $(p - c)/p$ averaged over the same cells, to equal the data's at that cost level. Margins vary across cells, so their mean, 0.2324, lies below the margin implied at the mean log markup. The dispersion is thus read from an external measurement of how much retailers mark up, mapped through the pricing conditions into how readily shoppers switch stores. The benchmark is a ratio of activity aggregates and the calibration treats it as a mean over cells, and it enters as a fixed input, so the estimate of λ and its standard error are conditional on it.

5.4 Sorting dispersion

I estimate the sorting dispersion on the national 2010 census, whose weighting areas are the finest unit at which it reports both income groups and rent. The unit is coarser than the model's census tract, and the groups are within-city quartiles of household income per capita instead of household income, so the estimate transfers to the model's tracts and groups under the assumption that the response of location shares to access depends on neither. The parameter governs how far an income group shifts across neighborhoods when consumption access differs, so it is read from a regression of income-group shares on access, with local rent held fixed at its own coefficient and city-by-group fixed effects absorbing whatever distinguishes one city from another.

The national register records where grocery stores are, not what they charge or stock, so access in that regression is the model's consumption value with every store's basket index set equal, $\lambda \log \sum_b S_b e^{-(\tau/\lambda)d_{ab}}$, summed over the weighting areas b of the city, where S_b counts area b 's grocery establishments weighted by the mean number of barcodes their size class carries in the Maceió data and d_{ab} is the distance between area centroids. It keeps the proximity and store-variety margins of (12) and drops the within-store ones. The regression is re-evaluated at every λ the estimation visits, so $\hat{\kappa}_\varepsilon$ is the value consistent with $\hat{\lambda}$. Least squares would confound the two directions of that relationship, since stores open where the customers they want already live and any amenity that draws both stores and richer households would be read as a response to access.

The instrument is the neighborhood's grocery stock in 1995, reconstructed from the establishment register. Retail location is persistent, so where shops stood a generation earlier still shifts how reachable groceries are in 2010, and that stock was in place before the income distribution the regression explains. The parameter is therefore read from the part of today's access that a neighborhood inherited from a retail geography laid down before today's residents sorted into it.

The comparison runs between neighborhoods of the same city that share a terrain, a distance to the centre and a rent. Slope and mean elevation enter from the digital elevation model, and they hold fixed the features of a site that keep it desirable over decades. The exclusion restriction is that among such neighborhoods, where grocery shops stood in 1995 affects the income composition of 2010 only through the access those shops leave behind. Appendix D reports the first stage, the weak-instrument-robust confidence set, and the estimate with the 2000 stock as the instrument.

Within the city, amenities are inverted group by group, so the model reproduces the observed income composition of every neighborhood by construction. That fit is exact before any moment is formed, which is why the parameter is read from the national cross-section.

Table 4: Where each parameter is identified

Parameter	Moment	Data
<i>Estimated on the city, inside the GMM objective</i>		
λ , store-taste dispersion	<i>Markup level.</i> The model's mean markup set against the margin grocery retailers report earning	Price data; Annual Survey of Commerce
α , how the taste for quality rises with income	<i>Shelf composition.</i> The barcodes a store carries against the model's probability of carrying them, the score of (18)	Price data; 2022 Census
<i>Re-estimated at each candidate λ, inside the objective</i>		
κ_ε , sorting dispersion	Income-group shares on consumption access, instrumented by the neighborhood's grocery stock in 1995	2010 Census, national; firm registry
<i>Estimated before the objective is formed and held fixed</i>		
σ_c , elasticity across categories	Feenstra–Soderbery double difference of expenditure shares and unit values across regions, (47)	Consumption survey
σ_b , elasticity within a category	The same first stage run across the product types within a category	Consumption survey

Notes: The top panel holds the parameters the city identifies, fitted by matching the two moments named in italics. Consumption access is built at λ , so the sorting cross-section is re-run at every candidate value and κ_ε is its value at $\hat{\lambda}$; it fits no moment of the city's own. The elasticities are first stages, held fixed while the objective is minimized, whose sampling uncertainty is propagated into the standard errors of Table 5. Each subsection above gives the reason for one row.

5.5 Estimator and inference

I estimate the 2 city parameters by GMM on the 2 moments of the subsections above. Writing $m(\theta)$ for the vector that holds them at a candidate θ , the markup-level gap and the shelf-composition score, $\hat{\theta}$ minimizes $m(\theta)'Wm(\theta)$. There are as many moments as parameters, so both moments are set to zero, any positive-definite weight matrix W returns the same $\hat{\theta}$, and the fitted moments provide no over-identification test.

Standard errors come from the GMM sandwich, with the moment covariance estimated by a block bootstrap over establishments. Both elasticities are fixed first stages, and their sampling uncertainty is propagated into the reported standard errors. That term dominates the sampling variation of the city moments, by more than two orders of magnitude for the taste for quality.

The cost–quality exponent is not estimated. Equation (48) splits the observed price–quality slope into what quality costs and what it earns in margin, and the margin part is a function of the demand primitives, so the cost part follows from them with a delta-method standard error. Section 6.1 reports it with the parameters it is derived from, and Appendix D reports the checks the estimated model is not fitted to.

6 Results

The grocery environment of a richer neighborhood is worth more, and the stores a household can reach are the source of its advantage. With respect to neighborhood income, the between-store block of (12) has an elasticity of -0.158 and the within-store block $+0.031$, so the grocery price index has an elasticity of -0.128 .

The margin that works against the household in a richer neighborhood is the price it pays once inside a store. A given amount of grocery utility costs more to buy where incomes are higher, for every income group, once the quality and variety of what is on the shelf are netted out. That price is a markup over what the good costs to put on the shelf, and Section 6.3 decomposes it into market power over the barcode, the substitution a store internalises within and across its own categories, and the traffic it loses when it raises a price.

6.1 Estimates

The two demand elasticities in Table 5 are $\hat{\sigma}_b = 5.550$ within a category and $\hat{\sigma}_c = 4.700$ across categories. The within-category elasticity sets the residual elasticity Δ_{ifn} that Section 6.3 decomposes.

Richer households value quality more, at $\hat{\alpha} = 0.509$. A barcode’s share of its shelf carries the exponent $\alpha \log y_g$ on quality by (11), so for a household one log point richer a barcode one log point higher in quality gains as much shelf share as a price cut of 0.112 log points would give

Table 5: Estimated structural parameters

Parameter	Interpretation	Value
<i>Panel A. Demand block</i>		
σ_b	Within-category (barcode) elasticity (first stage)	5.550 (2.418)
σ_c	Across-category elasticity (first stage)	4.700 (0.255)
α	Income gradient of the taste for quality	0.509 (0.283)
λ	Store-taste dispersion	0.288 (0.005)
ϕ	Cost–quality exponent (derived)	1.008 (0.003)
<i>Panel B. Spatial equilibrium</i>		
κ_ϵ	Location-choice (Fréchet) dispersion	6.189 (1.465)
τ	Trip cost per km (calibrated)	0.150
η	Housing supply elasticity (calibrated)	0.300

Notes: Standard errors in parentheses. α and λ are estimated jointly by GMM. σ_b and σ_c are fixed first stages from the consumption survey, whose sampling uncertainty is propagated into the standard errors of α and λ . The cost–quality exponent ϕ follows from (48) and the Fréchet dispersion κ_ϵ from the cross-section of Section 5 at $\hat{\lambda}$, each with a delta-method standard error. τ and η are calibrated.

it. The tilt is large enough to change which barcodes a store finds worth stocking, which is the channel Section 6.2 measures as the quality margin.

The store-taste dispersion $\hat{\lambda} = 0.288$ is the more precisely estimated of the two GMM parameters. It sets how much traffic a store loses when it raises a price, and so the size of the loss-leader force that Section 6.3 measures. It also multiplies the store-variety term of (12), which accounts for more than the whole elasticity of the between-store block, so it scales the headline elasticity above.

Households respond to consumption access and rents when they choose a neighborhood, but not deterministically. The sorting dispersion $\hat{\kappa}_\epsilon = 6.189$, common across income groups, sits within the range of the Fréchet dispersions of residence and workplace choice estimated in quantitative models of cities (Monte et al., 2018; Heblich et al., 2020; Ahlfeldt et al., 2015; Tsivanidis, 2026). It is identified on a national cross-section rather than within this city, and the sorting magnitudes of Section 8.2 are proportional to it. The consumption-value gains reported there depend on it only through the re-sorting that moves stores’ customers, and so their prices and shelves.

What a household pays for quality is close to what quality costs. The price–quality slope splits into what quality costs and what it earns in margin, (48). The pricing conditions give a markup share of $\hat{\Gamma} = 0.0007$ and hence a cost share of $\hat{\phi} = 1.008$. A barcode priced above its category mean therefore costs very nearly that much more to put on the shelf and carries no

wider margin for being premium. Section 6.3 shows where the flat markup comes from in the composition of markups.

The per-kilometer trip cost $\tau = 0.150$, from Miyauchi et al. (2025), and the housing supply elasticity $\eta = 0.300$ are calibrated. Beyond these scalars the estimation recovers quality shifters ζ_i for 122,427 barcodes, HS-8 cost shifters $\bar{c}_n$ for 859 categories, and neighborhood amenities $A_{\ell g}$ for all 1,725 tracts in each of the four income groups. Appendix D.7 puts the markups these estimates imply against the data by store size, a dimension no moment targets.

6.2 Decomposing the grocery price index

A household in a richer neighborhood is better off in groceries because of the stores it can reach, not because of what it meets once it is inside one. Table 6 reports the elasticity of every term of (12) with respect to neighborhood income, grouped into what a household settles before it enters a store and what it meets inside, and Figure 6 plots the two blocks. Getting to stores carries an elasticity of -0.158 and what is inside them $+0.031$, and the two add to -0.128 . A neighborhood one percent richer than another faces a grocery price index 0.128 percent lower. The terms add to $\log \mathcal{P}_{\ell g}$ neighborhood by neighborhood and group by group, and that is a log, so their elasticities add to the elasticity of the whole.

Table 6: Decomposition of the grocery price index, by block and term

Term	Elasticity	Share of total
<i>Between-store</i>	-0.158	124.1%
Proximity	$+0.010$	-7.8%
Store variety	-0.168	131.9%
<i>Within-store</i>	$+0.031$	-24.1%
Category variety	$+0.009$	-7.2%
Barcode variety	$+0.006$	-4.4%
Quality	-0.049	38.6%
Price level	$+0.065$	-51.1%
Total ($\log \mathcal{P}_{\ell g}$)	-0.128	100%

Notes: Each entry is the coefficient from a projection of one term of (12) on log neighborhood income. The terms sum to $\log \mathcal{P}_{\ell g}$ tract by tract and group by group, so the total is an identity. A negative entry means the term makes groceries cheaper in richer neighborhoods. Indented rows add to the block above them up to rounding, and every share is of the total.

Store variety accounts for the entire between-store block. Residents of richer neighborhoods effectively choose among more stores; their walks are no shorter. The term is the effective number of stores a household chooses among, the exponential of the entropy of its store-choice shares, so it rises both with the number of stores in reach and with how evenly distance and prices spread spending across them. The size of that term depends on store tastes being drawn independently across stores. Each store within reach then brings a fresh draw, and N otherwise identical stores lower the index by $\lambda \log N$. Tastes correlated across nearby or same-format stores




Figure 6: The two blocks of the decomposition, by neighborhood income

Notes: The two blocks of (12) in price-index terms, so a falling line means groceries are cheaper in richer neighborhoods, each indexed to 100 at the poorest bin. Markers are means over equal-count bins of log neighborhood income with 95% confidence intervals, and the lines are OLS projections fitted on the neighborhood-level data. Labels report each block's elasticity. Table 6 gives the terms inside each block.

would make an added store a closer substitute for those already in reach and the store-variety term smaller (Akerberg and Rysman, 2005). The level of markups identifies λ through how many shoppers a store loses when it raises its prices, but not through where they go, whether to a similar store nearby or to any store in reach, and it is the second that the correlation governs. Inside the store, the category and barcode variety terms are positive, so the shelves households meet, weighted by where they shop, are no deeper in richer neighborhoods. The barcode counts of Section 3 rise with income because they count every distinct barcode a neighborhood's stores sell between them, which grows with the number of stores and which the decomposition records as store variety. The barcode variety term measures how many barcodes each category holds at the stores households choose, and that number does not rise with income.

The lower index of richer neighborhoods comes from which stores a household can reach. Store entry changes the cost of living (Atkin et al., 2018), rich and poor households buy from different stores and firms (Faber and Fally, 2022), and travel costs shape where within a city households shop (Davis et al., 2019; Miyauchi et al., 2025). Corrections for new-variety and quality bias, in the tradition of Feenstra (1994) and Handbury and Weinstein (2015), operate on the goods a household buys and so belong to the within-store block, whose elasticity is +0.031. The between-store block measures the set of sellers a household can reach, and its elasticity of -0.158 is 124.1 percent of the total.

The within-store block is the store price index a household faces, and it is where the income groups separate. Its four terms are defined store by store rather than neighborhood by neighborhood, so aggregating them with the choice shares s_{flg} and projecting on log neighborhood income preserves the identity. Appendix C.1 gives the algebra. Store choice depends on the group only through that index, so proximity and variety are all but identical across income quartiles

while the index itself is not. The elasticity of the grocery price index with respect to neighborhood income runs from -0.122 for the poorest quartile to -0.135 for the richest. Q4 households face a grocery price index of 1.089 against 0.880 for Q1, but each group's index weights barcodes by its own taste for quality, set with the income normalization of Section 4, so the levels are not comparable across groups and the groups are compared through their elasticities.

Most of the block is the price level, and it is the term that separates the groups. Table 7 reports the four terms by household income quartile, in the same signs. Every group faces a more expensive store index in a richer neighborhood. The price level steepens with income, from 0.0501 for the poorest quartile to 0.0845 for the richest, so once quality is netted out a given amount of grocery utility costs the richest households the most to buy where incomes are high.

Table 7: The within-store block, by household income quartile

Term	Q1	Q2	Q3	Q4
Price level	0.0501	0.0596	0.0668	0.0845
Barcode variety	0.0085	0.0061	0.0046	0.0033
Quality	-0.0376	-0.0434	-0.0489	-0.0670
Category variety	0.0128	0.0104	0.0086	0.0047
Within-store block	0.0338	0.0327	0.0311	0.0255
Shelf quality $\log \mu$	0.1711	0.1976	0.2224	0.3047

Notes: The within-store block of Table 6, by income group, as elasticities with respect to log neighborhood income and aggregated to neighborhoods with the model's choice shares. Signs are those of a price index, as in Table 6, so a positive entry means the term makes a store in a richer neighborhood more expensive. The final row is not a term of the block but the object its quality term is built from, the average quality of the barcodes on the shelf in the group's own valuation.

The block total runs the other way, falling from 0.0338 to 0.0255. The quality term, the only one working in the household's favour, runs from -0.0376 for the poorest quartile to -0.0670 for the richest, so the same improvement in shelves is worth progressively more as one moves up the income distribution. It offsets nearly all of the rise in the price level across groups, and the two variety terms, which shrink across groups, account for the rest of the fall in the total.

The last row of Table 7 gives the object the quality term is built from. Shelf quality $\log \mu$ in a group's own valuation rises with neighborhood income more steeply the richer the group, from 0.1711 for the bottom quartile to 0.3047 for the top. A richer neighborhood offers better shelves to every group and offers them most to the group that values quality most, which is the income tilt α of Section 6.1 appearing on the supply side. The price level exceeds the quality term in every group, so better shelves do not offset what they cost, which is why the within-store block is positive in Table 6.

The decomposition has two properties that matter for how the grocery environment and a neighborhood's income mix respond to each other. The advantage of a richer neighborhood comes from what its spending supports, the stores within reach and the shelves they carry, so

a change in the income mix can change the grocery environment, and Section 8.3 measures by how much. There that change runs mostly through the shelves, through barcode and category variety, and little through the number of stores. And the advantage is nearly the same for every income group, so an improvement in the grocery environment is worth about as much to a poorer household as to a richer one. Sections 8.2 and 8.3 measure each direction.

6.3 Decomposing markups

The mean model markup, $(p - c)/p$, is 0.2324. Marginal cost is what the good costs to put on the shelf. Labour, rent and other operating costs enter the model as fixed costs, so this is a gross margin. Section 5 anchors that level outside the sample. The estimates determine its composition, and Table 8 reports it at $\hat{\theta}$.

Table 8: Decomposition of the markup

Term	Mean	Median	S.d.	% of total
Simple markup $1/\Delta_{ifn}$	0.183	0.181	0.006	78.6
Within-category cannibalisation	0.033	0.034	0.008	14.3
Cross-category substitution	0.152	0.151	0.019	65.5
Store traffic (loss-leader)	-0.136	-0.139	0.012	-58.5
Total, $(p - c)/p$	0.232	0.227	0.022	100.0

Notes: The pricing condition (17) evaluated at $\hat{\theta}$ across 1,561,010 barcode–store observations, one row per term. The first three raise the markup, the first through market power over the barcode itself and the other two through the store internalising substitution toward its own products within and across categories. The store-traffic term is the loss-leader force and is the only one that lowers it. Shares are of the total, so the terms sum to 100 percent, and the three that raise the markup together exceed it because the store-traffic term offsets them. 92 observations with $L \notin (-1, 1)$ are excluded.

The estimated margin is the net of forces considerably larger than itself. Market power over the barcode, the simple term $1/\Delta_{ifn}$, supplies 78.6 percent of it. A store internalises that raising one price pushes buyers onto its own close substitutes, so it loses less demand than a single-product seller would and within-category cannibalisation supplies 14.3 percent. The same logic across categories supplies 65.5 percent. The three sum to more than the total, and the store-traffic term absorbs the difference.

Store traffic is the only force in (17) that lowers the markup, and it is large. Cutting a price draws shoppers to the store and lifts sales of everything else, which is worth -0.1358 , or -58.5 percent of the total. Its size is governed by $\mathcal{T}_{ifn}$, which acts on the same aggregate as cross-category substitution and against it, at a rate $(1 - s_{f\ell g})/\lambda$. The loss-leader piece scales with $1/\lambda$, so the store-taste dispersion decides whether footfall disciplines pricing. At $\hat{\sigma}_c = 4.700$ and $\hat{\lambda} = 0.288$ the two forces are of comparable magnitude, so substitution dominates on net without overwhelming the traffic channel.

Along the quality distribution the total moves from 0.2334 in the bottom decile to 0.2291

in the top. Premium barcodes hold smaller shares of their own shelf, so the residual elasticity Δ_{ifn} is larger by (35) and the simple term moves from 0.1832 to 0.1817. A smaller own share also leaves more of the category with its shelfmates, so cannibalisation, measured against the barcode's own revenue, rises from 0.0329 to 0.0339. Most of the small fall in the total comes from the cross-category term, which moves from 0.1533 to 0.1487, while the traffic term moves from -0.1361 to -0.1352 . The total is close to flat, which is why the markup share of the quality price premium in (48) is small.

This composition governs how prices move in the counterfactuals of Sections 8.2 and 8.3. An entrant changes the set of stores a neighborhood can reach, and incumbent prices respond through the traffic term and through the shares that set Δ_{ifn} and the two internalisation corrections. Both channels are estimated here, so the price responses reported there are a property of the estimated first-order condition.

7 The cost of unequal access to stores

Unequal access to stores costs the poorest households the most. If every household could shop at every store in the city at no trip cost, the poorest quartile of households would gain 12.3 percent of consumption, the richest 6.8 percent and the average household 10.5 percent (Table 9). The ratio of the richest quartile's real consumption to the poorest's would fall by 5.1 percent.

The benchmark changes only which stores households can use. Every household can reach every store in the city and the trip cost τ is zero. The reachable set $\mathcal{F}_\ell$ and the distances then drop out of (8), and each income group faces the same grocery price index in every neighborhood. Prices, each store's shelf and neighborhood amenities stay at their estimated values, so the benchmark values access to the stores the city already has. Households re-sort across neighborhoods and rents clear, with 77 percent of a shift in housing demand passing into rent. Amenities are inverted once at the baseline and held fixed, which attributes the reallocation to access (Appendix E).

The gain falls with income because poorer households spend a larger share of their budget on groceries and live where access is worse. A household that stays in its neighborhood gains its grocery budget share β_g times the fall in its grocery price index, and both fall with income (Table 9, lower panel). The poorest quartile's index falls by 65 percent. At the population's mean budget share the gain still falls with income, by less, and budget shares account for 76 percent of the log difference between the poorest and the richest quartile's gain.

Most of the gain comes from reaching more stores and not from shorter trips. Removing the trip cost while each neighborhood keeps the stores it already reaches gains the poorest quartile 2.6 percent, against 12.3 percent when every store comes within reach. In the cross-section of Table 6, likewise, the effective number of stores lowers the index of richer neighborhoods and proximity does not.

Table 9: The value of equal access to stores, by household income quartile

	Q1 (poorest)	Q2	Q3	Q4 (richest)
<i>Welfare</i>				
Consumption equivalent (%)	12.3	11.4	10.3	6.8
Change in log utility	0.116	0.108	0.098	0.065
at baseline rents	0.123	0.113	0.102	0.065
Trip cost removed only (%)	2.6	2.5	2.3	1.7
<i>Households in place</i>				
Fall in the grocery price index, $-\Delta \log \mathcal{P}_{\ell g}$	1.057	1.043	1.027	0.913
Grocery budget share, β_g	0.110	0.104	0.096	0.069
Consumption gain (%)	12.4	11.5	10.3	6.5
at the mean budget share (%)	10.8	10.7	10.5	9.3

Notes: Changes between the estimated equilibrium and the equilibrium in which every household can reach every store at zero trip cost, after households re-sort and rents clear. Q1 to Q4 are household income quartiles. The consumption equivalent is the proportional change in income that changes a group's expected utility by as much as equal access does. The third row evaluates the same change at baseline rents. The fourth row sets the trip cost to zero while every neighborhood keeps the stores it reaches at baseline. The lower panel averages over each group's baseline locations. The consumption gain is $\exp(-\beta_g \Delta \log \mathcal{P}_{\ell g}) - 1$, and the last row replaces β_g with its population mean, 0.097.

Equal access lowers inequality in real consumption by more than unequal access adds to it. Real consumption is income deflated by the neighborhood's grocery price index. Unequal access raises its Gini 0.42 percent above the Gini of nominal income, and equal access changes it by -2.36 percent, so removing the added gap accounts for 17.87 percent of the change. The rest comes from the level of the gain, which every household receives and which is worth more to poorer households.

Rents take back part of the gain where it is largest. For a household of type g that starts in neighborhood ℓ and stays there, every term of (5) except the grocery price index and rent is unchanged, so its change in log utility is

$$\Delta \log u_{\ell g} = \underbrace{-\beta_g \Delta \log \mathcal{P}_{\ell g}}_{\text{consumption gain}} - \underbrace{\rho_g \Delta \log r_{\ell}}_{\text{rent}} \quad (21)$$

with β_g running from 0.110 for the poorest quartile to 0.069 for the richest and ρ_g from 0.238 to 0.174. Households move toward the neighborhoods whose access improves most, which are the poorer ones, so rent rises there and falls in richer neighborhoods (Figure 7). Across neighborhood and income-group cells, each log point of consumption gain raises the rent take-back by 0.361 log points, and rents compress the dispersion of the gain by 33.7 percent. Population is fixed, so rent rises in some neighborhoods as it falls in others, and the population-weighted rent term averages -0.001 log points against a consumption gain of 0.100. Once rents adjust, 100.0 percent of households are better off in place.

Equal access moves 4.0 percent of households and leaves the sorting of income groups across




Figure 7: The change in log utility of households who stay in their neighborhood under equal access, split as in (21) into the consumption gain, the change from rent, which is negative where rent rises, and their sum, against the neighborhood’s baseline income implied by its composition. One line per household income quartile; binned scatterplot as in Figure 6.

the city unchanged. The dissimilarity index between the richest and the poorest quartiles goes from 0.631 to 0.632, and the slope of the richest quartile’s neighborhood share on neighborhood income from 0.352 to 0.358.

8 Grocery retail and neighborhood change

Grocery retail and the income mix of a neighborhood respond to each other, and two experiments on the same neighborhood measure each direction. The first opens a high-end supermarket and measures the income mix that follows. Each percent fall in the location price index raises mean income by 0.270 percent, as richer households move in and the poorest quartile’s population falls. The second moves higher-income households in and measures the grocery environment that follows. A store opens, the shelves residents reach lengthen and move up in quality, and the index moves by -0.164 percent per percent of mean income.

When richer households arrive, the placement and the rents it produces account for the change in who lives in the neighborhood, and the retail response slightly offsets it. Solving the same placement with every store held fixed shows that the retail response reduces the rise in mean income by 1 percent of it. The response raises the welfare of incumbents in every income group by a similar amount, so more of each group stays, and the incumbents it keeps are poorer than the households the placement brings.

8.1 The experiments

Both experiments act on one treated area, so they measure the two directions of the same relationship. The area is centred on a middle-income, populated tract that reaches few barcodes and lies within a shopping radius of the top income quartile, where a change in composition is plausible. It grows outward from that site until it holds 3000 households or reaches 3.0 kilometers, and at this site the radius binds first, leaving the 12 nearest populated tracts and 2715 households within 2.96 kilometers. Appendix E gives the screen.

The first experiment forces open a clone of an observed high-end supermarket at the treated tract nearest the site. The clone copies the store with the highest mean shelf quality among those at or above the 75th percentile of the city in both the categories and the barcodes they carry and at or above the 25th percentile in floorspace. The store re-solves its prices and keeps its shelf, since the shelf is the treatment. The second experiment places top-quartile households in the treated area, a number equal to a share 0.10 of each treated tract's population. They come from the 287 tracts of above-median income whose poorest quartile has grocery access in the top 33 percent, in proportion to each tract's top-quartile population. The same number of each treated tract's residents, in proportion to its income mix, move to the rest of the city. The placement adds no housing units, so the shock itself changes who lives in the treated area and not how many.

After either shock every other margin responds. Incumbents re-price and re-solve their shelves at the income mix and spending of their customers. At each treated tract the 3 store formats whose profit a richer income mix raises most are planted as candidate entrants, the same in both experiments, and each opens if the change in its profit covers the change in its rent bill. Incumbents may close, rents clear, and every household not placed re-sorts, so the treated area's population and composition at the end are outcomes. The location price index $\mathcal{P}$ is averaged over the treated area at baseline population weights, so its change is not confounded with a change in who faces it. Welfare is the consumption equivalent $-\beta_g \Delta \log \mathcal{P} - \rho_g \Delta \log r$, from (5), of a household that stays in its tract, split into the gain from the grocery index and the loss to rent. Appendix E gives each margin and the solution method.

The experiments measure the retail channel alone. The location amenity is held fixed, so they leave out everything else a new store or a change in composition does to a neighborhood. The store of the first experiment is held open whatever it earns, and its absence from the observed city says it would not cover its costs unaided, so it describes a supported store whose support is not measured. Entry is evaluated for the planted candidates, each held to the most permissive overhead the data admit. After the placement, the closed-form screen applied to every observed format at every treated tract finds 146 format-by-tract pairs above their thresholds, against the 36 candidates planted, so the second experiment can understate the entry a richer income mix attracts. Prices move only through markups, with each barcode-store pricing residual held at its observed value, so the second experiment understates the price increase a richer income mix

brings (Appendix D.7). Each experiment is one shock at one site, solved by continuation from the observed allocation, and Appendix E treats the shelf approximation, the barcode pools and the selection of equilibria.

8.2 Opening a high-end supermarket

Table 10: The neighborhood's response to a high-end supermarket

	All	Q1	Q2	Q3	Q4
<i>The retail environment</i>					
Location price index, $\Delta \log \mathcal{P}$ (log points)	−1.40				
Incumbent barcodes added / dropped	2636 / 230				
Incumbent shelf quality, mean $\log \xi$	−0.273 to −0.273				
<i>The neighborhood (%)</i>					
Population	0.08	−0.57	0.08	0.38	0.68
Mean income	0.38				
Rent	0.25				
<i>Welfare, consumption equivalent (%)</i>					
Treated-area incumbents		0.06	0.07	0.08	0.11
from the grocery index		0.10	0.11	0.14	0.18
lost to rent		0.04	0.04	0.05	0.06
City		0.15	0.05	0.02	0.00
$d \log \bar{y} / d \log \mathcal{P}$	−0.270				

Notes: Changes between the observed equilibrium and the equilibrium with the store open, at shelf-cost draw 1. Q1 to Q4 are household income quartiles. No other store opens (0) and none closes (0). Incumbent barcodes are counted over every incumbent store in the city against its observed shelf; shelf quality is the mean $\log \xi$ of the incumbents the treated area reaches. Quality is centred on each category's barcode mean, and this mean counts a barcode once for each store that carries it; it lies below zero because the barcodes carried by more stores are of lower quality. Rent is the population-weighted change in log rent. Welfare is defined in Section 8.1, split into its two terms; the City row averages over every tract. $\bar{y}$ is the mean income of the treated area's residents.

The store lowers the treated area's location price index by −1.40 log points and raises its mean income by 0.38 percent (Table 10). The population of the poorest quartile falls and that of the other three rises, most in the top one. No candidate entrant clears its threshold and no incumbent closes, so the forced store is the only change to the network. An incumbent closes only once its operating profit falls below the most permissive floor the observed city admits (Appendix E), and the absence of closures holds under that rule.

The incumbents lengthen their shelves without changing their quality. Across the city they add 2636 barcodes and drop 230, and the mean quality of the incumbent shelves the treated area reaches is unchanged.

Every income group already living in the treated area gains, and the gain rises with income. The gain from the grocery index rises from the poorest quartile to the richest, as expected of a high-end store, and the rent the store capitalizes takes back a smaller part of it that also rises

with income. The store also occupies floorspace, so part of that rise in rent comes from its own claim on the stock and not from the index. Across the city the gain falls with income, from 0.15 percent for the poorest quartile to 0.00 percent for the richest.

The city-wide gain shapes who moves in. A household weighs the treated area against every other tract, so what draws it there is its gain in the treated area relative to its gain elsewhere. The poorest quartile gains more across the city than in the treated area, and its population there falls. The richest quartile gains only in the treated area, and it moves in most. The sorting response, and the elasticity built from it, therefore reflect the incumbents' response across the city as well as the store itself.

8.3 Moving higher-income households in

Table 11: The response to a richer income mix

	All	Q1	Q2	Q3	Q4
<i>The treated area (log points)</i>					
Mean income, equilibrium	22.2				
Grocery expenditure	12.9				
<i>The retail environment</i>					
Stores opened / closed	1 / 0				
Barcodes carried, stores the area reaches	164130 to 169944				
Incumbent shelf quality, mean $\log \xi$	-0.273 to -0.256				
Incumbent barcodes added / dropped	776 / 5905				
Location price index, $\Delta \log \mathcal{P}$ (log points)	-3.63				
$d \log \mathcal{P} / d \log \bar{y}$	-0.164				
<i>The neighborhood (%)</i>					
Rent	14.11				
Population, net of the placement		-19.71	-16.23	-12.88	-7.83
<i>Welfare, consumption equivalent (%)</i>					
Treated-area incumbents		-3.57	-2.87	-2.23	-1.26
from the grocery index		0.32	0.34	0.37	0.33
lost to rent		3.89	3.21	2.60	1.59
Origin tracts		0.03	0.03	0.03	0.03
Rest of the city		0.04	0.04	0.04	0.03
Placed households					7.23

Notes: Changes between the observed equilibrium and the equilibrium after the placement, at shelf-cost draw 1. Q1 to Q4 are household income quartiles. Mean income is its change in the new equilibrium, after re-sorting, of which the placement itself accounts for 17.8 log points (Table 13). Barcodes carried cover every store the treated area reaches, entrants included, and shelf quality covers the incumbents among them; incumbent barcodes are counted over every incumbent store in the city against its observed shelf. $\bar{y}$ is the treated area's mean income. Rent is the population-weighted change in log rent. Population net of the placement is the change in each quartile's population in the treated area less the households placed, relative to its baseline, and it includes the share 0.10 of every quartile's residents that the placement itself removes. Welfare is defined in Section 8.1; for the placed households it compares the treated area after the shock with the origin tracts before it.

A richer income mix makes the neighborhood worth serving. A candidate store opens in the

treated area, and the stores its residents reach carry more barcodes of higher average quality (Table 11). No incumbent closes under the exit floor of Section 8.2. Grocery expenditure rises by less than income, since the budget share of groceries falls with income, and the location price index falls.

The entrant crowds out the incumbents' variety. Incumbent stores across the city drop 5905 barcodes and add 776, since the entrant takes part of their customers and the spending that supports a long shelf goes with them. The treated area's residents still face more barcodes than before, because the entrant's shelf more than replaces what the incumbents drop.

Table 12: The change in the location price index, by margin

Margin of (12)	$\Delta \log \mathcal{P}$ (log points)
Proximity	−0.58
Store variety	−0.70
Category variety	−2.64
Barcode variety	−2.50
Quality and price	2.79
Total	−3.63

Notes: The margins of (12), each averaged over the treated area at baseline population weights and differenced between the observed and the counterfactual equilibrium. The quality and price terms are reported net of each other; Appendix E gives them separately. The identity holds to machine precision at both ends, so the margins sum to the total.

Category and barcode variety, in similar parts, lower the index by more than its total fall, and store variety and proximity lower it further (Table 12). Households shift their purchases toward the entrant's barcodes, which are of higher quality and, with marginal cost rising in quality, of higher price, and the quality and price terms together raise the index. The gain from the entrant therefore comes through the variety households reach and not through the quality of what they switch to. The entrant opens inside the treated area and takes 13 percent of its grocery spending, so the effective number of stores within reach rises and the trip to them shortens.

Rents rise by 14.11 percent in the treated area and change by −0.02 percent in the rest of the city, and an incumbent of every income group who stays in the neighborhood loses. The gain from the grocery index is close to equal across quartiles, while the loss to rent falls with income, since poorer households spend a larger share of their income on housing. The losses are therefore regressive. The model treats every household as a renter, so these are renters' losses; an owner who stays pays the higher rent to themselves and is insulated from that term. The households placed in the treated area gain relative to the tracts they came from, and welfare changes by 0.03 percent in the origin tracts and by 0.03 percent in the rest of the city.

Displacement follows the same order. The placement itself removes a share 0.10 of every quartile's residents. Relative to the allocation it leaves, the three lower quartiles lose further residents, the poorest the most, while the top quartile gains (Table 11). The neighborhood the shock produces is therefore richer than the placement alone made it, because the rent it capitalizes

prices the poorer incumbents out.

8.4 What the retail response adds to sorting

Table 13: The placement with and without the retail response

	Retail frozen		Retail responds	
Mean income of the treated area (log points)	22.4		22.2	
of which the placement itself	17.8		17.8	
Rent (%)	13.61		14.11	
	Population	Welfare	Population	Welfare
Q1	−20.67	−3.80	−19.71	−3.57
Q2	−17.34	−3.12	−16.23	−2.87
Q3	−14.16	−2.50	−12.88	−2.23
Q4	−8.97	−1.47	−7.83	−1.26

Notes: The placement of Section 8.3 solved twice. Under “retail frozen” every store keeps its observed prices, shelf and existence, so households respond only to the placement and to rents; under “retail responds” every margin of Table 11 is free. Population is each quartile’s change in the treated area net of the placement and welfare is that of the treated area’s incumbents, both in percent and defined as in Table 11.

The two experiments measure the two directions on the same treated area. A one-percent fall in the index raises mean income by 0.270 percent, while the placement moves the index by −0.164 percent per percent of mean income. They are finite responses to different shocks and are compared, not multiplied (Appendix E).

The retail response does not drive the change in who lives in the neighborhood. Solving the same placement with every store’s prices, shelf and existence held at the observed ones separates three parts of the rise in the treated area’s mean income (Table 13). The placement itself accounts for 80 percent of it, the sorting that rents produce for 20 percent, and the retail response for −1 percent. With the shops free to respond, mean income rises by 22.2 log points, against 22.4 with them frozen.

The retail response keeps incumbents of every income group, and so slightly dilutes the placement. Against the frozen solution, it raises the welfare of the treated area’s incumbents by a similar amount in every quartile, and each quartile keeps a similar additional share of its residents. Those incumbents are poorer on average than the households the placement brings, so keeping them lowers the rise in mean income. The first experiment differs in whom the retail change favors. The high-end store raises the grocery gain of the richest quartile most, and the poorest gains more elsewhere in the city, so the store draws richer households in.

9 Conclusion

Local grocery markets make cities more unequal. Richer neighborhoods pay more for the same barcode, yet face a lower grocery price index, and the difference comes from the stores they can reach. A neighborhood one percent richer faces a price index 0.128 percent lower, and the effective number of stores within reach lowers it by 0.168 percent, more than the total difference. Inside the store, higher prices outweigh the higher product quality of what the stores of richer neighborhoods stock.

Unequal access to stores widens the gap in real consumption between rich and poor households. If every household could shop at every store at no trip cost, consumption would rise by 10.5 percent on average and by 12.3 percent for the poorest quartile, and the ratio of the richest quartile's real consumption to the poorest's would fall by 5.1 percent. The poorest gain the most because groceries are a larger share of their budget and because they live where access is worse. Most of the gain comes from reaching more stores and not from shorter trips to the stores already within reach.

The results concern how differences in the cost of living within a city are measured. Comparing the price of the same barcode across neighborhoods finds richer neighborhoods more expensive, while the index that counts the stores within reach finds them cheaper. A price index built from the goods households buy measures the first and misses the second.

Grocery retail and a neighborhood's income mix also respond to each other. A high-end supermarket opened in a middle-income neighborhood raises its mean income by 0.270 percent for each percent fall in the location price index, as richer households move in and the poorest quartile's population falls. Moving richer households into the same neighborhood brings new store entry and higher-quality products to its shelves. The rent increases that accompany the change in residents' incomes leave incumbents of every income group worse off, while the better stores raise the welfare of incumbents in every group.

References

- Abowd, John M., Francis Kramarz, and David N. Margolis**, “High Wage Workers and High Wage Firms,” *Econometrica*, 1999, 67 (2), 251–333.
- Ackerberg, Daniel A. and Marc Rysman**, “Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects,” *RAND Journal of Economics*, 2005, 36 (4), 771–788.
- Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf**, “The Economics of Density: Evidence from the Berlin Wall,” *Econometrica*, 2015, 83 (6), 2127–2189.
- Allcott, Hunt, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell**, “Food Deserts and the Causes of Nutritional Inequality,” *The Quarterly Journal of Economics*, 2019, 134 (4), 1793–1844.
- Almagro, Milena and Tomás Domínguez-Iino**, “Location Sorting and Endogenous Amenities: Evidence from Amsterdam,” *Econometrica*, 2025, 93 (3), 1031–1071.
- Ashenfelter, Orley and Alan Krueger**, “Estimates of the Economic Return to Schooling from a New Sample of Twins,” *American Economic Review*, 1994, 84 (5), 1157–1173.
- Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro**, “Retail Globalization and Household Welfare: Evidence from Mexico,” *Journal of Political Economy*, 2018, 126 (1), 1–73.
- Bajari, Patrick, Han Hong, and Stephen P. Ryan**, “Identification and Estimation of a Discrete Game of Complete Information,” *Econometrica*, 2010, 78 (5), 1529–1568.
- Baum-Snow, Nathaniel and Lu Han**, “The Microgeography of Housing Supply,” *Journal of Political Economy*, 2024, 132 (6), 1897–1946.
- Behrens, Kristian, Brahim Boualam, Julien Martin, and Florian Mayneris**, “Gentrification and Pioneer Businesses,” *Review of Economics and Statistics*, 2024, 106 (1), 119–132.
- Berry, Steven and Joel Waldfogel**, “Product Quality and Market Size,” *Journal of Industrial Economics*, 2010, 58 (1), 1–31.
- Berry, Steven T.**, “Estimation of a Model of Entry in the Airline Industry,” *Econometrica*, 1992, 60 (4), 889–917.
- Borraz, Fernando, Felipe Carozzi, Nicolas Gonzalez-Pampillon, and Leandro Zipitria**, “Local Retail Prices, Product Variety, and Neighborhood Change,” *American Economic Journal: Economic Policy*, 2024, 16 (1), 1–33.

- Bresnahan, Timothy F. and Peter C. Reiss**, “Entry and Competition in Concentrated Markets,” *Journal of Political Economy*, 1991, 99 (5), 977–1009.
- Broda, Christian and David E. Weinstein**, “Product Creation and Destruction: Evidence and Price Implications,” *American Economic Review*, 2010, 100 (3), 691–723.
- Bronnenberg, Bart J., Jean-Pierre Dubé, Matthew Gentzkow, and Jesse M. Shapiro**, “Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium,” *The Quarterly Journal of Economics*, 2015, 130 (4), 1669–1726.
- Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline**, “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labor Economics*, 2018, 36 (S1), S13–S70.
- Chen, Zhijun and Patrick Rey**, “Loss Leading as an Exploitative Practice,” *American Economic Review*, 2012, 102 (7), 3462–3482.
- Cho, Clare and Richard Volpe**, “Independent Grocery Stores in the Changing Landscape of the U.S. Food Retail Industry,” Economic Research Report ERR-240, U.S. Department of Agriculture, Economic Research Service 2017. Store counts and sales from Nielsen TDLinx; an independent is a grocery store whose owner operates fewer than four stores. Accessed 2026-09.
- Ciliberto, Federico and Elie Tamer**, “Market Structure and Multiple Equilibria in Airline Markets,” *Econometrica*, 2009, 77 (6), 1791–1828.
- Conselho Nacional de Política Fazendária**, “Convênio ICMS 134/2016,” 2016. Requires financial and payment institutions to report to the state tax administrations the card, Pix and other electronic payment transactions of sellers identified by CNPJ or CPF, including sellers not enrolled in the state ICMS taxpayer register. In force since 1 February 2017; upheld by the Supreme Federal Court in ADI 7276 (2024). Accessed 2026-09-11.
- Couture, Victor and Jessie Handbury**, “Urban Revival in America,” *Journal of Urban Economics*, 2020, 119, 103267.
- , **Cécile Gaubert, Jessie Handbury, and Erik Hurst**, “Income Growth and the Distributional Effects of Urban Spatial Sorting,” *Review of Economic Studies*, 2024, 91 (2), 858–898.
- Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales**, “How Segregated Is Urban Consumption?,” *Journal of Political Economy*, 2019, 127 (4), 1684–1738.
- Davis, Peter**, “Spatial Competition in Retail Markets: Movie Theaters,” *RAND Journal of Economics*, 2006, 37 (4), 964–982.

- DellaVigna, Stefano and Matthew Gentzkow**, “Uniform Pricing in U.S. Retail Chains,” *The Quarterly Journal of Economics*, 2019, 134 (4), 2011–2084.
- Diamond, Rebecca**, “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000,” *American Economic Review*, 2016, 106 (3), 479–524.
- **and Enrico Moretti**, “Where Is Standard of Living the Highest? Local Prices and the Geography of Consumption,” Working Paper 29533, National Bureau of Economic Research 2021.
- Draganska, Michaela, Michael Mazzeo, and Katja Seim**, “Beyond Plain Vanilla: Modeling Joint Product Assortment and Pricing Decisions,” *Quantitative Marketing and Economics*, 2009, 7 (2), 105–146.
- Eizenberg, Alon**, “Upstream Innovation and Product Variety in the U.S. Home PC Market,” *Review of Economic Studies*, 2014, 81 (3), 1003–1045.
- **, Saul Lach, and Merav Oren-Yiftach**, “Retail Prices in a City,” *American Economic Journal: Economic Policy*, 2021, 13 (2), 175–206.
- Ellickson, Paul B.**, “Does Sutton Apply to Supermarkets?,” *RAND Journal of Economics*, 2007, 38 (1), 43–59.
- **, Paul L. E. Grieco, and Oleksii Khvastunov**, “Measuring Competition in Spatial Retail,” *RAND Journal of Economics*, 2020, 51 (1), 189–232.
- Estado de Alagoas**, “Decreto n. 43.606, de 30 de julho de 2015: institui a Nota Fiscal de Consumidor Eletronica (NFC-e), modelo 65,” *Diario Oficial do Estado de Alagoas* 2015. Accessed 2026-09-07.
- European Space Agency and Airbus**, “Copernicus DEM: Global and European Digital Elevation Model (GLO-90),” 2022. 90 m global raster. Accessed 2026-09.
- Faber, Benjamin and Thibault Fally**, “Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data,” *Review of Economic Studies*, 2022, 89 (3), 1420–1459.
- Fan, Ying and Chenyu Yang**, “Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market,” *American Economic Journal: Microeconomics*, 2020, 12 (2), 99–134.
- Feenstra, Robert C.**, “New Product Varieties and the Measurement of International Prices,” *American Economic Review*, 1994, 84 (1), 157–177.
- George, Lisa and Joel Waldfogel**, “Who Affects Whom in Daily Newspaper Markets?,” *Journal of Political Economy*, 2003, 111 (4), 765–784.

- Glaeser, Edward L., Hyunjin Kim, and Michael Luca**, “Nowcasting Gentrification: Using Yelp Data to Quantify Neighborhood Change,” *AEA Papers and Proceedings*, 2018, 108, 77–82.
- , **Michael Luca, and Erica Moszkowski**, “Gentrification and Retail Churn: Theory and Evidence,” *Regional Science and Urban Economics*, 2023, 100.
- Google**, “Google Places API,” 2025. Accessed 2025.
- , “Places Aggregate API, Google Maps Platform,” 2026. Count of operating establishments by place type within the municipality of Maceio, retrieved 2026-09-10. Query and response cached in the replication package.
- Guedes, Ricardo, Felipe S. Iachan, and Marcelo Sant’Anna**, “Housing Supply in the Presence of Informality,” *Regional Science and Urban Economics*, 2023, 99.
- Handbury, Jessie**, “Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities,” *Econometrica*, 2021, 89 (6), 2679–2715.
- **and David E. Weinstein**, “Goods Prices and Availability in Cities,” *Review of Economic Studies*, 2015, 82 (1), 258–296.
- , **Ilya Rahkovsky, and Molly Schnell**, “Is the Focus on Food Deserts Fruitless? Retail Access and Food Purchases Across the Socioeconomic Spectrum,” NBER Working Paper 21126, National Bureau of Economic Research 2015.
- Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm**, “The Making of the Modern Metropolis: Evidence from London,” *The Quarterly Journal of Economics*, 2020, 135 (4), 2059–2133.
- Hoelzlein, Matthias**, “Two-Sided Sorting and Spatial Inequality in Cities,” Working Paper, University of Notre Dame 2023. Second revision at the Journal of Political Economy.
- **and Sam Miller**, “JUE Insight: The Gentrification Effect of Retail Amenities: Evidence from Whole Foods Market,” SSRN Working Paper 5009323 2024.
- Hottman, Colin J., Stephen J. Redding, and David E. Weinstein**, “Quantifying the Sources of Firm Heterogeneity,” *The Quarterly Journal of Economics*, 2016, 131 (3), 1291–1364.
- Houde, Jean-François**, “Spatial Differentiation and Vertical Mergers in Retail Gasoline Markets,” *American Economic Review*, 2012, 102 (5), 2147–2182.
- Instituto Brasileiro de Geografia e Estatística**, “Censo Demografico 2010: Microdados da Amostra,” 2012.
- , “Pesquisa de Orcamentos Familiares (POF),” 2018. Accessed 2026-02-23.

- , “PNAD Continua,” 2019. Accessed 2026-02-23.
- , “Cadastro Nacional de Enderecos para Fins Estatisticos (CNEFE),” 2022. Accessed 2026-09-07.
- , “Pesquisa Anual de Comercio (PAC),” 2024. SIDRA table 10645. Accessed 2026-09.
- , “Censo Demografico 2022: Microdados da Amostra - acesso controlado,” 2026. Released 2026 under a controlled-access model requiring institutional accreditation and a signed confidentiality agreement. Accessed 2026-09-12.
- Jaravel, Xavier**, “The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector,” *Quarterly Journal of Economics*, 2019, 134 (2), 715–783.
- Kilts Center for Marketing**, “NielsenIQ Retail Scanner and Consumer Panel Datasets,” 2025. Dataset documentation. Accessed 2026-09.
- Kline, Patrick**, “Firm Wage Effects,” in “Handbook of Labor Economics,” Vol. 5, Elsevier, 2024.
- , **Raffaele Saggio**, and **Mikkel Sølvsten**, “Leave-Out Estimation of Variance Components,” *Econometrica*, 2020, 88 (5), 1859–1898.
- Lal, Rajiv and Carmen Matutes**, “Retail Pricing and Advertising Strategies,” *Journal of Business*, 1994, 67 (3), 345–370.
- Mazzeo, Michael J.**, “Product Choice and Oligopoly Market Structure,” *RAND Journal of Economics*, 2002, 33 (2), 221–242.
- Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J. Redding**, “The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data,” *The Quarterly Journal of Economics*, 2025, 140 (4), 2507–2570.
- Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg**, “Commuting, Migration, and Local Employment Elasticities,” *American Economic Review*, 2018, 108 (12), 3855–3890.
- Moretti, Enrico**, “Real Wage Inequality,” *American Economic Journal: Applied Economics*, 2013, 5 (1), 65–103.
- Nash, John F.**, “Equilibrium Points in n -Person Games,” *Proceedings of the National Academy of Sciences*, 1950, 36 (1), 48–49.
- NielsenIQ**, “Estudo do pequeno varejo brasileiro,” 2025. Survey of 400 establishments. Traditional retail and independent self-service together are 33% of Brazilian points of sale and 30% of billing; traditional alone is 11% and independent self-service 19%. Reported June 2025.

- Nocke, Volker and Nicolas Schutz**, “Multiproduct-Firm Oligopoly: An Aggregative Games Approach,” *Econometrica*, 2018, 86 (2), 523–557.
- Open Food Facts**, “Open Food Facts Product Database,” 2026. Open Database License. Accessed at a pinned revision; see Appendix A.
- OpenStreetMap contributors**, “OpenStreetMap,” 2025. Open Database License.
- Quan, Thomas W. and Kevin R. Williams**, “Product Variety, Across-Market Demand Heterogeneity, and the Value of Online Retail,” *RAND Journal of Economics*, 2018, 49 (4), 877–913.
- Receita Federal do Brasil**, “Cadastro Nacional da Pessoa Juridica (CNPJ): dados publicos,” 2025. Accessed 2025.
- , “Microempreendedor Individual (MEI): limite de receita bruta anual,” 2025. Ceiling set at R\$81,000 by Lei Complementar 155/2016, in force since 2018. Accessed 2026-09-07.
- Redding, Stephen J. and David E. Weinstein**, “Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences,” *Quarterly Journal of Economics*, 2020, 135 (1), 503–560.
- and **Esteban Rossi-Hansberg**, “Quantitative Spatial Economics,” *Annual Review of Economics*, 2017, 9, 21–58.
- Rhodes, Andrew**, “Multiproduct Retailing,” *Review of Economic Studies*, 2015, 82 (1), 360–390.
- Rusmevichientong, Paat, David Shmoys, Chaoxu Tong, and Huseyin Topaloglu**, “Assortment Optimization under the Multinomial Logit Model with Random Choice Parameters,” *Production and Operations Management*, 2014, 23 (11), 2023–2039.
- Saiz, Albert**, “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*, 2010, 125 (3), 1253–1296.
- Secretaria de Estado da Fazenda de Alagoas**, “Instrucao Normativa SEF n. 23, de 3 de maio de 2017: cronograma de obrigatoriedade da NFC-e,” 2017. Replaces Instrucao Normativa SEF n. 46/2015. Accessed 2026-09-07.
- SEFAZ-AL**, “Economiza Alagoas - Desenvolvedor,” Economiza Alagoas 2026. Accessed 2026-02-23.
- , “Manual de Orientacao do Desenvolvedor,” Economiza Alagoas 2026. Accessed 2026-02-23.
- , “Sobre o Economiza Alagoas,” Economiza Alagoas 2026. Accessed 2026-02-23.

- Seim, Katja**, “An Empirical Model of Firm Entry with Endogenous Product-Type Choices,” *RAND Journal of Economics*, 2006, 37 (3), 619–640.
- Smith, Howard**, “Supermarket Choice and Supermarket Competition in Market Equilibrium,” *Review of Economic Studies*, 2004, 71 (1), 235–263.
- Soderbery, Anson**, “Estimating Import Supply and Demand Elasticities: Analysis and Implications,” *Journal of International Economics*, 2015, 96 (1), 1–17.
- Sorkin, Isaac**, “Ranking Firms Using Revealed Preference,” *The Quarterly Journal of Economics*, 2018, 133 (3), 1331–1393.
- Stroebe, Johannes and Joseph Vavra**, “House Prices, Local Demand, and Retail Prices,” *Journal of Political Economy*, 2019, 127 (3), 1391–1436.
- Tamer, Elie**, “Incomplete Simultaneous Discrete Response Model with Multiple Equilibria,” *Review of Economic Studies*, 2003, 70 (1), 147–165.
- Thomassen, Øyvind, Howard Smith, Stephan Seiler, and Pasquale Schiraldi**, “Multi-Category Competition and Market Power: A Model of Supermarket Pricing,” *American Economic Review*, 2017, 107 (8), 2308–2351.
- Tsivanidis, Nick**, “Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá’s TransMilenio,” *American Economic Review*, 2026, 116 (2), 418–463.

Appendix

A Data Appendix

Table 14 inventories every source the paper uses, its reference period, its unit of observation, and where it enters. The sections that follow document acquisition and cleaning, then the three measures the analysis constructs from the price data: what counts as available, how tract income composition is estimated, and how barcode quality is recovered. Depending on the object being constructed, the unit of analysis is the establishment, the census tract, or the census weighting area (área de ponderação).

A.1 Price data: collection

Economiza Alagoas is a state-run price-transparency program operated by SEFAZ-AL (SEFAZ-AL, 2026c). Its public API (SEFAZ-AL, 2026b) accepts product queries (description or barcode), municipality codes, and a look-back window. Each call uses a 10-day look-back and requests up to 5,000 records per page. Timestamps are reported in UTC.

The API is designed for a shopper looking up one product, and returns only records whose description matches a query string. Recovering the full set of products on sale therefore requires querying the description space exhaustively instead of requesting a catalogue. I do this by enumerating all $26^3 = 17,576$ three-letter strings and issuing each as a separate query, for every municipality. Any description containing three consecutive unaccented letters matches at least one such string, so the enumeration returns every product whose description has such a run, which a description lacks only if it is very short or broken up by accents, digits or punctuation. A single sweep of the study area is therefore roughly $14 \times 17,576 \approx 246,000$ calls, run concurrently with rate limiting. Any single call looks back only 10 days, so coverage is a property of the sweep calendar and not of the nominal window. The analysis data is assembled from 21 sweeps run between 27 May 2025 and 21 November 2025, with a median of 8 days and a maximum of 20 days between consecutive sweeps. No call reaches the days between two sweeps more than 10 days apart, nor the days after the last sweep. Between 23 May 2025 and 23 November 2025 these add up to 22 days, or 12% of the window, in 4 runs of at most 9 days, the last of which follows the final sweep.

A store's assortment is the set of barcodes that sold there during the window, since a barcode-establishment pair enters the data only when a sweep's look-back catches a sale of it. A barcode on the shelf that never sold in six months is not observed, and it is not part of what the store's customers bought. The definition has two consequences. The union over the window counts as one shelf barcodes that were not all on sale on the same day. And a slow seller can go undetected if its few sales fall in the uncovered days, which would bias the income elasticities of Section 3 if slow sales were more common in poorer neighborhoods, for instance for a premium barcode

Table 14: Inventory of data sources

Source	Period	Unit	Role in the paper
<i>Core</i>			
Economiza Alagoas (NFC-e)	2025	Barcode $\times$ store	Transaction prices and assortments; the paper’s core data
<i>Household demand</i>			
Consumption survey	2017–18	Household	Expenditure shares β_g, ρ_g by income quartile; first stages for σ_b and σ_c from diary unit values
IBGE Census	2022	Census tract	Population, households, mean and variance of household-head income; tract centroids
IBGE Census, sample	2022	Weighting area	Observed income composition nationally; estimates the tract income-composition map. Controlled access
IBGE Census	2010	Weighting area	Mean rents in Maceió, the most recent census to measure rent; income and rent for weighting areas in 125 cities nationally
PNAD Contínua	2019	Person	CNPJ registration of food-retail businesses, the lower boundary of invoice coverage
<i>Establishments</i>			
Receita Federal	2025	Establishment	Legal names, opening dates and firm size for identifier standardization in Maceió; nationally, the register’s opening and closing dates reconstruct historical grocery stocks
Google Places	2025	Establishment	Geocoding of establishments to coordinates
Google Places Aggregate	2026	Municipality	Independent count of operating grocery retailers, for the coverage check
IBGE CNEFE	2022	Address point	Geocoding fallback where no confident Google match exists
OpenStreetMap	2025	Building / zone	Building footprints; store floorspace h_f
Copernicus GLO-90	2021	90 m raster	Elevation and terrain slope, controls for the sorting-elasticity instrument
IBGE Annual Survey of Commerce	2024	Retail activity	Margin over cost of goods sold, used to set the level of marginal costs
<i>Products</i>			
Open Food Facts	2026	Barcode	External validation of the quality index

Notes: The three census rows are separate entries because each supplies a different object: the 2022 tract aggregates measure each neighborhood, the 2022 sample estimates how income composition relates to those measurements, and the 2010 sample is used for rent, which the 2022 questionnaire did not collect, and nationally for the sorting elasticity. Section 5 states which supplies each model input.

stocked in a low-traffic store. The number of sweeps that record a pair measures how close it came to being missed. The median pair among the 1,565,228 in the descriptive sample is recorded by 7 of the 21 sweeps; 11.8% are recorded by one sweep only, and 18.1% of those at or above the 80th percentile of their category’s quality distribution. The share recorded by a single sweep does not move with the income of the establishment’s tract. A linear probability regression on log tract income gives 0.11 percentage points (s.e. 0.73, clustered by tract) over all pairs and -0.79 (0.88) over premium ones, and 0.35 (0.67) and -0.60 (0.70) with the store’s barcode count held fixed. Differential detection of this size is small next to the premium-end variety elasticity of 1.029.

I handle pagination and retries. Within a sweep the same barcode–establishment record is returned by many different query strings, since a product description contains several three-letter substrings, so I keep one copy of each, keyed on barcode and Firm ID. Records from different sweeps are retained separately: each reports the last transaction within that sweep’s look-back window, and it is these repeated observations that the cell-level collapse described below reduces to one price per pair. The resulting raw file contains barcode, HS-8, the Global Product Classification (GPC) code, description, unit of measure, last sale date, listed and sale prices, and establishment identifiers and addresses.

Table 15 shows what a record looks like.

Table 15: Illustrative records from the price data

Barcode	HS-8	Product description	Price (US\$)	Establishment
7896259410164	0405.10.00	Manteiga Camponesa Com Sa...	1.79	Uni Compra Supermercado...
7896481130137	1104.19.00	Flocao De Milho Coringa 500g	0.32	Super Varejo Marechal Ltda
7897261800011	1701.99.00	Acucar Caete 1kg	0.59	Mercadinho O Precinho Ltda
7892840822255	1904.10.00	Snack Cebolitos 36gr	0.71	Conveniencia Via Norte ...
7896945400813	2103.20.10	Catchup Tambau	1.06	Atacadao S.a.
7894900593709	2202.10.00	Kapo Uva 200ml	0.43	Hipermercado Marechal

Notes: Six records, one from each of six HS-2 chapters, chosen to show the range of the product space, not as a random draw; within a chapter the record shown is the most recent. Descriptions are establishment-supplied, and are reproduced here with capitalisation normalised and long strings truncated to fit the column. Prices are converted to 2025 US dollars at 5.59 reais per dollar. Establishment names are published by the program alongside prices. Street addresses, also published, are omitted here for space.

A.2 Price data: cleaning and product harmonization

I retain a consistent set of fields (barcode, HS-8, GPC, unit, sale date, declared and sale values, description, and Firm ID), pad Firm ID to 14 digits, parse dates, and treat zero-valued numeric fields as missing. Missing or inconsistent identifiers are corrected in two steps: (i) each barcode is assigned its modal HS-8, and a record without a barcode keeps the HS-8 on its invoice; (ii) missing barcodes are filled using the most common barcode for the same HS-8 and normalized

description. The second step recovers 15,373 of the 40,401,283 records collected. Records whose barcode survives unfilled are dropped. Missing HS-8 and GPC codes are then filled with the modal values for each barcode.

A.3 Establishments: identification, geocoding, and linkage

Establishment addresses from the price data are normalized using a common set of street replacements (e.g., "AVENIDA" → "AV") and de-duplicated by Firm ID. The Receita Federal Firm ID files are merged to attach legal names, opening dates, capital, and firm size indicators (Receita Federal do Brasil, 2025a).

Geocoding proceeds in stages. Establishments are matched by name and address to Google Places (Google, 2025) with a confidence threshold ≥ 0.25 ; 79% of the 8,103 establishments are geocoded this way, at a median match confidence of 0.78. The threshold is permissive by design, but few accepted matches sit near it: the tenth percentile of accepted confidence is 0.35 and the twenty-fifth is 0.54. A match that is nonetheless wrong is caught downstream, not at the threshold, since a store placed in the wrong city fails the census-tract requirement below. Establishments without a confident Google match are geocoded against official IBGE address points (CNEFE) (Instituto Brasileiro de Geografia e Estatística, 2022) in three further stages. The first matches the street address exactly within the establishment's CEP. The second, for addresses that fail it, relaxes the CEP restriction and takes the closest address anywhere in the metropolitan file by Jaro–Winkler string distance, subject to a maximum distance of 0.05. The third assigns whatever remains to its CEP centroid. The second stage is the permissive one, since dropping the CEP restriction allows a match to a similarly named street in another part of the city.

After geocoding, 3.2% of establishments still cannot be placed in a census tract and drop out of any tract-level exhibit. These are failed matches, not borderline locations. Their coordinates fall in other Brazilian cities entirely, which happens when the name-and-address search finds a same-named establishment elsewhere, and requiring a census tract inside the study area screens them out. The CNEFE stages search only the metropolitan address file, so a wrong relaxed match places a store on a similarly named street inside the study area, where the tract requirement does not catch it.

Geocoding is pinned to a snapshot so that repeated runs return identical coordinates. 33 establishments entered the price data after that snapshot and were therefore never sent to the geocoder. They carry no coordinates, cannot be assigned to a census tract, and so appear in Table 1 but in no tract-level exhibit and not in the estimation sample.

Every establishment attribute beyond price comes from a match to an outside source, and each match has its own coverage, reported in Table 16. The binding one is OpenStreetMap, whose building footprints cover only 15% of establishments, which is why commercial floorspace is built at the weighting-area level, not per store.

Table 16: Establishment linkage coverage

Linked source	Attribute used	Matched (%)
Google Places	Coordinates, name	79
IBGE census tract	Neighborhood assignment	97
OpenStreetMap	Building footprint area	15

Notes: Share of the 8,103 establishments sent to the geocoder, which is the roster the matching was performed on. It differs from the establishment count in Table 1 in both directions, and the two reconcile exactly: 33 establishments appear in the price window but entered after the geocoding snapshot and were never sent to the geocoder, while 11 were geocoded but record no transaction inside the window. Google Places coordinates require a name-and-address match confidence of at least 0.25.

A.4 Sample construction

Cleaned prices are merged with geocoded establishments, census tracts, and weighting areas, and restricted to food categories (HS-2 chapters 01–24). When several records map to the same barcode–establishment cell, I collapse them to the median price, so that each cell carries one price. The collapse rule is set once, in the cleaning step, and the estimation reads the collapsed cell instead of re-aggregating the transactions for itself. This collapse is innocuous because prices are stable within a pair over the window. Table 17 reports the panel behind it. 88% of pairs are observed more than once, and among those, 52% never change price and the ninetieth-percentile within-pair coefficient of variation is 0.118. Both dispersion statistics are computed on pairs observed more than once: a pair seen a single time carries no information about whether its price moved, and scoring it as unchanged would let the sweep calendar’s coverage appear as price stickiness. The cross-sectional treatment of prices therefore discards little within-window time variation.

Table 17: Panel structure of the price data

Per barcode–establishment pair	Median	90th pct.
Transactions observed	7	20
Distinct transaction dates	6	18
Days since last transaction	15	128
Share of pairs observed more than once	88%	
Share of pairs with an unchanged price	52%	
Within-pair price coeff. of variation (p90)	0.118	

Notes: Pairs are distinct barcode–establishment combinations with a non-missing barcode and a positive price, over the six-month sample window, in establishments of the grocery sample. The last two rows are computed on the 88% of pairs observed more than once, which are the pairs whose price can be seen to move or not; a pair observed once is excluded from them rather than counted as unchanged. Days since last transaction is measured from the end of the window and is a staleness diagnostic, not a price statistic.

Table 18 traces the sample. Panel A is the descriptive grocery sample whose last row reproduces Table 1. Panel B is the sample the structural model is estimated on. Both panels cover

the same window, 23 May 2025 to 23 November 2025. Panel B starts from the cleaned cell file for every product and keeps the grocery barcodes with a quality estimate. It comes before the two-barcode rule of Panel A, so it holds one more cell for each single-item seller, and it loses the establishments none of whose barcodes has a quality estimate. Panel B applies five further restrictions. It uses the cleaned barcode–establishment cell price and not the underlying transactions; it keeps only barcodes for which a quality shifter and an HS-8 cost effect are both available, since the model cannot price a barcode without them; it drops pairs whose implied cost exceeds twice the recorded price; it drops establishments the model cannot place, those without a census tract and those farther than 3 km from every residential tract; and it retains only establishments carrying at least 20 grocery barcodes. The cost screen compares each recorded price with the marginal cost the hedonic decomposition implies at its reference level, before the cost level is set by the retail margin, and it is fixed there so that the sample does not move with that calibration. A cell whose implied cost is more than twice its price is most likely a price recorded in the wrong unit, a pack price entered for a single item or the reverse, though the rule drops any cell that far below cost. This leaves 1,958 establishments and 1,561,793 barcode–establishment pairs. The descriptive sample is therefore deliberately broader than the estimation sample. Panel B shows what the store restriction costs: it removes establishments in proportion far exceeding the barcode–establishment pairs that leave with them, because the establishments it removes carry few barcodes each.

The residential side of the model is a separate universe from the establishments in Panel B. It comprises the 1,725 census tracts that meet the population floor below, each of which is a location households live in and shop from, whether or not it contains a store of its own.

Table 18: From the collected data to the estimation sample

	Records	Establishments	Barcodes
<i>A. Grocery sample (descriptive)</i>			
Grocery records, HS-2 chapters 01–24	14,732,991	5,042	–
carrying a barcode	14,689,600	3,630	125,032
excluding single-barcode sellers	14,687,078	2,923	124,995
<i>B. Estimation sample</i>			
Barcode × establishment cells, all products	6,685,255	–	–
grocery barcodes with a quality estimate	1,656,480	3,488	–
less price-recording errors	1,607,975	–	–
less stores without a census tract	1,567,342	3,374	–
less stores under 20 barcodes	1,561,793	1,958	122,427

Notes: Panel A counts transaction records between 23 May 2025 and 23 November 2025. Its second row drops records carrying no barcode, which also removes the establishments whose grocery activity is entirely unbarcoded; its third row is the grocery sample of Table 1. Panel B counts barcode–establishment cells, not transactions, since the estimation reads the collapsed cell file, and then the cells that survive each restriction. Establishment counts in Panel B are the stores present at that step, and the barcode count is the number of distinct barcodes in the final estimation sample.

One further restriction applies on the residential side. Census tracts with fewer than 100 occupied households, or with no recorded income, are excluded from the model and from every tract-level exhibit, since tract income and its dispersion are not meaningfully measured in them.

The store network changes little inside the window, so the model treats it as fixed. The firm register dates both openings and closures, and of the 2,923 establishments in the grocery sample, 44 opened inside the window (1.5%) and 6 closed (0.2%).

A.5 Coverage against outside sources

Comparing prices and assortments across neighborhoods of one city requires the stores in that city to be enumerated and placed. Barcode-level research on retail prices typically uses two Nielsen products, and neither is built to do that.

The Retail Scanner data is supplied by participating retail chains, so a store enters when its chain agrees to take part. The sample is neither a census nor a random draw, the retailer's identity is masked, and the finest geography released to researchers is the store's three-digit postcode area and county (Kilts Center for Marketing, 2025). A store that belongs to no chain has no route in.

The Consumer Panel is built on a sample of households, who scan their purchases at home. A store is observed only when a panelist shops there, and goods a store stocks but no panelist bought leave no record, so what a store offers has to be inferred from what was bought.

Both restrictions cost more in Brazil than in the United States, where the independent segment is a far smaller share of grocery sales (Section 2).

The coverage of sellers is checked against Google's Places Aggregate service, which reports the number of operating establishments of a given type in a named municipality. It was queried in September 2026, after the price window closed, so it counts the stores operating at that date and the comparison is not of the same period. For Maceió it returns 1,421 across the supermarket, grocery and convenience categories (Google, 2026). The two sides are put on the same footing before the counts are set against each other. Google is filtered to establishments currently operating, the sample is cut to the municipality, not the metropolitan area, and the incidental sellers of Table 2 are dropped, since Google's grocery types exclude the gas stations and restaurants that ring up the occasional grocery item. That leaves 1,702 establishments, 120% as many. The two sources share no identifier, so this compares counts and never matches one store to another. The invoice count exceeds the directory's. Since the comparison is of counts for the city as a whole, it is not a match rate, it does not measure coverage neighborhood by neighborhood, and it says nothing about sellers that neither source records.

A fixed shop that takes electronic payment is also visible to the authority that administers the invoice. Payment institutions report each seller's card and instant-payment receipts to the state tax administrations, identified by tax number and including sellers absent from the state

taxpayer register (Conselho Nacional de Política Fazendária, 2016). This matters for the lower boundary of the seller universe discussed in Section 2.

A.6 Tract income composition

The 2022 tract file reports the mean and variance of the income of household heads but not its distribution, so the share of a tract’s households in each city-wide income quartile is estimated rather than observed. The quartile cutoffs are those of a household-weighted mixture of lognormal distributions fitted to each tract’s two moments (Appendix D), so that groups and tracts are defined on one distribution.

The relationship between the two moments a geography publishes and the composition it actually has is estimated on the national 2022 sample, which observes both for every weighting area in the country. Maceió’s own weighting areas are excluded from the fit, so the map that predicts the city’s tracts is estimated without them. For each group,

$$\text{logit}(s_g) = a_g + b_g \log\left(\frac{\bar{y}_\ell}{\bar{y}}\right) + c_g \sigma_\ell,$$

where $\bar{y}_\ell/\bar{y}$ is the area’s mean income relative to its municipality and σ_ℓ is the dispersion implied by its two moments, and each area counts equally in the fit. A Maceió tract enters the map the same way, relative to its own municipality. Fitted on 6,472 areas across 525 cities and scored leave-one-city-out, the map misses by a median of 2.1 percentage points per group, with a bottom-quartile bias of -1.8 points. Predicted shares are rescaled so that each group holds a quarter of the city’s households, which makes them quartile shares.

Households are allocated to the four groups within each tract in two steps. Group shares are predicted from the tract’s two moments by the map above, then tilted, by the minimum-information exponential reweighting consistent with the constraint, so that the group shares and the group mean incomes reproduce the tract’s recorded mean income. The tilt corrects something the prediction cannot. The model charges every household in group g the city-wide group mean $\bar{y}_g$, while a poor tract’s group- g households earn less than that, so predicted and recorded tract mean income do not agree until the shares are reconciled. The reconciliation moves shares by a median of 5 percentage points. For 132 tracts no mixture of the four group means reproduces the recorded mean, mostly because it lies outside their range, and those tracts keep their predicted shares. The reconciled shares therefore carry the differences in income within a group across tracts, which the model has no other place to put, so tracts differ more in their group shares than they would if group incomes varied by tract.

Three limits apply. The map is fitted on weighting areas and applied to census tracts, so it assumes the relationship transfers across unit size, while leave-one-city-out tests transfer across cities only. It is fitted on household income and applied to the moments of household-head

income, the only income the tract file reports, so it also assumes that the two measures place an area in the city’s distribution alike. And the shares enter the model through a single mean income per group, so the reconciliation, and not the imputation, is why realized group shares depart from a quarter each.

A.7 Product quality

This section documents the barcode quality index of Section 2: how it is constructed, what it correlates with outside the price data, and what the decomposition assumes.

Construction and connectedness. Let $G = (I, F, E)$ be the bipartite graph with nodes I (barcodes) and F (establishments) and an edge for each barcode observed at each establishment during the sample window. Barcode and establishment effects are identified only within connected components of G , since prices in disjoint components share no common reference. The largest connected set retains 99.8% of barcodes, 98.8% of establishments and 100.0% of observations, so the connected set and the estimation sample are effectively the same object. 54.8% of barcodes are stocked in more than one store, and those barcodes carry 98.6% of observations, so the graph is tied together by a majority of barcodes and an overwhelming majority of the price data.

Variance decomposition and reliability. Table 19 decomposes the within-category variance of log prices into barcode and establishment components. Each estimated effect is the true effect plus estimation error, so the variance of the estimates is the variance of the true effects plus the variance of that error. The plug-in columns substitute the estimated effects into the variance formula and so overstate how dispersed quality is. The error is largest for barcodes stocked in few establishments, whose effect is estimated from few prices, which is why the bias is known as limited-mobility bias. I report a correction alongside it. The correction is a split-sample estimator. The decomposition is run on two random halves of the data, and the cross-half covariance of the effect estimates recovers the bias-free variance, since estimation error is independent across halves (Sorkin, 2018). It is an alternative to the leave-out correction of Kline et al. (2020), which requires the same leave-one-out connected set, and Kline (2024) surveys both. The correction is modest, as the reliability of the index, computed on a single set of barcodes below, shows.

A cross-half covariance exists only for barcodes stocked in both halves, so the bias-corrected column is estimated on the barcodes carried by more than one store and the count rows of Table 19 report that smaller sample separately. The reliability of the index, 0.99, divides the bias-corrected variance by the plug-in variance over those same barcodes, so that numerator and denominator describe one population. It is therefore the reliability of the index where the index is identified, and it is silent about barcodes seen at a single store, whose effect is fitted to one price and carries no independent check.

Table 19: Variance decomposition of log prices

	Full sample	Connected set	
	Plug-in	Plug-in	Bias-corrected
Variance of barcode quality, $\text{Var}(\delta_i)$	0.525	0.525	0.502
Variance of establishment effect, $\text{Var}(\theta_f)$	0.019	0.019	0.018
Sorting covariance, $\text{Cov}(\delta_i, \theta_f)$	0.004	0.004	0.004
Sorting correlation, $\text{Corr}(\delta_i, \theta_f)$	0.043	0.043	0.045
Variance of log price, within category	0.611		0.611
Barcode–establishment cells (millions)	1.65		1.65
Underlying transactions (millions)	14.7		14.7
Barcodes (millions)	0.12	0.12	0.06
Establishments (thousands)	2.9	2.9	2.7

Notes: Variance components of equation (1), weighted by the number of underlying price observations per barcode–establishment cell. The connected set is the largest connected component of the barcode–establishment graph. Plug-in moments use the estimated fixed effects directly and overstate dispersion because of estimation error (limited-mobility bias). The bias-corrected column uses a split-sample estimator in which the decomposition is estimated on two random halves of the data and moments are formed from cross-half covariances, which are free of estimation-error bias because sampling error is independent across halves.

What the index implies for store sorting. The correlation between barcode and establishment effects is 0.04, so the stores that charge more for a given good are not the stores that carry better goods. That is a statement about price levels, and it is compatible with the sorting the motivating facts document, which runs through a different margin. Correlating each store’s mean stocked quality with its estimated price level gives 0.04, and with the income of its neighborhood 0.13, while the price level itself rises with neighborhood income at 0.20. Richer neighborhoods are served by stores that charge more for the identical barcode, and only slightly by stores whose shelves are on average better. The variety difference across neighborhood income that Section 3 documents is in the number of premium barcodes on offer and not in the average quality of a shelf, and that count enters the model through both the barcode-variety and the quality terms of (12).

$\text{Corr}(\delta_i, \theta_f)$ is also an observation-weighted moment over barcode–establishment cells, and 98.6% of observations sit on barcodes stocked in more than one store. Barcodes that every store carries contribute the same quality at every price level, so they pull the correlation toward zero mechanically.

Rank stability. The establishment sample is split at random into two halves and the decomposition is estimated separately on each, so the two resulting barcode effects are statistically independent. Figure 8 plots the resulting quality percentiles against each other for the 55,812 barcodes stocked in both halves. Points cluster around the 45-degree line, with a rank correlation of 0.902 on the connected set (0.912 on the full sample) and a linear correlation of 0.884, so the ranking is not driven by any particular set of establishments. The split is repeated five times

and the figures are means across repetitions; the rank correlation varies by 0.011 across them, and the bias-corrected variance of barcode quality by 0.0003. That variation is the benchmark for the store-format and price-level splits below, since it is what two arbitrary halves of the same data produce.




Figure 8: Rank stability of barcode quality (connected sample). Rank correlation 0.902.

External validation. I match barcodes to Open Food Facts and relate the estimated quality to attributes that play no part in its construction. The matched set skews toward branded packaged goods, and its mean quality index across barcodes, 0.32, is close to the 0.33 of all barcodes; both are unweighted means of an index centred on its observation-weighted mean, so neither is zero. Table 20 reports the association between estimated quality and each attribute within HS-8 categories, so that products are compared only against close substitutes, and a second time controlling for the size of the package.

Products carrying an independent quality certification score well above their category peers. A barcode with an organic label scores 0.30 log points higher than an otherwise comparable barcode in the same category, a price premium of about 36%. Against a within-category standard deviation of 0.82 log points in the matched set, that is a premium of 0.37 standard deviations. Protected-origin, fair-trade, and artisanal designations score almost as high, and branded products command about 19% over unbranded ones. The certification premia survive the size control, at 0.32 log points for organic, while the brand premium falls to 0.11 log points once package size is held fixed, so part of what brand carries in the index is package format and not the name.

An index built only from prices therefore lines up with the attributes an independent certifier judges a product on. What this cannot settle is the split between cost and margin, since certification raises what a product costs to make as well as what buyers will pay for it. It establishes that the index tracks something real about the product, not that the premium is willingness to pay.

Two properties of the index follow from its being a premium on the package as posted. Package size alone accounts for 29% of its within-category variance, with an elasticity of 0.43, so a larger pack of the same good scores higher. And because some lines are billed per unit and others by weight, a within-category comparison can mix the two bases. Both matter for cardinal quality. Mixed billing does not reorder products within a billing basis, but package size does, so the percentile descriptions rank a larger pack of the same good above a smaller one.

Table 20: External validation of the quality index against Open Food Facts attributes

	Effect on log quality index $\hat{\delta}_i$	
	Within category	+ package size
Carries an organic label	0.30*** (0.07)	0.32*** (0.09)
Carries a premium designation	0.32*** (0.07)	0.35*** (0.09)
Is branded (vs. unbranded)	0.17*** (0.04)	0.11** (0.05)

Notes: Each row is a separate regression of the barcode quality index $\hat{\delta}_i$ on the listed Open Food Facts attribute, with HS-8 product-category fixed effects; standard errors clustered by category in parentheses. The sample is the 15,143 barcodes matched to Open Food Facts; a barcode whose Open Food Facts record carries no label list is treated as not carrying the label, not as missing. The attributes are carried by 69 barcodes (organic), 99 (premium designation) and 10,848 (branded), so the organic and premium estimates come from the fewest barcodes in the table. All three attributes are indicators, so their coefficients are the difference in log quality between barcodes with and without the attribute. A premium designation is an organic, protected-designation-of-origin, protected-geographical-indication, fair-trade, or artisanal tag. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

What the index contains. The index carries what makes a barcode more expensive for reasons other than the good itself. Manufacturer identity accounts for 43% of its within-category variance, and barcodes issued abroad command 26% over domestic ones in the same category. Brand stays in the measure because the index records what households pay for a barcode, and demand in the model is over barcodes. Bronnenberg et al. (2015) find that better-informed shoppers pay much smaller brand premia, so part of the brand component reflects what shoppers believe about a product and not what it contains. The index is the premium a barcode commands in prices, not quality inferred from quantities within an estimated demand system. The components that can be removed do not distort the decomposition for the rest of the sample. Store-assigned barcodes, the prefix-2 codes a store issues itself, denote different products at different stores, so their comparability across stores is fictional; re-estimating without them, or on domestically issued barcodes alone, reproduces the index with rank correlations of 0.999 and 0.994. Manufacturer identity admits no such check, since dropping brand means dropping the

products. Whether the import premium is quality or landed cost is not settled here, and to the extent that it is cost it overstates the welfare the model assigns to imported goods.

Separability across store types. The decomposition requires that a store raise the price of every barcode it carries by the same proportion, whatever the barcode's quality, so that the barcode effect is the same wherever the product is stocked. What matters for the measure is not whether interactions between quality and location exist, but whether they reorder products, since the paper describes assortments through quality percentiles. This paragraph and the next put that question to the estimates.

To assess additive separability across establishment types, I classify establishments by their primary CNAE activity code into food retail (471, general food retail, and 472, specialized food retail), food service (56) and other retail (pharmacies, mixed stores and the like), and estimate the decomposition separately within each format's largest connected subgraph. Table 21 reports pairwise rank correlations of the quality index for barcodes appearing in both subgraphs. The informative comparison is food retail against other retail, the two formats sharing a large common assortment of 27,008 barcodes, where the two rankings agree at 0.80, against the 0.902 that two random halves of the same data produce. The lower correlations involving food service are computed on a few thousand common barcodes with a narrow and atypical assortment, so they are noisier and say little about separability among the formats that carry the bulk of the data.

The last row splits establishments by their estimated price level instead of their activity code, which is the high-end against ordinary contrast rather than a difference in what the store sells, and the two rankings there agree at 0.77.

Ranks are silent about scale, so the table also reports the slope of one split's quality on the other's. If one kind of store amplified quality premia, the two indices would be proportional rather than equal, and a rank correlation is one for any factor of proportionality; the slope is the statistic that moves. A raw slope is attenuated by estimation error in the regressor. Instrumenting one group's estimate with an independent second estimate from the same group removes that attenuation and puts the factor at 0.90 (0.01) on 21,386 barcodes, against 0.99 for the same design run within a group. Quality differences are compressed by about a tenth in high-price stores relative to low-price ones, in the opposite direction to the amplification the concern supposes. The rank correlations across store types are below the random-split benchmark, so percentile placements agree less across kinds of store than across random halves of the sample.

Additivity of store-pair gaps. This is the direct test of additive separability. Under an additive decomposition the average price gap between two stores, taken over the barcodes they both carry, depends only on the two stores' price levels and not on which barcodes those happen to be, since the barcode effect is common to the two listings and drops out. A markup that varied with quality would break that, by making the gap depend on the shared assortment, so the fit of

Table 21: Barcode quality recovered from different kinds of store

Split A	Split B	Rank correlation	Slope	Common barcodes
Food retail	Food service	0.674	0.743	5,561
Food retail	Other retail	0.801	0.767	27,008
Food service	Other retail	0.755	0.699	4,754
High price level	Low price level	0.769	0.768	38,733

Notes: The decomposition is estimated separately within each split, on that split’s largest connected subgraph, and the two resulting barcode effects are compared on the barcodes common to both. The first three rows split establishments by CNAE activity class; the last splits them at the median estimated establishment price level, which is the high-end against ordinary contrast rather than a difference in activity. Rank correlation is Spearman’s. Slope is from a regression of the first split’s barcode effect on the second’s, and is attenuated by estimation error in the second, so it is read against the 0.90 that the same comparison gives when instrumented and not against one. Two random halves of the establishment sample agree at 0.902.

the observed gap to the predicted one measures how much of the cross-store variation a single per-store shift can carry.

For every unordered pair of establishments (f, f') connected by at least k common barcodes, the observed log-price gap $\bar{\Delta}_{ff'}$ is the average across common barcodes of $\log p_{ifn} - \log p_{if'n}$, in which the barcode effect cancels algebraically. Figure 9 plots that observed gap against the difference the estimated establishment effects predict, for every pair of stores linked by common barcodes; Table 22 reports the noise-corrected R^2 at minimum edge sizes of two, three and four barcodes. The correction subtracts the estimated within-edge sampling variance from both numerator and denominator, following Kline (2024). To keep the all-pairs computation tractable, very high-degree barcodes are capped at a random subsample of thirty listings, which leaves each edge mean unbiased in expectation, though it also changes which pairs clear the minimum common-barcode threshold.

The gap itself can also be regressed on the quality of the barcodes it is measured on, which tests the alternative directly. Under a separable decomposition that slope is zero for every pair, whatever the shared barcodes happen to be. Estimating it for every pair of stores sharing at least three barcodes, the slopes are dispersed by 1.453, against 1.474 from permuting quality among each pair’s shared barcodes, which leaves the gaps untouched and breaks only the relationship being tested. The observed dispersion is no larger than the permuted one, so this test finds no quality-dependence in store-pair gaps. Slopes estimated on as few as three shared barcodes are noisy, so the test detects only dispersion well beyond sampling noise. The comparison of high- and low-price stores above has more power against the same alternative, and it finds quality premia compressed in high-price stores. A single per-store shift accounts for 0.389 of the variation in the gaps themselves, unchanged at minimum edge sizes of two, three and four shared barcodes.




Figure 9: Observed versus predicted log-price gaps between stores. For each pair of stores that share barcodes, the observed gap is the average, over those shared barcodes, of how much more one store charges than the other for the same product. The predicted gap is the difference in the two stores’ estimated establishment effects. Each dot is the mean observed gap within a quantile-spaced bin of the predicted gap, and the dashed line is the 45-degree reference implied by additivity, not a fit.

Table 22: Fit of the additive prediction to store-pair price gaps

Min. common barcodes	Edges	R^2 (raw)	R^2 (corrected)	RMSE	SD(edge)
2	404,565	0.389	0.389	0.422	0.519
3	300,522	0.403	0.404	0.369	0.463
4	240,200	0.411	0.412	0.338	0.431

Notes: One observation per pair of establishments sharing at least the stated number of barcodes, weighted by that number. The raw R^2 is from regressing the observed mean log-price gap on the gap the estimated establishment effects predict. The corrected column subtracts the sampling variance of each pair’s gap, built from the within-cell variance of log price over the transactions behind each cell median, from both numerator and denominator. That correction is small because posted prices move little within a barcode–establishment cell, so the two columns are close. RMSE is of the residual gap and SD(edge) the dispersion of the observed gap, both in log points. Establishment effects are estimated on the full connected sample, so this is an in-sample fit.

The store-format split and the edge test together address a second assumption, that which barcodes a store carries is uninformative about the price residual. If a premium barcode entered a lower-end outlet only on atypically favorable wholesale terms, its price there would reflect that arrangement, not a premium the barcode carries everywhere. Taking medians over the sample window absorbs the transient form of this. The persistent form cannot be told apart from a quality-dependent markup, since with one observation per barcode–establishment cell a

component specific to the pairing is the same object whether it originates in the match or in the stocking decision. The event-study diagnostics that assess exogenous mobility in panel designs (Card et al., 2018) require a panel and are unavailable here.

Establishment price levels. The establishment effect θ_f absorbs everything about where a product is sold, including the store’s costs, its markups, its amenities, and the price level of its neighborhood. Each establishment sits at a fixed location, so a separate neighborhood effect would not be identified from it, and neighborhood income enters here only as a diagnostic and never as a regressor.

Figures 10 and 11 plot the establishment price level against the income of the surrounding neighborhood and against the number of barcodes carried. Store price levels rise with neighborhood income, with stores in the top income ventile charging 0.18 log points above those in the bottom ventile, the establishment-level analog of the firm-pays-high-wages fact in worker–firm settings. It relates to the response of retail prices to local housing wealth that Stroebe and Vavra (2019) document. That gap is an unweighted comparison across establishments, so it describes the typical store and not the typical transaction, and it is not the same object as the observation-weighted $\text{Var}(\theta_f)$ in Table 19. Price levels also fall with store size. Location and store identity therefore carry substantial pricing information, which the decomposition has to absorb before the barcode effect can be read as an attribute of the product itself.




Figure 10: Establishment price levels by neighborhood income ventile (weighting-area mean income). Each dot is the mean establishment price level within one of 20 equal-count income bins, placed at the bin midpoint; vertical bars are 95% confidence intervals for the bin mean. Higher-income neighborhoods support higher establishment price levels.




Figure 11: Establishment price levels by establishment size ventile (number of barcodes carried, log scale). Each dot is the mean establishment price level within one of 20 equal-count size bins; vertical bars are 95% confidence intervals for the bin mean.

Attenuation from the generated regressor. The quality index is estimated, so any regression placing it on the right-hand side inherits classical measurement error, which attenuates the coefficient toward zero by the reliability of the measure. A split-sample instrumental-variables design measures that attenuation directly. Establishments are split into two random halves, the decomposition is re-estimated on each half's largest connected set, and the resulting independent half-sample estimates of barcode quality ($\hat{\delta}_i^A$ and $\hat{\delta}_i^B$) are available for every barcode stocked in both (58,683 barcodes). Since estimation error is independent across halves, one half's estimate is a valid instrument for the other's, and the IV coefficient is free of attenuation bias (Ashenfelter and Krueger, 1994; Sorkin, 2018).

Table 23 implements the check on a barcode-level regression of the barcode's income exposure on its quality,

$$\bar{r}_i = a + b \hat{\delta}_i + u_i, \quad \bar{r}_i \equiv \frac{\sum_f n_{if} r_f}{\sum_f n_{if}}, \quad (22)$$

where r_f is the percentile rank of mean income in establishment f 's weighting area, n_{if} is the number of price observations of barcode i at establishment f , and both sums run over the establishments that carry i . The dependent variable is a percentile rank, so in the level columns b is the shift in the average income rank of a barcode's stockists per log point of quality, while in the percentile columns the quality measure is replaced by its own within-sample rank and both sides are ranks. Instrumenting undoes the attenuation in the half-sample measure. The IV coefficient (0.031) rises above the half-sample OLS coefficient (0.029). The half-sample measure is the

noisier of the two, being estimated on half the establishments, while the full-sample index has the higher reliability reported in Table 19 for barcodes stocked in more than one establishment. For those barcodes the gap therefore bounds the attenuation from above instead of measuring it. In this linear regression the error biases the coefficient toward zero. The argument does not sign the bias in the percentile profiles of Section 3, where estimation error changes which barcodes fall in each percentile, and the seller-count and price elasticities do not use quality at all. The counterfactuals use cardinal quality through a nonlinear transformation, so the same error does not attenuate there and the bias is not signed by this argument. Cardinal quality is also estimated least precisely for barcodes stocked in few establishments, which the reliability correction does not cover, so that imprecision matters for the counterfactuals more than for the descriptive results.

Table 23: Split-sample IV check for attenuation from the generated quality regressor

	Quality in levels			Quality in percentiles		
	OLS full	OLS split A	IV A inst. B	OLS full	OLS split A	IV A inst. B
Coefficient on quality	0.0296 (0.0008)	0.0286 (0.0008)	0.0308 (0.0009)	0.0993 (0.0028)	0.1003 (0.0028)	0.1061 (0.0030)
First-stage coefficient			0.9225 (0.0029)			0.9232 (0.0017)

Notes: The dependent variable is the barcode’s income exposure, the listing-weighted mean income percentile of the establishments carrying it. “Full” uses the quality measure estimated on the full sample, “split A” uses the measure estimated on a random half of establishments, and “A inst. B” instruments the half-A measure with the independent half-B measure. Percentile columns use within-sample percentile ranks of each measure. Heteroskedasticity-robust standard errors in parentheses. The sample comprises 58,683 barcodes listed in both halves.

A.8 Complementary sources

Census and spatial layers. I use 2022 census tract polygons (*setores*), official address points, and 2010 weighting areas (*áreas de ponderação*) for Alagoas. All spatial layers are transformed to EPSG:4326. Tract-level aggregates from the 2022 Census supply population, households, and the mean and variance of income.

Household income is released only at the weighting area, never at the tract, in either census year; the tract file reports the income of household heads. This is why the tract income composition is estimated rather than read. The 2022 sample microdata, released in 2026 under controlled access (Instituto Brasileiro de Geografia e Estatística, 2026), report it for all weighting areas in the country and are aggregated here to one record per area: households, mean and variance of household income, and the household count in each within-city income quartile. The income-composition map is fitted on that national file.

The 2010 sample is used for two purposes. The 2022 questionnaire did not ask what rent a household pays, so 2010 remains the most recent census that measures rent, and the weighting-area rents the model uses are taken from it (Instituto Brasileiro de Geografia e Estatística, 2012). Its national file also supplies the weighting-area income and rent on which the sorting elasticity is estimated (Section 5).

Consumption survey. The 2017–2018 consumption-survey microdata are parsed from fixed-width files. I construct household income, deflate and annualize expenditures, aggregate to product codes, and compute expenditure shares by income quartile, including the housing share ρ_g . The food diary additionally records a standardized quantity in kilograms, which yields unit values by product and region; Section 5 uses these to estimate the two demand elasticities in a first stage.

Mapping the consumption survey to HS-8. Survey product codes are linked to HS-8 categories through two official translators, one from survey codes to the national accounts classification and one from that classification to HS-8. The first is published for the 2009 survey rather than the 2017–2018 one, so 2017–2018 codes are carried onto it by matching product descriptions; this bridge, not the translator itself, is the weakest link in the chain. Products the translators do not reach are mapped by description-based matching against the official HS-8 list. Mapping weights are normalized within HS-8, two HS-8 codes that sit under metal-products headings but carry grocery items in the price data are assigned by hand to their survey food codes, and one survey code covering real estate is dropped with the remaining weights renormalized. The result is a weighted mapping from HS-8 to survey codes with confidence flags.

National establishment register. The Receita Federal publishes the full establishment register in bulk. Each record carries the firm identifier, the activity code (CNAE 2.0), the postcode, the opening date, and, for establishments no longer active, the date they closed. I extract the candidate set with the activity prefixes 47, 1091, 463 and 469, then define grocery retail as fourteen subclasses, namely hypermarkets and supermarkets (4711301, 4711302), minimarkets and grocery stores (4712100), bakeries selling at retail (4721102) and producing their own goods (1091102), dairy and cold cuts (4721103), confectionery (4721104), butchers (4722901), fishmongers (4722902), drinks (4723700), fruit and vegetables (4724500), other food retail (4729699), convenience stores (4729602), and food-predominant general wholesale, the *atacarejo* format that sells to households (4691500). Food wholesale proper (4639701, 4639702) is excluded, since those classes are populated by manufacturers’ distribution arms, not by shops a household can walk into.

Establishments are placed in weighting areas through their postcode, using the official address points. The register’s public files keep establishments whose registration has been closed,

with the date of closure, so the set of shops trading in any earlier year is recovered by keeping those opened on or before that year and not closed before it, which avoids the survivor bias that a register of currently-active firms would introduce.

PNAD Contínua. The PNAD Contínua 2019 microdata measure how much of Maceió’s food retail is run by registered businesses. The question is put to the people who answer for a business, employers and own-account workers in the food-retail activity classes, and asks whether the business holds a CNPJ. The two groups are reported separately, since they describe different sellers. An employer runs a business with employees, which is what an establishment in the price data is, and 92.4% of the 36 employer respondents report a CNPJ. Own-account workers in food retail are mostly street vendors and market traders, and 15.3% of the 189 respondents report one. Pooling the two would describe neither.

Commercial floorspace. Store size in square metres is not recorded in any of the administrative sources. I construct it from OpenStreetMap building multipolygons for Brazil (OpenStreetMap contributors, 2025). Where an establishment is matched to its own building footprint, the nearest such polygon within 100 meters (15% of cases, Table 16), that footprint is used. Otherwise, commercial and shop polygons are summed within the establishment’s weighting area and divided by the number of shop polygons the area contains, giving the average mapped shop in that area. The divisor is the count of mapped shops and not the count of sample establishments. The measure therefore gives the average size of a shop OpenStreetMap knows about, not a share of commercial space allocated across the sample. Establishments in weighting areas with no mapped shop polygon at all receive a single dwelling-unit equivalent, the smallest size the model’s housing units admit.

The measure is a weighting-area average and not a store-level measurement, so it carries no within-area variation in store size except through the matched footprints. It enters the paper in two places: as each store’s footprint h_f , which sets the commercial claim on floorspace in (20) and the rent a store must cover in the counterfactuals, and in the comparison of markups by floorspace quartile in Appendix D.

Open Food Facts. The product dump is pinned to a fixed revision, so the matched attribute set is reproducible as of the download date. Matching is on the raw barcode with no harmonization, since a record that does not match on the barcode cannot be attributed to a product with confidence. Appendix A.7 reports the validation these attributes support.

A.9 Data availability and replication

The price data is collected from a public API operated by SEFAZ-AL and carries no personal identifiers (SEFAZ-AL, 2026b). The consumption-survey and PNAD Contínua microdata are

published by IBGE, and the Receita Federal firm registry is public. The 2022 census sample microdata are the one exception to an otherwise fully public set of sources: IBGE releases them under a controlled-access tier, and they are therefore not redistributable. The deposit ships the weighting-area aggregate estimated from them, together with the city’s income-quartile cutoffs and group mean incomes, which are the only objects any result depends on, so a replicator reproduces every number without applying for the raw files. Open Food Facts is released under an open database license, and OpenStreetMap under the Open Database License.

Three constraints on exact replication apply. First, the collection requires an API token issued by SEFAZ-AL and the geocoding requires a Google Maps key; neither is distributed with the code, and both are read from the environment and not from any file in the repository. Second, the API returns the transactions recorded at the time of the call, so a sweep run today will not reproduce a sweep run during the sample window: the collected data, not the collection code, is the reproducible object. Third, Google Places responses are subject to Google’s terms, so the deposit carries the coordinates and match confidences derived from them rather than the responses themselves.

The deposit therefore contains the collection, cleaning, and estimation code; the sweep calendar and the manifest identifying which sweeps enter the analysis; and the assembled price panel, on which every result in the paper can be reproduced without re-running the collection.

B Motivating Facts: Additional Results

B.1 Seller counts at other radii

The seller count of Section 3 uses a 1.0 km radius. The model of Section 4 has no cutoff. Distance enters continuously through the trip cost τ , which is not set until Section 5, so the 1.0 km count stands in for a distance-weighted aggregate. Figure 12 repeats the exercise at 0.5 and 2.0 km. The income relationship holds at every cutoff. Tracts in the highest-income quartile have 2.4 times as many sellers as those in the lowest within 0.5 km, 2.5 within 1.0 and 2.2 within 2.0 km.

B.2 Variety and neighborhood income

Section 3 reports the income elasticity of the total number of barcodes a tract carries, estimated by Poisson pseudo-maximum likelihood like the profiles, 0.310 with a robust standard error of 0.075 across the 857 tracts that hold at least one seller. Figure 13 plots the count behind it.

B.3 Measurement error in the quality percentile

The quality percentile q sits on the right-hand side of (3) and is itself estimated, so the profile is attenuated. Pooling the percentiles and interacting log income with q puts the rise from the bottom of the quality distribution to the top at 0.602 (0.057). Splitting the establishment sample




Figure 12: Grocery sellers within three radii, by neighborhood income

Notes: The exercise of Figure 3 repeated at 0.5, 1.0 and 2.0 km. Each point is a bin mean across tracts with a 95% confidence interval; the line is a loess fit. Vertical scales differ across panels because the counts do. The income elasticity is positive at every cutoff.

in half yields two percentile measures built from disjoint sets of stores, and instrumenting one with the other raises it to 0.695 (0.094).⁷ Instrumenting roughly doubles the interaction on the cells that hold a barcode.

The price interaction is unmoved by the same correction. Instrumenting q the same way leaves it at 0.005 (0.019) over 1,164,990 barcode–tract cells, against 0.001 (0.019) uninstrumented, so the absence of a quality dimension in prices is not measurement error in q . That estimate is a linear instrumental-variables regression, since its outcome is a log price rather than a count.

B.4 The price elasticity in detail

I estimate (4) separately for barcodes in each quality percentile q ,

$$\log p_{i\ell} = \gamma_q^p \log y_\ell + \omega_i + u_{i\ell}, \quad i \in G_q, \quad (23)$$

where G_q is the set of barcodes at quality percentile q , and Figure 14 plots $\hat{\gamma}_q^p$ against q .

The elasticity is positive at every percentile and drifts downward. It is 0.060 (standard error 0.008) at the tenth percentile, 0.050 (0.008) at the median and 0.081 (0.028) at the ninetieth. The drift is within sampling error throughout. Interacting log income with the quality percentile over the whole distribution puts the interaction at 0.001 (0.019), and the top percentiles are the

⁷Over 61,476 tract–percentile cells. Poisson carries no instrumental-variables estimator, so the interaction is projected on its instrument and the first-stage residual enters the count regression; standard errors come from a bootstrap that resamples tracts. Both estimates use cells holding at least one barcode, since an empty cell has no barcode from which to build the instrument. One half is a valid instrument for the other to the extent that their measurement errors are independent, which disjoint store sets make plausible but do not guarantee, since both rank the same barcodes.




Figure 13: Barcodes sold in a neighborhood, by neighborhood income

Notes: Distinct barcodes sold anywhere in a census tract against log tract income. Each point is the log of the mean count across an equal-count bin of income, the scale on which the Poisson elasticity is a slope; bars are 95% confidence intervals and the line is a loess fit. The count is the union across every establishment in the tract, so a barcode stocked by several local stores is counted once.

least precisely estimated part of the profile, since a quality percentile contains few barcodes that are also widely stocked.

The elasticity also survives a control for local competition. To show this, I examine the within-barcode price elasticity after conditioning on the number of competing sellers. Let $n_{i\ell}$ denote the number of distinct retailers carrying barcode i in tract ℓ , and estimate

$$\log p_{i\ell} = \beta \log n_{i\ell} + \omega_i + \eta_{i\ell}, \quad (24)$$

where ω_i is a barcode fixed effect, which absorbs category effects since each barcode belongs to one HS-8 category. The residual $\hat{\eta}_{i\ell}$ isolates the component of the log price that cannot be attributed to the product's identity or to the number of local competitors. The regression is descriptive. The number of local retailers is itself endogenous to neighborhood income and demand, so the coefficient on $\log n_{i\ell}$ is not a causal effect of competition on prices. The exercise asks only whether the price-income elasticity survives after conditioning on an observable proxy for local market structure.

Binned the same way as Figure 5, the residual $\hat{\eta}_{i\ell}$ rises by 5.3 log points from the lowest income bin to the highest, against 5.1 without the competition control. The count is of rival stores within the tract, not within shopping reach, and the comparison is of the binned residual




Figure 14: Income elasticity of price, by quality percentile

Notes: Each point is $\hat{\gamma}_q^p$ from equation (23), estimated separately on the barcodes at each quality percentile, net of barcode fixed effects and clustered by census tract; bars are 95% confidence intervals and the line is a loess fit. The elasticity is positive throughout and drifts down across the quality distribution, but the drift is within sampling error. The interaction with the percentile over the whole distribution is 0.001 (0.019).

profiles, not the income coefficient of a regression that includes the count.

The elasticity also survives a restriction to barcodes observed across the income distribution. A barcode is priced only where it is stocked, and availability rises with income, so the estimation sample is selected on a variable that moves with the regressor. The barcode fixed effect removes basket composition but not selection in which tracts carry a given barcode. Requiring a barcode to appear in both the bottom and the top quartile of tract income, which rules out an elasticity driven by barcodes sold at only one end of the income distribution, keeps 30,264 barcodes and 84.3% of the cells and gives 0.056 (0.006); requiring all four quartiles keeps 23,399 barcodes and 81.3% of the cells and gives 0.054 (0.006), against 0.056 (0.006) on the full sample. The restriction is on income span rather than on the number of tracts a barcode reaches, since only part of the city's tracts hold a sampled seller and a barcode confined to one end of the income distribution carries no information about the slope however many stores stock it.

B.5 The price and variety facts among independent stores

A concern about the price and variety facts is that they reflect which retail chains locate where, with prices and product menus set centrally and not in response to local demand. I address it by restricting the sample to independent stores. The first eight digits of the Brazilian CNPJ

identifier, the *radical*, identify the legal entity, so establishments sharing a radical belong to the same firm. An establishment is independent if its radical carries exactly one active retail establishment in the national firm register, so there is no headquarters that could coordinate prices or assortments across locations. The count is national rather than local, because a chain with a single store in Maceió still prices it centrally. Independents are 80.5% of the establishments that enter these regressions, against 80% of the grocery sample as a whole (Section 2).

Figure 15 re-estimates the two facts on this subsample, using the same specifications, equations (23) and (3), as the body. The two facts fare differently, and the difference is informative.

The price elasticity survives intact, at 0.069 (standard error 0.007) pooled and 0.056 (0.010) at the median quality percentile, against 0.056 (0.006) and 0.050 (0.008) in the full sample. The level requires no chain. Among independents the quality interaction is -0.068 (0.014), negative and precisely estimated, so within independent stores prices rise faster with neighborhood income for lower-quality barcodes. The absence of a quality dimension in prices holds for the full sample and not for independents.

The variety elasticity keeps its shape across quality and loses its level. Pooled over the quality distribution the independents elasticity is 0.035 (0.081) at the tract level, against 0.310 (0.075) in the full sample, and the profile still runs from -0.269 (0.087) at the fifth percentile to 0.768 (0.160) at the ninety-ninth. Within stores it runs from -0.574 (0.105) to 0.640 (0.181).

The level of the tract elasticity is not robust to dropping chains, while its quality profile is. A tract's variety is the union of what its establishments carry, so the subsample removes part of each tract's assortment, and this appendix does not measure how that removal varies with tract income.

C Model Appendix

C.1 The Price Index Decomposition

The grocery price index $\mathcal{P}_{\ell g}$ that a group- g household in neighborhood ℓ faces is built in Section 4.1 from three nested constant-elasticity aggregators, over stores, over categories within a store, and over barcodes within a category. This appendix shows that its log splits exactly into the trip cost the household expects to pay, three variety terms counting the stores it reaches, the categories those stores carry and the barcodes in each category, the average quality of what sits on those shelves, and a quality-weighted mean price. That split is (12). The tiers also group the terms by the decision they belong to, since the store tier delivers the two between-store terms and leaves a single index, the store price index, which the next two tiers open into the four within-store terms.

The three tiers are taken in turn. Each step replaces an index by an average of what it aggregates plus a variety term, and leaves a single index for the next step to expand. The last step also separates quality from price. Every step is an identity, so the terms add to the index exactly.




Figure 15: The price and variety facts among independent stores

Notes: The price and variety profiles re-estimated on independents alone, from the same code path and specifications as the full sample. The variety panel is Poisson on the completed grid and carries a tract and an establishment series, as in Figure 4; the price panel is a within-barcode regression on log price, where an unstocked barcode has no price and there are no zeros to keep. An establishment is independent if its CNPJ root carries one active retail establishment in the national firm register, so a chain with a single Maceió store counts as a chain. The price elasticity survives in level. The variety elasticity keeps its shape across quality and loses its level, for the compositional reason given in the text.

The step used at each tier. Let $P = (\sum_k x_k^{1-\sigma})^{\frac{1}{1-\sigma}}$ with $\sigma > 1$ and expenditure shares $s_k = x_k^{1-\sigma} / \sum_j x_j^{1-\sigma}$. The denominator of the share is $P^{1-\sigma}$, so $s_k = (x_k/P)^{1-\sigma}$, and taking logs inverts it, $\log x_k = \log P + \log s_k / (1 - \sigma)$. Weighting by s_k , summing over k and using $\sum_k s_k = 1$,

$$\log P = \sum_k s_k \log x_k + \frac{1}{\sigma - 1} \sum_k s_k \log s_k. \quad (25)$$

The index is the expenditure-weighted average of what it aggregates, plus a term that, for $\sigma > 1$, is negative whenever expenditure is spread across more than one alternative and zero otherwise. The condition holds at each tier, since $\sigma - 1 = 1/\lambda > 0$ for stores and $\sigma_b > \sigma_c > 1$ within them. That term is the variety gain, since $-\sum_k s_k \log s_k$ is the entropy of the shares and its exponential is the effective number of alternatives.

Stores. The store index (8) is constant-elasticity in the delivered index $e^{\tau d_{f\ell}} P_{fg}$ with $\sigma - 1 = 1/\lambda$, and its shares are the store-choice probabilities (9). The delivered index is a product, so its log is a sum and trip cost separates from the store's own prices. Applying (25),

$$\log \mathcal{P}_{\ell g} = \sum_f s_{f\ell g} \tau d_{f\ell} + \lambda \sum_f s_{f\ell g} \log s_{f\ell g} + \sum_f s_{f\ell g} \log P_{fg}. \quad (26)$$

The first term is proximity and the second is store variety, and together they are the between-store block. The third is the store price index the household faces once it has chosen, the within-

store block, and it is the only term that still holds an index.

Categories. The store's own index (10) is constant-elasticity in the category indices with elasticity σ_c and shares $s_{n|fg}$, so (25) applies again,

$$\log P_{fg} = \sum_n s_{n|fg} \log P_{fn,g} + \frac{1}{\sigma_c - 1} \sum_n s_{n|fg} \log s_{n|fg}.$$

Substituting into (26),

$$\begin{aligned} \log \mathcal{P}_{\ell g} &= \sum_f s_{f\ell g} \tau d_{f\ell} + \lambda \sum_f s_{f\ell g} \log s_{f\ell g} \\ &+ \frac{1}{\sigma_c - 1} \sum_{f,n} s_{f\ell g} s_{n|fg} \log s_{n|fg} + \sum_{f,n} s_{f\ell g} s_{n|fg} \log P_{fn,g}. \end{aligned} \quad (27)$$

The new term is category variety. Every barcode-tier term from here carries the weight $s_{f\ell g} s_{n|fg}$, the chance a group- g household in ℓ shops at f times the share of its grocery budget it spends on category n there.

Barcodes. Applying (25) a third time would leave the number of barcodes on the shelf buried inside their expenditure shares. Since that count is one of the margins to be measured, the category sum in (10) is factored instead,

$$P_{fn,g}^{1-\sigma_b} = N_{fn} \mu_{fn,g} \bar{p}_{fn,g}^{1-\sigma_b}, \quad (28)$$

where N_{fn} is the number of barcodes carried and

$$\mu_{fn,g} = \frac{1}{N_{fn}} \sum_{i \in \mathcal{I}_{fn}} \zeta_i^{\alpha \log y_g}, \quad \bar{p}_{fn,g}^{1-\sigma_b} = \frac{\sum_{i \in \mathcal{I}_{fn}} \zeta_i^{\alpha \log y_g} p_{ifn}^{1-\sigma_b}}{\sum_{i \in \mathcal{I}_{fn}} \zeta_i^{\alpha \log y_g}}$$

are the average taste weight the shelf carries and a quality-weighted average price. The factoring is an identity, since N_{fn} cancels the denominator of $\mu_{fn,g}$ and $\sum_i \zeta_i^{\alpha \log y_g}$ cancels the denominator of $\bar{p}_{fn,g}^{1-\sigma_b}$. Taking logs and dividing by $1 - \sigma_b$,

$$\log P_{fn,g} = \frac{\log N_{fn}}{1 - \sigma_b} + \frac{\log \mu_{fn,g}}{1 - \sigma_b} + \log \bar{p}_{fn,g}. \quad (29)$$

Substituting (29) into the last term of (27) splits it into barcode variety, quality and price, and the result is (12).

All three variety terms weakly lower the index, the two share terms because shares are at most one and the count term because $\sigma_b > 1$, and each is zero when there is a single store, category or barcode. They measure the effective number of stores a household reaches and of categories within a store, and the number of barcodes within a category. Each step is exact on its own, so the margins can be reported separately.

C.2 The Pricing First-Order Condition

A store in Section 4.2 sets one price per barcode, taking its rivals' prices and the residential allocation as given. Each establishment prices on its own account, including outlets of the same chain, so a chain does not internalize the shoppers one of its outlets sends to another. Establishments that belong to no chain are 80% of the grocery establishments in the sample (Section 2), so the assumption matters only for the remainder. This appendix derives the condition those prices satisfy. The markup on a barcode is the inverse of the elasticity of its own demand, widened by what the store recaptures on shelfmates in the same category and in its other categories, and narrowed by the shoppers a price rise sends to another store. That condition is (17), and the appendix gives closed forms for all four objects in it.

The route is to write the store's first-order condition in terms of a single object, the elasticity of each barcode's revenue to the price of barcode i . That object is computed in its three cases and substituted back, and the three cases become the three corrections.

The condition in revenue terms. Write X_{ifg} for what group g spends on barcode i at store f , the store's revenue from that barcode and group, which is the spending X_{fg} the store draws from the group split by the category share and the within-category share,

$$X_{ifg} = X_{fg} s_{n|fg} s_{i|fn,g}, \quad X_{fg} = \left(\sum_{\ell} N_{\ell g} s_{f\ell g} \right) E_g, \quad (30)$$

and $L_j = (p_{jfm} - c_{jm})/p_{jfm}$ for the markup on a barcode j in category m , written L_{ifn} in the body and abbreviated here because j ranges over the whole store. Variable profit in (13) is $\sum_g \sum_m \sum_{k \in \mathcal{I}_{fm}} L_k X_{kfg}$.

Raising p_{ifn} moves both factors of that product. It moves the markup on barcode i , and writing $L_i = 1 - c_{in}/p_{ifn}$ gives that response as

$$p_{ifn} \frac{\partial L_i}{\partial p_{ifn}} = \frac{c_{in}}{p_{ifn}} = 1 - L_i.$$

It also moves what every barcode in the store sells. Multiplying the first-order condition in p_{ifn}

through by p_{ifn} ,

$$\sum_g \left[(1 - L_i) X_{ifg} + L_i X_{ifg} \eta_{ii,g} + \sum_{j \neq i} L_j X_{jfg} \eta_{ji,g} \right] = 0, \quad \eta_{ji,g} \equiv \frac{\partial \log X_{jfg}}{\partial \log p_{ifn}}. \quad (31)$$

Only $\eta_{ji,g}$ is left to compute.

The revenue elasticities. By (30) each $\log X_{jfg}$ is the sum of three logs, the spending the store draws from the group, the category share and the within-category share, so $\eta_{ji,g}$ is the sum of three responses. Writing s_i for $s_{i|fn,g}$ and s_n for $s_{n|fg}$, and taking the three in that order,

$$\begin{aligned} \frac{\partial \log X_{fg}}{\partial \log p_{ifn}} &= -s_n s_i \frac{1 - \bar{s}_{f,g}}{\lambda}, \\ \frac{\partial \log s_{n|fg}}{\partial \log p_{ifn}} &= -(\sigma_c - 1)(1 - s_n) s_i, \quad \frac{\partial \log s_{n'|fg}}{\partial \log p_{ifn}} = (\sigma_c - 1) s_n s_i, \\ \frac{\partial \log s_{i|fn,g}}{\partial \log p_{ifn}} &= -(\sigma_b - 1)(1 - s_i), \quad \frac{\partial \log s_{j|fn,g}}{\partial \log p_{ifn}} = (\sigma_b - 1) s_i. \end{aligned} \quad (32)$$

The first row is the store-traffic response. A higher p_{ifn} raises P_{fg} , which lowers the store's choice probability (9) in every neighborhood the store serves. By (9), $\partial \log s_{f\ell g} / \partial \log P_{fg} = -(1 - s_{f\ell g}) / \lambda$, holding E_g , the allocation and rivals' indices fixed, and summing over neighborhoods with (30) gives the row, where $\bar{s}_{f,g} \equiv \sum_{\ell} N_{\ell g} s_{f\ell g}^2 / \sum_{\ell} N_{\ell g} s_{f\ell g}$ is the visitor-weighted mean of $s_{f\ell g}$ over the neighborhoods from which f draws group- g shoppers.

The other two rows are constant-elasticity derivatives of the category share in (14) and of (11), using Shephard's lemma in the form $\partial \log P_{fn,g} / \partial \log p_{ifn} = s_i$ and $\partial \log P_{fg} / \partial \log p_{ifn} = s_n s_i$. Their left column is the barcode's own category and its own place on the shelf, their right column every other category n' and every shelfmate j , and the coefficient in the last row does not depend on which shelfmate.

Barcode j carries the first row always, the second row for whichever category it sits in, and the third row only if it shares a shelf with i . Adding them,

$$\eta_{ji,g} = -s_n s_i \frac{1 - \bar{s}_{f,g}}{\lambda} + \begin{cases} -(\sigma_c - 1)(1 - s_n) s_i - (\sigma_b - 1)(1 - s_i), & j = i, \\ -(\sigma_c - 1)(1 - s_n) s_i + (\sigma_b - 1) s_i, & j \text{ on the same shelf,} \\ (\sigma_c - 1) s_n s_i, & j \text{ in another category.} \end{cases} \quad (33)$$

The traffic term stands outside the cases because a price rise costs the store shoppers and every barcode it sells loses them. The second line differs from the first only in the within-category

response, $+(\sigma_b - 1)s_i$ for a shelfmate against $-(\sigma_b - 1)(1 - s_i)$ for barcode i itself, since what barcode i loses on that margin its shelfmates gain.

Substituting back. Dividing (31) by $X_{if} \equiv \sum_g X_{ifg}$ turns the outer sum into an average over income groups with weights

$$\omega_g = \frac{X_{ifg}}{\sum_{g'} X_{ifg'}} \propto X_{fg} s_{n|fg} s_{i|fn,g}, \quad (34)$$

group g 's share of what the store sells of this barcode. The visitor mix therefore enters the pricing condition directly, and a store whose shoppers are richer prices against those shoppers' elasticities and shares. The three lines of (33) now give Δ_{ifn} , C_{ifn}^{cat} , and the pair C_{ifn}^{store} and $\mathcal{T}_{ifn}$ of (17), in turn.

The first line defines the residual elasticity,

$$\begin{aligned} \Delta_{ifn} \equiv 1 - \sum_g \omega_g \eta_{ii,g} &= 1 + \sum_g \omega_g \left[(\sigma_b - 1)(1 - s_{i|fn,g}) + (\sigma_c - 1)(1 - s_{n|fg}) s_{i|fn,g} \right. \\ &\quad \left. + s_{n|fg} s_{i|fn,g} \frac{1 - \bar{s}_{f,g}}{\lambda} \right], \end{aligned} \quad (35)$$

so the two own terms of (31) collapse to $1 - L_i \Delta_{ifn}$. For a shelfmate, $X_{jfg}/X_{ifg} = s_{j|fn,g}/s_{i|fn,g}$ cancels the s_i that multiplies the second line and the traffic term, leaving one coefficient on one share-weighted markup,

$$C_{ifn}^{\text{cat}} = \sum_g \omega_g (\sigma_b - B_{fn,g}) \sum_{j \neq i} L_j s_{j|fn,g}, \quad B_{fn,g} \equiv 1 + s_{n|fg} \frac{1 - \bar{s}_{f,g}}{\lambda} + (\sigma_c - 1)(1 - s_{n|fg}). \quad (36)$$

The same step one level up handles a barcode in another category. The ratio

$$\frac{X_{jfg}}{X_{ifg}} = \frac{s_{n'|fg} s_{j|fn',g}}{s_{n|fg} s_{i|fn,g}}$$

cancels the $s_n s_i$ multiplying both the third line and the traffic term, and leaves the single coefficient $(\sigma_c - 1) - (1 - \bar{s}_{f,g})/\lambda$ on the markup of everything the store sells outside category n . Its

two halves pull in opposite directions, so they are kept apart rather than netted,

$$C_{ifn}^{\text{store}} = \sum_g \omega_g (\sigma_c - 1) \sum_{n' \neq n} s_{n'|fg} \sum_{j \in \mathcal{I}_{fn'}} L_j s_{j|fn',g}, \quad (37)$$

$$\mathcal{T}_{ifn} = \sum_g \omega_g \frac{1 - \bar{s}_{f,g}}{\lambda} \sum_{n' \neq n} s_{n'|fg} \sum_{j \in \mathcal{I}_{fn'}} L_j s_{j|fn',g}. \quad (38)$$

The condition is therefore $1 - L_i \Delta_{ifn} + C_{ifn}^{\text{cat}} + C_{ifn}^{\text{store}} - \mathcal{T}_{ifn} = 0$, which is (17).

Three features carry into the rest of the paper. Δ_{ifn} is an elasticity of quantity and not of expenditure, since the bracketed terms are the response of $p_{ifn} q_{ifn}$ and quantity is that divided by price, which is where the leading unit comes from and why a barcode with a negligible share of its category, whose elasticity Δ_{ifn} is then σ_b , carries $L = 1/\sigma_b$ rather than $1/(\sigma_b - 1)$. The rate at which the store recaptures on its shelfmates is σ_b net of $B_{fn,g}$, everything a price rise leaks out of the category, namely the unit response, the traffic lost and the substitution toward the store's other categories. And C_{ifn}^{store} and $\mathcal{T}_{ifn}$ multiply the same aggregate with opposite signs, so which dominates is a quantitative question that Section 6.3 settles at the estimates.

The traffic response is not confined to $\mathcal{T}_{ifn}$. The same $(1 - \bar{s}_{f,g})/\lambda$ sits in Δ_{ifn} , as the sales the barcode loses on its own account, and in $B_{fn,g}$, as the sales its shelfmates lose. What $\mathcal{T}_{ifn}$ isolates is the part falling on the store's other categories, which is the part that survives as a separate force once the within-category terms are netted, and it is the loss-leader channel Section 6.3 reports.

C.3 The Assortment Condition

A store in Section 4.2 also chooses which barcodes to carry. Stocking is a choice of set, not of items one at a time, because a barcode added to a shelf takes sales from the shelfmates already on it. This appendix shows that under the constant within-category elasticity the displacement is proportional, every incumbent losing the same fraction of its share, and that the set problem therefore collapses to a condition that can be checked barcode by barcode. A store carries a barcode when its markup beats the average markup of the shelf it displaces, scaled by the share it would capture, and credited with the spending the longer shelf pulls into the category and the traffic it brings in, by enough to cover the stocking cost. That condition is (18). The appendix then shows that the condition characterizes, to first order, a shelf that no single addition or deletion improves (Lemma 1), and gives the version with barcode-specific stocking costs that the estimator and the counterfactuals use (Proposition 1).

Cannibalization is proportional. Fix a store, a category and an income group, and write $\mathcal{S}$ for the set already carried, $d_k = \zeta_k^{\alpha \log y_g} p_{kfn}^{1-\sigma_b}$ for each barcode's weight in (11) and $D = \sum_{k \in \mathcal{S}} d_k$ for

the category denominator, so that $s_{k|fn,g} = d_k/D$. Adding barcode i replaces D by $D + d_i$, so

$$s_{i,g}^{\text{post}} = \frac{d_i}{D + d_i}, \quad s_{k,g}^{\text{post}} = s_{k|fn,g} \frac{D}{D + d_i} = s_{k|fn,g} (1 - s_{i,g}^{\text{post}}) \quad \text{for } k \in \mathcal{S}. \quad (39)$$

Every incumbent loses the same fraction $s_{i,g}^{\text{post}}$ of its share, whatever share it held before. That proportionality is what makes the value of a set separable into barcode-by-barcode contributions.

The value of adding a barcode, inside the category. The shares $s_{k|fn,g}$ of (11) are shares of *spending*, so by (14) the quantity of barcode k sold to group g is $X_{nfg} s_{k|fn,g} / p_{kfn}$ and the profit it earns is $X_{nfg} s_{k|fn,g} L_{kfn}$ with $L_{kfn} = (p_{kfn} - c_{kn}) / p_{kfn}$. Category profit at group- g expenditure $X_{nfg} = X_{fg} s_{n|fg}$ is therefore $\Pi_{fn}(\mathcal{S}) = \sum_g X_{nfg} \sum_{k \in \mathcal{S}} L_{kfn} s_{k|fn,g}$. Substituting (39),

$$\Pi_{fn}(\mathcal{S} \cup \{i\}) - \Pi_{fn}(\mathcal{S}) = \sum_g X_{nfg} s_{i,g}^{\text{post}} \left[L_{ifn} - \sum_{k \in \mathcal{S}} L_{kfn} s_{k|fn,g} \right],$$

since the incumbents' shares are scaled by $1 - s_{i,g}^{\text{post}}$ and so lose exactly $s_{i,g}^{\text{post}}$ of their contribution. That is the first term of the bracket of (18). The same expression evaluated at $\mathcal{S} \setminus \{i\}$ is the value of a barcode the store already carries.

The spending the category draws from the rest of the store. A longer shelf also lowers the category index. From (11), $\log P_{fn,g} = \frac{1}{1-\sigma_b} \log \sum_k d_k$, so adding i moves it by $\frac{1}{1-\sigma_b} \log(1 - s_{i,g}^{\text{post}})^{-1}$, which to first order is $-s_{i,g}^{\text{post}} / (\sigma_b - 1)$. The category and traffic terms below use this first-order change, which understates the exact one when the added barcode would take a large share, as in categories with few barcodes. At a fixed store budget X_{fg} the category shares are CES in the category indices with elasticity σ_c , so to first order $ds_{n|fg} = s_{n|fg} (1 - s_{n|fg}) (\sigma_c - 1) s_{i,g}^{\text{post}} / (\sigma_b - 1)$ and $ds_{n'|fg} = -s_{n'|fg} s_{n|fg} (\sigma_c - 1) s_{i,g}^{\text{post}} / (\sigma_b - 1)$ for every other category n' . Store profit at a fixed budget is $X_{fg} \sum_{n'} s_{n'|fg} \bar{L}_{fn',g}$, so the reallocation changes it by

$$X_{fg} \sum_{n'} \bar{L}_{fn',g} ds_{n'|fg} = X_{nfg} s_{i,g}^{\text{post}} \frac{\sigma_c - 1}{\sigma_b - 1} (\bar{L}_{fn,g} - \bar{L}_{f,g}) = X_{nfg} s_{i,g}^{\text{post}} L_{fn,g}^{\text{cat}},$$

the second term of the bracket. It is zero when the category's margin equals the store's average, and otherwise it is of the same order as the within-category term.

The traffic the shelf brings in. The lower category index passes into the store index through the category share, $d \log P_{fg} = s_{n|fg} d \log P_{fn,g}$, and the store-choice rule (9) turns a lower index

into more shoppers at the rate $(1 - \bar{s}_{f,g})/\lambda$. The extra spending is $X_{fg} s_{n|fg} s_{i,g}^{\text{post}} (1 - \bar{s}_{f,g})/[\lambda(\sigma_b - 1)]$, and it comes from group- g shoppers, who spread it across the store's categories and barcodes in group g 's own shares $s_{n|fg}$ and $s_{k|fn',g}$, so it earns group g 's spending-weighted markup $\bar{L}_{f,g}$. Since $X_{nfg} = X_{fg} s_{n|fg}$, the gain is $X_{nfg} s_{i,g}^{\text{post}} L_{fg}^{\text{traffic}}$ with L_{fg}^{traffic} as in (19). Adding the three terms and collecting the common factor gives (18).

The reallocation and traffic terms are first order in $s_{i,g}^{\text{post}}$, the same order as the cannibalization term they offset, which is why neither can be dropped. They are the parts of the condition that reach outside the category, and the traffic term is what makes a barcode worth carrying at a store whose shoppers are worth keeping. A barcode that would attract no demand at the store has $s_{i,g}^{\text{post}} = 0$, contributes nothing and fails the condition on its own account, so no separate rule is needed to keep the candidate set finite.

The condition is the optimality condition of the set problem. The within-category term of (18) is exact for a discrete addition at fixed category spending; the reallocation and traffic terms are exact to first order in the share the barcode would capture. The estimator conditions on the condition and the counterfactuals move shelves by it, so the result is stated formally.

Lemma 1 (The stocking condition characterizes an unimprovable shelf). *Fix prices, rivals' assortments and the visitor mix, and let the store's profit respond to its shelf in category n through the category's share of the store budget and through the store's traffic. At fixed category spending X_{nfg} , adding $i \notin \mathcal{S}$ changes category profit, net of stocking costs, by exactly $\sum_g X_{nfg} s_{i,g}^{\text{post}} [L_{ifn} - \bar{L}_{fn,g}] - F_n$. With the category share and the traffic free, the change in store profit is $\sum_g X_{nfg} s_{i,g}^{\text{post}} [(L_{ifn} - \bar{L}_{fn,g}) + L_{fn,g}^{\text{cat}} + L_{fg}^{\text{traffic}}] - F_n$ up to terms of higher order in $s_{i,g}^{\text{post}}$. A shelf that no single addition or deletion improves therefore satisfies (18) with the stated sign for every barcode in the category, to first order.*

Proof. By (39) the new shelf's mean markup is the convex combination $(1 - s_{i,g}^{\text{post}})\bar{L}_{fn,g}(\mathcal{S}) + s_{i,g}^{\text{post}} L_{ifn}$, so the change in category profit at fixed X_{nfg} is an identity; deletion is the same statement at $\mathcal{S} \setminus \{i\}$. The reallocation and traffic terms are the first-order moves in the category share and the store's index derived above. $\square$

A shelf that no single addition or deletion improves need not maximize the store's profit. Several income groups make the objective a mixture of logit assortment problems, for which optimization is computationally hard in general (Rusmevichientong et al., 2014); a positive stocking cost and the store-traffic response add further room for several unimprovable shelves; and across stores one store's drop raises a rival's residual demand for the same barcode. The shelf the model works with is therefore the unimprovable shelf of the lemma, and Appendix C.4 explains how the counterfactuals select among such shelves.

Stocking costs with a barcode-specific draw. A cost common to every barcode in a cell would require each shelf to follow a strict ranking of the index, every carried barcode above every

uncarried one, and observed shelves do not. The estimated model therefore lets the stocking cost be common within a store–category cell up to a barcode- and store-specific draw,

$$i \in \mathcal{I}_{fn} \iff x_{ifn} - v_{fn} \geq \varepsilon_{ifn}, \quad x_{ifn} = \log \Delta \pi_f(i), \quad \varepsilon_{ifn} \sim \text{Logistic}(0, \varsigma_{ifn}), \quad (40)$$

so that $\Pr(i \in \mathcal{I}_{fn}) = G((x_{ifn} - v_{fn})/\varsigma_{ifn})$ with G the logistic distribution function. The draws are independent across barcodes, stores and categories, and independent of the indices x_{ifn} . Since $x_{ifn} = \log \Delta \pi_f(i)$, the stocking cost of barcode i is $F_{ifn} = \exp(v_{fn} + \varepsilon_{ifn})$. The scale is one for a barcode the store carries, whose index is built from its observed price. A barcode it does not carry has no price there, so its index is built from the price the pricing conditions assign it and inherits the dispersion of the pricing residual, and its scale is $\varsigma_{ifn} = (1 + (3/\pi^2)(\sigma_b - 1)^2 s_e^2)^{1/2}$, with s_e the within-cell standard deviation of the residual and $3/\pi^2$ converting a variance to a logistic scale. The threshold v_{fn} carries the part of the stocking cost common to the cell, and it is set so that the cell's expected number of barcodes carried equals the observed one. The level of F_n is therefore solved cell by cell from the size of the shelf, and the composition of the shelf, given that threshold, identifies α (Section 5). The index x_{ifn} is built from the expenditure shares across income groups, so v_{fn} also absorbs the level of the cell's expenditure. The draw ε_{ifn} carries what is specific to the barcode at that store.

Proposition 1 (Every observed assortment is rationalized, and bounds its own draws). *Under (40), for any observed assortment $\mathcal{I}_{fn}$ and any finite $\{x_{ifn}\}$ and v_{fn} , the set of cost draws consistent with what is observed is*

$$\varepsilon_{ifn} \leq x_{ifn} - v_{fn} \text{ for } i \in \mathcal{I}_{fn}, \quad \varepsilon_{ifn} > x_{ifn} - v_{fn} \text{ for } i \notin \mathcal{I}_{fn}, \quad (41)$$

which has positive probability under the logistic. The conditional law of ε_{ifn} given the observed carriage is the logistic truncated to the relevant side, and drawing from it rationalizes the observed shelf cell by cell, with no cell excluded.

Proof. Fix the observed shelf $\mathcal{I}_{fn}$. Each index is evaluated at that shelf. For a carried barcode, x_{ifn} is the gain from keeping i given the rest of the shelf, and for an uncarried one it is the gain from adding i to the shelf. Once the shelf is fixed, $\{x_{ifn}\}$ and v_{fn} are therefore numbers that do not depend on the draws. Each event in (41) restricts only its own draw, and the draws are independent, so the events are independent. Each is a half-line of the logistic support, so each has positive probability and so does their intersection over a finite cell. Conditioning a random variable on a positive-probability event gives the truncated law. $\square$

The probability of (41), $\prod_{i \in \mathcal{I}_{fn}} G((x_{ifn} - v_{fn})/\varsigma_{ifn}) \prod_{i \notin \mathcal{I}_{fn}} [1 - G((x_{ifn} - v_{fn})/\varsigma_{ifn})]$, is therefore the probability that the observed shelf is unimprovable in the sense of Lemma 1: every

carried barcode passes its deletion test and every uncarried one fails its addition test. The shelf-composition moment is the score of this probability with respect to α . When several shelves are unimprovable at the same draws, the probability of observing a given one also depends on which of them the store selects, and the moment omits that factor. It is the score of the likelihood of the observed shelf when the unimprovable shelf is unique or when the selection among unimprovable shelves does not depend on α , and I assume the latter. This is the assumption under which discrete games with multiple equilibria are estimated by maximum likelihood (Tamer, 2003; Bajari et al., 2010). Without it, the probability of being unimprovable bounds the probability of being observed from above, and α is identified only as a set (Ciliberto and Tamer, 2009).

The proposition takes a finite threshold as given. A cell whose threshold has no solution at a candidate parameter vector contributes nothing to the shelf-composition moment at that point, and the counterfactuals hold such a cell at its observed shelf (Appendix E).

The estimator uses the threshold to condition on each cell's count, and the counterfactual shelf response uses the truncated law to turn a shock into the probability that each barcode crosses its threshold (Appendix E).

C.4 Existence and Uniqueness of Equilibrium

This appendix shows that the equilibrium of Section 4.4 exists and identifies which parts of it are unique. The paper needs three things from it. The estimator evaluates the pricing conditions at observed prices and needs them to return a single vector of margins. The counterfactuals solve for a new equilibrium and need one to exist at every configuration of stores and shelves the solver visits. And since the counterfactuals report one equilibrium, the paper has to say where multiplicity can arise and which equilibrium it reports. Table 24 lists the results in the order they are proved, with what each delivers and where it is used. The results on a store's choice of shelf concern the store's own problem and are in Appendix C.3.

The equilibrium as a fixed point. Fix primitives $\{A_{\ell g}\}, \{N_g\}$, the operating stores and their assortments. Two objects remain, the population allocation $N = \{N_{\ell g}\}$ and the vector of margins $L = \{L_{ifn}\}$, and every other object in the equilibrium definition is a function of the pair. Margins give prices through $p_{ifn} = c_{in}/(1 - L_{ifn})$, prices give the category and store indices of (10) and the store-choice shares of (9), and those with the allocation give the grocery price index $\mathcal{P}_{\ell g}$ of (8) and the visitor mix each store faces. The allocation on its own gives rents $r_\ell(N)$ from floorspace-market clearing, and the two together give the deterministic indirect utility $\bar{u}_{\ell g} = A_{\ell g} y_g \mathcal{P}_{\ell g}^{-\beta_g} r_\ell^{-\rho_g}$. Fréchet sorting maps utilities back into populations,

$$N_{\ell g} = T_{\ell g}(N, L) \equiv N_g \frac{\bar{u}_{\ell g}(N, L)^{\kappa_\epsilon}}{\sum_{j \in \mathcal{L}} \bar{u}_{jg}(N, L)^{\kappa_\epsilon}}, \quad (42)$$

Table 24: Existence and uniqueness results and their use

Result	What it delivers	Where it is used
Lemma 2	A unique and bounded rent at every allocation	Every result below; the counterfactual rents of (50)
Lemma 3	The row sums of the pricing system in closed form	Proposition 2; the markup-level moment that identifies λ (Section 5)
Proposition 2	Unique margins at a given demand block	The estimator, which evaluates the pricing conditions at observed prices and shares
Proposition 3	An equilibrium at any fixed set of stores and shelves, with prices, rents and sorting solved jointly	The counterfactual solver, at every configuration it visits
Proposition 4	An equilibrium in mixed strategies with the stores and shelves also chosen, condition (iii) holding in expectation over the mixing	Nothing directly, since the counterfactuals check condition (iii) at the configuration they compute
Proposition 5	Unique rents and sorting when the retail side is held fixed, for every parameter value	The population step of the counterfactual solver; the frozen-retail comparison of Table 13; locating multiplicity in the retail response

and the pricing conditions (17) map margins back into margins. An equilibrium at a given configuration is a joint fixed point of the two maps on $\mathcal{X} \times [0, L^{\max}]^{\mathcal{I}}$, where $\mathcal{X} \equiv \{N \geq 0 : \sum_{\ell} N_{\ell g} = N_g \forall g\}$ is a product of scaled simplices and $[0, L^{\max}]^{\mathcal{I}}$ is a cube of margins bounded away from one. The two blocks are solved jointly and not in sequence, since solving them in sequence would require the pricing block to return a single value at every allocation, which it is not shown to do.

Rents.

Lemma 2 (Rents). *Let $\underline{r} > 0$. For each location ℓ and any $N \in \mathcal{X}$, (20) has a unique solution $r_{\ell}(N) \in [\underline{r}, \bar{r}]$ with $\bar{r}$ independent of N , and r_{ℓ} is continuous and non-decreasing in the local population of every group.*

Proof. Write $D_{\ell} = \sum_g \rho_g y_g N_{\ell g}$ and $E_{\ell}(x; D_{\ell}) = \bar{H}_{\ell}(x/r_{\ell}^0)^{\eta} - D_{\ell}/x - H_{\ell}^{\text{firm}}$ for the excess supply of floorspace at rent x . For $\eta > 0$, E_{ℓ} is strictly increasing in x and non-increasing in D_{ℓ} , and commercial demand does not depend on x . As $x \rightarrow \infty$, $E_{\ell} \rightarrow \infty$ since $\bar{H}_{\ell} > 0$, and as $x \downarrow 0$, $E_{\ell} \rightarrow -\infty$ when $D_{\ell} > 0$ and $E_{\ell} \rightarrow -H_{\ell}^{\text{firm}} \leq 0$ otherwise, so E_{ℓ} has a single root $\tilde{r}_{\ell}$, which is zero exactly when the location has neither residents nor a store. The root is bounded above by its value when the whole population of the city lives in ℓ , and $\bar{r}$ is the largest such value across locations, or $\underline{r}$ if that is larger. It depends on the operating stores through H_{ℓ}^{firm} . At a positive root, $\partial E_{\ell}/\partial x = \eta \bar{H}_{\ell} x^{\eta-1}/(r_{\ell}^0)^{\eta} + D_{\ell}/x^2 > 0$, so the implicit function theorem gives continuity and monotonicity of $\tilde{r}_{\ell}$ in N wherever $\tilde{r}_{\ell} > 0$. The zero root arises only where $\tilde{r}_{\ell} < \underline{r}$, where

the floor binds and $r_\ell = \underline{r}$ is constant, and the maximum of two continuous non-decreasing functions is continuous and non-decreasing. Hence r_ℓ has both properties and lies in $[\underline{r}, \bar{r}]$ at every allocation, and since $\underline{r} > 0$, the factor $r_\ell^{-\rho_g}$ in utility is bounded above and below by positive constants. $\square$

Prices at a given demand block. The pricing conditions are one equation per barcode, coupled across a store's shelf. The estimator evaluates them at observed prices and shares, and at a fixed demand block they are affine in the margins. The margins they return are those the pricing conditions imply at the observed demand block; they differ from the observed margins $(p_{ifn} - c_{in})/p_{ifn}$ by a residual for each barcode-store pair, which the counterfactuals hold at its observed value. Evaluate the demand block at given prices and treat the margins as the unknowns, with costs following as $c_{in} = p_{ifn}(1 - L_i)$. Then Δ_{ifn} , ω_g , $s_{i|fn,g}$, $s_{n|fg}$, $B_{fn,g}$ and the store-traffic term $\Omega_{fg} \equiv (1 - \bar{s}_{f,g})/\lambda$ that appears in (32) do not depend on L , and inspection of (35)–(38) shows that the margins enter the three corrections only through the category sums $\sum_{j \in \mathcal{I}_{fn'}} L_j s_{j|fn',g}$, for the barcode's own category and for the store's other categories, linearly. The pricing conditions (17) are therefore the affine system

$$L_i = L_i^0 + \sum_{j \neq i} C_{ij} L_j, \quad L_i^0 \equiv \frac{1}{\Delta_{ifn}}, \quad C_{ij} \equiv \frac{1}{\Delta_{ifn}} \frac{\partial (C_{ifn}^{\text{cat}} + C_{ifn}^{\text{store}} - \mathcal{T}_{ifn})}{\partial L_j}. \quad (43)$$

The intercept L_i^0 is the markup barcode i would carry if the store ignored what its shelfmates recapture, and C_{ij} is the net weight on barcode j 's margin, the sales recaptured less the traffic lost, which is non-negative under (45). C_{ij} is zero when i and j sit in different stores, since a barcode's corrections involve only its own store's assortment.

Lemma 3 (Row sums). *For every barcode, the coefficients in (43) satisfy the identity*

$$\sum_{j \neq i} C_{ij} = \frac{\Delta_{ifn} - 1 - \bar{\Omega}_{ifn}}{\Delta_{ifn}}, \quad \bar{\Omega}_{ifn} \equiv \sum_g \omega_g \Omega_{fg}. \quad (44)$$

Proof. Only barcodes of store f enter the sum, since $C_{ij} = 0$ across stores. Sum the same-category coefficients using $\sum_{j \neq i} s_{j|fn,g} = 1 - s_{i|fn,g}$ and the cross-category ones using $\sum_{n' \neq n} s_{n'|fg} = 1 - s_{n|fg}$ and $\sum_{j \in \mathcal{I}_{fn'}} s_{j|fn',g} = 1$, which hold since shares within a category and across a store's categories sum to one. Writing s_i and s_n for the two shares, the group- g term is

$$\underbrace{(\sigma_b - B_{fn,g})(1 - s_i)}_{\text{same category}} + \underbrace{[(\sigma_c - 1) - \Omega_{fg}](1 - s_n)}_{\text{other categories}}.$$

Substitute $B_{fn,g} = 1 + s_n \Omega_{fg} + (\sigma_c - 1)(1 - s_n)$ and subtract the group- g integrand of $\Delta_{ifn} - 1$ in (35), which is $(\sigma_b - 1)(1 - s_i) + (\sigma_c - 1)(1 - s_n)s_i + s_n s_i \Omega_{fg}$. Every term in $\sigma_b - 1$ and in $\sigma_c - 1$ cancels, and the terms in Ω_{fg} leave

$$-\Omega_{fg} [s_n(1 - s_i) + (1 - s_n) + s_n s_i] = -\Omega_{fg},$$

the last entry in the bracket coming from the traffic piece of Δ_{ifn} itself. The group- g term is therefore that integrand less Ω_{fg} . Averaging over g with the weights ω_g and dividing by Δ_{ifn} gives (44). $\square$

Proposition 2 (Margins at a given demand block). *Hold the demand block $\{s_{i|fn,g}, s_{n|fg}, \Delta_{ifn}, B_{fn,g}, \Omega_{fg}, \omega_g\}$ at the values implied by a given price vector, with $\sigma_b > 1$ and $\sigma_c > 1$, and suppose that at every barcode and income group*

$$\sigma_b \geq B_{fn,g}, \quad \sigma_c - 1 \geq \Omega_{fg} \geq 0. \quad (45)$$

Then $C \geq 0$ elementwise,

$$\|C\|_\infty = \max_i \left[1 - \frac{1 + \bar{\Omega}_{ifn}}{\Delta_{ifn}} \right] < 1,$$

and the pricing conditions have the unique solution $L^* = (I - C)^{-1}L^0$, to which iterating (43) with the demand block held fixed converges from any starting vector, geometrically at rate $\|C\|_\infty$.

Proof. Under (45) the same-category and cross-category coefficients

$$\frac{1}{\Delta_{ifn}} \sum_g \omega_g s_j [\sigma_b - B_{fn,g}], \quad \frac{1}{\Delta_{ifn}} \sum_g \omega_g [(\sigma_c - 1) - \Omega_{fg}] s_{n'|fg} s_{j|fn',g}$$

are both non-negative, so $C \geq 0$ and the absolute row sums equal the row sums of Lemma 3. Non-negativity forces $\Delta_{ifn} \geq 1 + \bar{\Omega}_{ifn}$, so each row sum is at least zero, and it is below one since every bracketed term of (35) is non-negative when $\sigma_b > 1$, $\sigma_c > 1$ and $\Omega_{fg} \geq 0$, which makes $(1 + \bar{\Omega}_{ifn})/\Delta_{ifn} > 0$. Hence $\rho(C) \leq \|C\|_\infty < 1$, and invertibility of $I - C$, the Neumann representation $L^* = \sum_{k \geq 0} C^k L^0$ and geometric convergence of the iteration follow. Since $C \geq 0$ and $L^0 > 0$, the representation gives $L_i^* \geq 1/\Delta_{ifn} > 0$. Iterating from zero, the bound in the proof of Proposition 3 keeps every iterate below $1/(1 + \min_j \bar{\Omega}_j)$, so $L_i^* < 1$ whenever every $\bar{\Omega}_j > 0$, and the costs $c_{in} = p_{ifn}(1 - L_i^*)$ are positive. $\square$

Remark 1 (Sign conditions at the estimates). The sign conditions hold at every price vector at the estimates, not only at observed prices. $B_{fn,g} = s_{n|fg}(1 + \Omega_{fg}) + (1 - s_{n|fg})\sigma_c$ is a convex combination of $1 + \Omega_{fg}$ and σ_c , and $0 \leq \Omega_{fg} \leq 1/\lambda$ since $\bar{s}_{f,g} \in [0, 1]$, so it suffices that $\sigma_b \geq \max(\sigma_c, 1 + 1/\lambda)$ and $\sigma_c - 1 \geq 1/\lambda$. With $\sigma_b = 5.55$, $\max(\sigma_c, 1 + 1/\lambda) = 4.70$, $\sigma_c - 1 = 3.70$ and

$1/\lambda = 3.47$, both inequalities hold. They hold at the point estimates and not across the sampling uncertainty of the first stage. σ_b is estimated with a standard error of 2.418, so values below 4.70, at which the first inequality fails, lie within its confidence interval, and the same is true of the maintained ordering $\sigma_b > \sigma_c$. The guarantees that require (45) are therefore statements at the estimates.

The row-sum identity is what lets Section 5 read λ off the level of markups. At a given demand block, if $\bar{\Omega}_{ifn}$ takes a common value $\bar{\Omega}_f$ across every barcode the store sells, $L_i = 1/(1 + \bar{\Omega}_f)$ for every barcode solves (43) by (44) whatever the values of σ_b and σ_c , since substituting $L_j = L$ gives $L = 1/\Delta_{ifn} + L(\Delta_{ifn} - 1 - \bar{\Omega}_f)/\Delta_{ifn}$, that is $L(1 + \bar{\Omega}_f) = 1$. By Proposition 2 it is the only solution whenever the elasticities satisfy (45), as they do at the estimates. The within-category and across-category elasticities therefore set how margins spread across a shelf, through the variation in $\bar{\Omega}_{ifn}$ that arises when the barcode-specific weights ω_g of (34) load differently on groups whose traffic terms Ω_{fg} differ, and the level is set by the store-choice margin. The unit in $1 + \bar{\Omega}_{ifn}$ comes from Δ_{ifn} being an elasticity of quantity and not of expenditure; the store-traffic term is the loss of traffic on everything else the store sells, which no shelfmate recovers.

Existence at a given configuration. Once prices are free to move, the demand block moves with them and (43) is no longer affine. Write $\vartheta(p, N)$ for the demand block, which the allocation reaches through each store's visitor mix and the spending it draws, and $\Phi(L, N) \equiv \mathcal{B}(L; \vartheta(p(L), N))$ for the right-hand side of (43) at the demand block the margins themselves imply. Existence requires the demand block to be continuous in prices and Φ to keep margins away from one.

Lemma 4 (Regularity of demand). *Suppose $\sigma_b > 1$ and $\sigma_c > 1$. Then the category indices $P_{fn,g}$, the shares $s_{i|fn,g}$, $s_{n|fg}$ and $s_{f\ell g}$, and hence each store's variable profit, are continuously differentiable in prices on $\{p : p \gg 0\}$. On allocations at which every tract holds households of every group, the group weights ω_g , the visitor-weighted shares $\bar{s}_{f,g}$, and with them Ω_{fg} , Δ_{ifn} and $B_{fn,g}$, are continuous in prices and the allocation jointly.*

Proof. Under the constant within-category elasticity $P_{fn,g}$ is the closed-form aggregator in (10), a positively weighted CES (power-sum) aggregator of positive prices with exponent $1 - \sigma_b \neq 0$, so it is C^1 in p on $p \gg 0$. Composition with the across-category aggregator in the same equation, itself C^1 since $1 - \sigma_c \neq 0$, and with the logit map (9), which is C^1 in $P_{fg} > 0$, gives the indices. The shares are ratios and powers of these positive C^1 functions, and variable profit, $\sum_{g,n,i} L_i X_{fg} s_{n|fg} s_{i|fn,g}$ with $L_i = 1 - c_{in}/p_{ifn}$, is a finite sum of products of them.

The weights ω_g and the shares $\bar{s}_{f,g}$ are ratios whose numerators and denominators are sums of products of the shares above with the allocation, so they are continuous in (p, N) wherever their denominators are positive. Each denominator is a store's group- g visitors, $\sum_{\ell} N_{\ell g} s_{f\ell g}$, or the spending they bring. The logit share $s_{f\ell g}$ is positive at every tract the store reaches, every

store reaches at least one tract, and every tract holds group- g households by assumption, so each denominator is positive. Ω_{fg} , Δ_{ifn} and $B_{fn,g}$ are continuous functions of these ratios and of the shares. $\square$

Proposition 3 (Existence at a given configuration). *Fix the operating stores and their assortments. Suppose $\sigma_b > 1$, $\sigma_c > 1$ and (45) hold at every price vector (Remark 1 verifies this at the estimates), and that no store draws more than a share $s^{\max} < 1$ of any group's shoppers, $\bar{s}_{f,g} \leq s^{\max}$, so that $\Omega_{fg} \geq \underline{\Omega} \equiv (1 - s^{\max})/\lambda > 0$. Then conditions (i), (ii), (iv) and (v) of the equilibrium definition hold simultaneously at some (N^*, L^*) with $L^* \in [0, L^{\max}]^{\mathcal{I}}$ and $L^{\max} \equiv 1/(1 + \underline{\Omega}) < 1$.*

Proof. The domain $\mathcal{X} \times [0, L^{\max}]^{\mathcal{I}}$ is non-empty, compact and convex, and the candidate map is $(N, L) \mapsto (T(N, L), \Phi(L, N))$.

Take the pricing coordinate first. Call a demand block admissible if it satisfies (45) and $\Omega_{fg} \geq \underline{\Omega}$ for every store and group; under the hypotheses, every demand block the domain reaches is admissible. At any admissible demand block the map $\mathcal{B}(\cdot; \vartheta)$ is affine with intercept $1/\Delta_{ifn} > 0$, coefficients $C_{ij} \geq 0$ and, by Lemma 3, row sums $1 - (1 + \bar{\Omega}_{ifn})/\Delta_{ifn}$. For $L \in [0, L^{\max}]^{\mathcal{I}}$,

$$\mathcal{B}_i(L) \leq \frac{1}{\Delta_{ifn}} + \left[1 - \frac{1 + \bar{\Omega}_{ifn}}{\Delta_{ifn}}\right] L^{\max} \leq L^{\max} \iff 1 \leq L^{\max}(1 + \bar{\Omega}_{ifn}),$$

which holds since $\bar{\Omega}_{ifn} \geq \underline{\Omega}$, and $\mathcal{B}_i(L) \geq 1/\Delta_{ifn} > 0$. So $\mathcal{B}(\cdot; \vartheta)$ maps the cube into itself at every admissible ϑ . Prices $p_i = c_{in}/(1 - L_i)$ are continuous on the cube since $L^{\max} < 1$, and ϑ is continuous in prices and the allocation by Lemma 4 on allocations at which every tract holds households of every group, so Φ maps the product into the cube continuously on the domain fixed below.

Take the sorting coordinate. By Lemma 4 the store indices and store-choice shares, and with them $\mathcal{P}_{\ell g}$, are continuous in L on the cube. By Lemma 2 rents lie in $[\underline{r}, \bar{r}]$ and are continuous in N . Hence $\bar{u}_{\ell g}$ is continuous in (N, L) and bounded above and below by positive constants, the Fréchet shares are continuous, and T lands in $\mathcal{X}$ since they sum to one. The same bounds give every share a floor $\epsilon > 0$, so T maps into $\mathcal{X}_{\epsilon} \equiv \{N \in \mathcal{X} : N_{\ell g} \geq \epsilon N_g\}$, a compact convex set on which every tract holds households of every group, so by Lemma 4 Φ is defined and continuous on it. The map therefore sends $\mathcal{X}_{\epsilon} \times [0, L^{\max}]^{\mathcal{I}}$ into itself.

Brouwer's theorem on that set yields a fixed point (N^*, L^*) . Its first coordinate is condition (iv) and its second is condition (ii), the two holding at the same point. Lemma 2 at N^* gives condition (v), and condition (i) holds by construction of $\bar{u}_{\ell g}$ and the nested demand system. $\square$

The bound $L^{\max}$ comes from the store-choice margin. $\underline{\Omega} > 0$ says a price rise always costs a store some probability of being visited, and that possibility holds margins away from one; as $\underline{\Omega} \rightarrow 0$, $L^{\max} \rightarrow 1$ and the argument fails. The share bound holds for every store that shares at least one of the tracts it serves with another store. A store that is the only one within reach of

every tract it serves has a share of one and no traffic to lose, so it falls outside the bound, and the existence results do not cover its prices. No such store is within reach of the counterfactual site.

Existence with stores and shelves chosen. Which stores operate and what they carry, condition (iii), is a choice over finite sets. What a store earns from a choice is the profit left once prices, rents and sorting have settled, and Proposition 3 guarantees that this continuation exists at every configuration without guaranteeing that it is a single point. A selection rule assigns one continuation equilibrium to each configuration.

Proposition 4 (Existence at the configuration stage). *Let each store's action set consist of the shelves it may carry at its site and the option of not operating, both finite, and maintain the conditions of Proposition 3 at every configuration. For any selection rule, the configuration game in which each store earns the profit that rule assigns it has an equilibrium in mixed strategies.*

Proof. By Proposition 3 the set of continuation equilibria is non-empty at every configuration, so each store's profit under the rule is a real number at every profile of actions. The game is finite, so Nash (1950) applies; continuity in the configuration is not required, the mixed extension of a finite game being multilinear in the mixing probabilities. Non-negative expected profit at operating stores is implied, since staying out is available and pays zero. $\square$

The proposition is weaker than the equilibrium definition. Condition (iii) asks for one configuration at which every shelf is unimprovable and every operating store earns $\Pi_f \geq 0$; in a mixed equilibrium these hold in expectation over the rivals' mixing, and a realized configuration can leave a store with a loss. The counterfactuals do not rely on the proposition for condition (iii). They compute a single configuration and check at it that each shelf solves (18) at its threshold, that no rationalized store gains by exiting, and that no candidate entrant gains by entering (Appendix E).

Mixed strategies are needed since the assortment problem lacks the structure that delivers a pure equilibrium. Adding a barcode rescales every shelfmate's share by a common factor (Appendix C.3), so the value of a shelf is not the sum of the values of its barcodes, and no pure-strategy existence result is established for the configuration game.

Uniqueness and the equilibrium reported. Uniqueness holds for the household side of the model whenever the retail side is held fixed, and this needs no restriction on the parameters. The proposition allows part of the population to be held at given locations, which is the placement of Section 8.3.

Proposition 5 (Sorting and rents are unique given the retail side). *Fix the operating stores, their assortments and their prices, and let a part $\bar{N}_{\ell g} \geq 0$ of each group, with $\sum_{\ell} \bar{N}_{\ell g} \leq N_g$, be held at given locations, the remaining $M_g \equiv N_g - \sum_{\ell} \bar{N}_{\ell g}$ sorting by (42) with N_g replaced by M_g and added to $\bar{N}_{\ell g}$.*

For any $\kappa_\epsilon > 0$, $\rho_g \geq 0$ and $\eta > 0$ there is exactly one rent vector and one allocation at which conditions (iv) and (v) hold together.

Proof. For existence, rents are continuous in the allocation by Lemma 2, so the map that sends an allocation to $\bar{N}$ plus the sorting of the M_g households at the rents it implies is continuous. It sends the compact convex set $\{N \geq \bar{N} : \sum_\ell (N_{\ell g} - \bar{N}_{\ell g}) = M_g \forall g\}$ into itself, and Brouwer's theorem gives a fixed point. For uniqueness, note that with prices and assortments fixed the index P_{fg} depends only on the store's own prices and assortment by (10), and the store-choice shares (9) and the index $\mathcal{P}_{\ell g}$ of (8) depend on the allocation only through $\{P_{fg}\}$ and the distances. Commercial floorspace H_ℓ^{firm} is fixed with the stores. So $\bar{u}_{\ell g} = a_{\ell g} r_\ell^{-\rho_g}$ with $a_{\ell g} > 0$ independent of the allocation, positive since amenities are positive and $\mathcal{P}_{\ell g}$ is finite, including at a tract that reaches no store, and the sorting households of group g in ℓ number $M_g a_{\ell g}^{\kappa_\epsilon} r_\ell^{-\kappa_\epsilon \rho_g} / \sum_j a_{jg}^{\kappa_\epsilon} r_j^{-\kappa_\epsilon \rho_g}$.

Suppose two rent vectors $r \neq r'$ both clear with their allocations N and N' . Relabelling the two if necessary, $m \equiv \max_\ell r'_\ell / r_\ell > 1$; let ℓ^* attain it. Since $r'_j \leq m r_j$ for every j and $\kappa_\epsilon \rho_g \geq 0$, $\sum_j a_{jg}^{\kappa_\epsilon} r_j'^{-\kappa_\epsilon \rho_g} \geq m^{-\kappa_\epsilon \rho_g} \sum_j a_{jg}^{\kappa_\epsilon} r_j^{-\kappa_\epsilon \rho_g}$, and so

$$N'_{\ell^*g} - \bar{N}_{\ell^*g} = \frac{M_g a_{\ell^*g}^{\kappa_\epsilon} (m r_{\ell^*})^{-\kappa_\epsilon \rho_g}}{\sum_j a_{jg}^{\kappa_\epsilon} r_j'^{-\kappa_\epsilon \rho_g}} \leq \frac{M_g a_{\ell^*g}^{\kappa_\epsilon} r_{\ell^*}^{-\kappa_\epsilon \rho_g}}{\sum_j a_{jg}^{\kappa_\epsilon} r_j^{-\kappa_\epsilon \rho_g}} = N_{\ell^*g} - \bar{N}_{\ell^*g}$$

for every g . Residential housing expenditure at ℓ^* is therefore weakly lower under the primed allocation, $D'_{\ell^*} \leq D_{\ell^*}$. Since $r'_{\ell^*} > r_{\ell^*} \geq \underline{r}$, the floor is slack at ℓ^* under the primed allocation and r'_{ℓ^*} is the root of $E_{\ell^*}(\cdot; D'_{\ell^*})$ in the proof of Lemma 2. As E_{ℓ^*} is non-increasing in expenditure, $E_{\ell^*}(r'_{\ell^*}; D_{\ell^*}) \leq E_{\ell^*}(r'_{\ell^*}; D'_{\ell^*}) = 0 = E_{\ell^*}(\tilde{r}_{\ell^*}; D_{\ell^*})$, and as it is strictly increasing in rent, $r'_{\ell^*} \leq \tilde{r}_{\ell^*} \leq r_{\ell^*}$, a contradiction. Hence $r' = r$, and the allocation, a function of rents alone, is unique. $\square$

The argument needs no bound on the sorting elasticity since, with the retail side fixed, the only channel from where households live back to what a location offers is rent, and rent falls on the arrivals. At the location where rents rise most in proportion, every group's share weakly falls, so housing demand cannot support the higher rent. In the counterfactuals the proposition covers the population and rent step of each pass of the solver, which solves sorting at given location values, and the frozen-retail solution of Table 13, which holds every store at its observed prices and shelf. The solver clears the floorspace market with (50), whose excess supply, divided by x_ℓ , is also strictly increasing in rent and non-increasing in housing demand, so the same argument applies.

Once the retail side responds, uniqueness is not available from the primitives, and each margin has its own source. Prices move the shares they are computed from, so the pricing conditions lose the affine structure of Proposition 2. A store with positive stocking costs can hold several shelves that no single addition or deletion improves (Appendix C.3). And discrete entry with strategic substitutes pins down how many stores operate without pinning down which when

firms are alike (Bresnahan and Reiss, 1991; Berry, 1992), and with heterogeneous firms it need not pin down either (Ciliberto and Tamer, 2009). By Proposition 5, these retail responses are the only route by which a neighborhood that gains residents can become more attractive, so any multiplicity in the model runs through them.

The counterfactuals report the equilibrium reached by continuation from the observed one. Each block is re-solved in turn starting from the observed allocation, prices, shelves and network (Appendix E), which is one selection rule in the sense of Proposition 4 and the standard stance in quantitative spatial models (Ahlfeldt et al., 2015; Redding and Rossi-Hansberg, 2017). The reported changes are therefore changes along the branch of equilibria that contains the observed city, and not claims about a unique equilibrium of the retail sector.

The estimator needs none of the retail-side uniqueness. It conditions on the observed network, shelves and prices, where Proposition 2 returns one vector of margins, and on the observed allocation, which the amenity inversion reproduces exactly for every group. The observed prices are an exact equilibrium of the full model once each barcode–store pair keeps its pricing residual, which the counterfactuals hold at its observed value (Appendix E). At the data the model with prices held fixed and the full model therefore coincide, and they differ only away from the observed equilibrium, which is to say in the counterfactuals.

D Estimation Appendix

D.1 Inputs fixed before the objective

Income groups are the four household income quartiles, the finest partition at which the consumption survey and the census both measure every group-specific object. The 2022 census supplies everything at the census tract: population, the mean and dispersion of tract income, and the centroids that define distances to stores. The quartile cutoffs and group mean incomes are those of a household-weighted mixture of lognormal distributions, one per tract, each matched to the tract’s mean and variance of household-head income, so groups and tracts are measured on one distribution. A tract with no recorded variance takes the median dispersion across tracts, and the log standard deviation is capped at 1.5. The 2022 tract file does not report the income distribution itself, so each tract’s group composition is predicted from its mean and dispersion by a map estimated on the national 2022 sample and reconciled with its recorded mean income (Appendix A.6). The 2010 census also supplies the one object the 2022 questionnaire did not collect, mean rents, which are released only for áreas de ponderação. A tract’s rent is its area’s mean rent scaled by the ratio of the tract’s mean income to the area’s, raised to the power 0.3. Relative rents across areas are therefore those of 2010. They are the baseline the amenity inversion fits and from which counterfactual rents move by (50), and they value the floorspace in each store’s rent bill. The travel cost comes from Miyauchi et al. (2025), who estimate the semi-elasticity of the cost of a non-commuting trip with respect to travel time at 0.69 per hour, in the exponential

form the model applies to distance. At 4.6 km/h, a walking pace, it is $\tau = 0.15$ per kilometre. The housing-supply elasticity is $\eta = 0.300$. For the Maceió metropolitan area, Guedes et al. (2023) estimate an inverse housing-supply elasticity of 1.24, a supply elasticity of about 0.8. That estimate describes a metropolitan area, which can expand at its edge, and so bounds from above the response of a single tract within it, which cannot. Across US neighborhoods, Baum-Snow and Han (2024) estimate average tract-level supply elasticities of 0.3 for housing units and 0.5 for floorspace, against a population-weighted metropolitan elasticity of 1.75 (Saiz, 2010). The value used is the tract-level elasticity for housing units. The budget shares come from the survey. Store choice sets are truncated at 3 km from the tract centroid, at which distance a store would need a basket 36% cheaper than one at the household's door to be chosen as often. Each store sits at the centroid of its census tract. A store whose tract is excluded for having fewer than 100 households, which is common in commercial districts, sits at the centroid of the included tract nearest its own coordinates.

Section 2 estimates barcode and category fixed effects from a hedonic decomposition. Only relative levels are identified, so the barcode effects are normalized within category and exponentiated into the quality of (11) and (16),

$$\tilde{\delta}_i = \delta_i - \frac{1}{|\mathcal{I}_n|} \sum_{j \in \mathcal{I}_n} \delta_j, \quad \tilde{\zeta}_i = \exp(\tilde{\delta}_i), \quad (46)$$

and the category effects into cost shifters the same way. Section 2 centres δ_i on the observation-weighted mean of its category. The model re-centres it on the unweighted barcode mean, so that the geometric mean of $\tilde{\zeta}_i$ within each category is one, and moves the difference into the category cost shifter, which leaves marginal cost reproducing the same prices. Quality is measured only within a category, so this choice sets its level. The average barcode of every category carries the same weight for every income group, and groups differ in how they split spending across categories only through the quality of what is on each shelf. The level of marginal cost is not identified by the hedonic, which recovers costs only relative to one another. The markup-level moment could fix it only given the store-taste dispersion, so it is set instead by the retail margin of Section 2, a single constant that scales every cost and leaves relative costs unchanged, and the moment then identifies λ at that level.

D.2 The two first stages

Both elasticities are estimated on the food diary of the household consumption survey, which records a standardized weight alongside the expenditure, so each product has an observed unit value in reais per kilogram in each region. The CES structure implies a log-linear relationship

between relative expenditure shares and relative prices,

$$\log\left(\frac{S_{nr}}{S_{hr}}\right) = (1 - \sigma) \log\left(\frac{P_{nr}}{P_{hr}}\right) + \mu_n + \varepsilon_{nr}, \quad (47)$$

for variety n in region r , measured against a reference variety h , with μ_n the variety's own level. Least squares cannot estimate it, since a variety's price and its demand shock are jointly determined. Identification instead uses the double difference of Feenstra (1994) as refined by Soderbery (2015). Differencing each variety against a reference variety within a region removes the regional price level, and demeaning each variety across regions removes its own level, so regions play the role that periods play in Feenstra (1994). The second moments of the doubly differenced shares and prices then identify the demand elasticity σ jointly with the inverse elasticity of supply, which is restricted to lie between zero and one. The estimator is identified by heteroskedasticity across varieties: the moment condition is a regression across them, so varieties must differ in the volatility of their demand and supply shocks for anything to be identified.

The two tiers differ in what a variety is and in how regions are defined. For the upper tier, varieties are product categories, the estimand is restricted to the model's domain of barcoded grocery categories and to the states of the Northeast, and regions are states. For the lower tier, varieties are the product types within a category, categories are the survey's own product groups, and regions are states crossed with the urban and rural strata, which roughly doubles the number of relative-price comparisons each type contributes. Finer regional partitions were examined and rejected: at the level of the survey's sampling strata the cells thin out, unit values grow noisy, and the requirement that a type appear in many regions cuts the sample by an order of magnitude.

The survey does not record barcodes, so the elasticities are estimated on its own units and applied to the model's. The model uses the substitution between product types within a survey group as the substitution between barcodes within an HS-8 category, and the substitution between survey categories, mapped to HS-8, as the substitution between HS-8 categories. A type's unit value averages over the products bought under it, so it moves with their quality mix as well as with their prices.

Standard errors are bootstrapped over categories rather than over varieties, since a category-level taste or supply shock moves every variety in the category together. Each standard error is the width of the 95 percent bootstrap interval divided by 2×1.96 , which treats the interval as symmetric. The lower tier keeps region-type cells with at least 5 purchases, types observed in at least 6 regions, and category-region pairs with at least two types.

D.3 Objects recovered by inversion

Each evaluation of the objective inverts two sets of equilibrium objects from the candidate parameters, the observed prices, and the observed assortments, so the model reproduces the data along those dimensions before moments are formed.

Rents are conditioned on their observed values, which enter sorting (5) directly and reach retail only through the fixed cost of holding a site.

Amenities are recovered by inverting the sorting equation (6) group by group, one amenity per tract and income group, each normalized to mean zero within its group. Inverting from a single reference group instead would force every other group's predicted allocation to be that group's spatial pattern raised to a power, which the data reject: the poorest and richest quartiles sort in opposite directions, so a common amenity fits the composition only by driving the richest group's dispersion to a corner.

The group-specific amenity has a cost that must be stated. Substituting it back into (6) reproduces the observed income composition of every tract identically, so the sorting block fits by construction and its fit is not evidence for the model. It is also why no composition moment can identify the sorting dispersion within this city.

D.4 The cost–quality exponent

The quality index is recovered from prices, so the cost–quality exponent and the taste weight (11) act on the same margin, how price varies with quality, and the model ties them together. Writing the markup factor as $\mu_{ifn} \equiv p_{ifn}/c_{in}$, taking logs of (16), averaging across the stores that carry a barcode, and differentiating across barcodes within a category gives

$$s_p = \phi + \Gamma, \quad s_p \equiv \frac{d \mathbb{E}_f[\log p_{ifn}]}{d \delta_i}, \quad \Gamma \equiv \frac{d \mathbb{E}_f[\log \mu_{ifn}]}{d \delta_i}. \quad (48)$$

The left-hand side is the slope of barcode-mean log price on the hedonic barcode effect. That effect is itself estimated from log prices, so the slope is a measured feature of the data and is close to one. The restriction says the quality price premium splits into a cost share and a markup share that must sum to the observed price–quality slope. A larger exponent means more of the premium reflects what the good costs to put on the shelf; a smaller one means more of it is margin.

Two consequences follow. The exponent is not a separate parameter to estimate: the markup share is a function of the demand primitives alone, so (48) determines it once those are known,

$$\phi(\theta) = s_p - \Gamma(\theta), \quad (49)$$

with the slope measured on the estimation sample. Fixing $\phi = 1$ would set the markup share to $\Gamma = s_p - 1$ and so fix the markup–quality slope from outside the model. The estimation instead derives ϕ from (48) at each parameter vector, so the markup–quality slope is the one the pricing conditions imply. The price–quality slope s_p is close to one by construction, so $\hat{\phi}$ near one is the statement that the pricing conditions imply a markup–quality slope near zero.

D.5 The sorting cross-section

Two features of the within-city problem close off the obvious routes to the sorting dispersion. The natural moment is a within-group orthogonality condition on observed income–quartile shares at the level of *áreas de ponderação*. That restriction is not stable across aggregations: imposed at the tract level it returns a negative dispersion, and imposed at the weighting-area level, on the same data and the same restriction, a positive one of similar magnitude. Two aggregations of one condition should not disagree in sign, and they do so because the restriction requires the amenity to be orthogonal to local rent while rents capitalize amenities. The group-specific amenity inversion closes the other route.

The parameter is not specific to this city, and it is estimated on the national 2010 sample described in the body. The identifying variation is historical. Consumption access is instrumented by the neighborhood’s grocery stock in 1995, reconstructed from the national establishment register. That register dates closures as well as openings, so the stock is the set of shops trading in 1995 and not only those that survived to be observed later, and it shifts where stores are without being a response to the 2010 income distribution. Each establishment in it is weighted by its firm’s size class in the current register, so a firm that grew after 1995 enters the 1995 stock at its later size. Observable geography that belongs in the amenity is controlled for directly: terrain slope and elevation from the Copernicus elevation model, and distance to the city centre. The regression weights each weighting area by the square root of its households, drops the areas with no grocery stock in 2010 or 1995, and measures distance to the centre from the household-weighted centroid of the city’s weighting areas. The first-stage F statistic is 2170.214 and the Anderson–Rubin set is [3.5, 8.75]. Standard errors and the first-stage statistic are clustered by city, over the 125 cities. Instrumenting with the 2000 stock instead delivers 6.412 (s.e. 1.435, $F = 3616.832$).

The regression leaves the rent coefficient free, so $\hat{\kappa}_\epsilon = 6.189$ is identified from the loading on consumption access. Two caveats apply to the estimate. The design is cross-sectional, so the exclusion restriction is that the historical location of commerce affects the 2010 income distribution only through current access; persistence in neighborhood amenities would violate it, which is why the geographic controls enter and why the estimate moves when they do. And in Almagro and Domínguez-Iino (2025), the closest structure in the literature, the coefficient on the log number of food stores is negative for every household type, which they define by age and family

structure, the opposite sign to the loading this model imposes.

D.6 Inference

The moment covariance is estimated by a block bootstrap over 200 draws that resamples establishments, which preserves the within-store barcode clusters. Both fitted moments are built from the store data, so this resampling carries their sampling variation; the census and the consumption survey enter through the first stages and the calibrated inputs, and their uncertainty enters through the terms below. Standard errors follow the sandwich formula $(J'WJ)^{-1}J'W\hat{\Sigma}WJ(J'WJ)^{-1}$, with the Jacobian evaluated by central differences at the estimate.

Both elasticities are fixed first stages, and each contributes a Murphy–Topel term $\text{Var}(\hat{\sigma})(\partial\theta/\partial\sigma)(\partial\theta/\partial\sigma)'$ to the covariance of the free parameters, with $\partial\theta/\partial\sigma$ obtained from how the moments respond to the same elasticity. The two are treated as independent. They come from different aggregations of one survey, so the covariance is not available, and omitting it can raise or lower the total depending on its sign and on the parameters' sensitivities to the two elasticities. The sorting dispersion enters none of the fitted moments, so it contributes no term to the covariance of the free parameters; it is instead evaluated at $\hat{\lambda}$ and inherits part of its sampling error. The retail-margin benchmark that fixes the cost level enters as a calibrated input and its uncertainty is not propagated, so the standard error of λ is conditional on it. The tract income shares, imputed from each tract's income moments (Appendix A.6), and the barcode qualities and HS-8 cost shifters $\bar{c}_n$ of Section 2 also enter as fixed inputs. The qualities and cost shifters are estimated from the same price records the bootstrap resamples but are not re-estimated within a draw, so the standard errors are conditional on all of these inputs and understate the uncertainty their estimation adds.

The population allocation is solved as a fixed point of (6) at each candidate parameter vector. With the group-specific amenity inverted at that vector, the sorting map returns the observed allocation whatever population it is given, so the fixed point is the data and the moments are formed at the observed allocation throughout. The sorting block therefore adds no information to the fit, which is why the sorting dispersion is estimated outside it.

D.7 Model fit

The moments in Table 4 discipline the level of the pricing condition and the stocking decision against the income mix of a store's customers. Store size enters neither. A larger store carries more barcodes, so it internalises more substitution when it prices, which widens its margin through C_{ifn}^{cat} and C_{ifn}^{store} , while its larger share of shoppers changes the traffic it loses through $\mathcal{T}_{ifn}$. Table 25 puts the model's margin against the data's by floorspace quartile, evaluating (17) at $\hat{\theta}$ for every barcode–store observation and taking the model-implied margin $L_{ifn}(\hat{\theta})$. The data runs from 0.251 in the smallest quartile to 0.221 in the largest, a gap of -0.030 , and the model

from 0.238 to 0.228, a gap of -0.010 . Both margins fall with size, the model's by less, so at the estimates the traffic term outweighs the internalisation terms, and the comparison tests the balance between them. A store without a matched building footprint is assigned the average shop size of its weighting area, so the quartiles sort stores partly by area and not only by their own size. The data row also uses the model's marginal cost, which carries no store component, so the comparison does not separate a margin that varies with size from a cost that does.

Table 25: Markup by store floorspace quartile: data against model

	Q1 (smallest)	Q2	Q3	Q4 (largest)
Data	0.251	0.235	0.220	0.221
Model	0.238	0.232	0.230	0.228

Notes: Mean markup $(p - c)/p$ across barcode-store observations, by quartile of the store's commercial floorspace. The data row uses observed prices against the model's marginal cost (16); the model row uses the margin the pricing conditions imply at $\hat{\theta}$. Their difference is the first-order-condition residual. Store size enters no fitted moment.

The pricing conditions do not reproduce how prices rise with neighborhood income in Section 3. Projecting the model's margin $L_{ifn}(\hat{\theta})$, converted to a log price at the barcode's marginal cost, on log neighborhood income within barcode, in each of 20 quality bins, gives -0.0066 on average below the premium end of the quality distribution against 0.0450 in the data, and -0.0036 against 0.0196 at the premium end. Marginal cost carries no store component, so the positive income elasticity of price in the data lies entirely in the first-order-condition residual, which therefore rises with neighborhood income. The counterfactuals hold each barcode-store residual at its observed value and re-solve only the margins, so the observed city is an exact equilibrium and prices move only through markups. A richer income mix raises prices in the data but not through any margin the model re-solves, so Section 8.3 understates the price increase a richer income mix brings and overstates the grocery gain to incumbents, which works against its conclusion that the grocery gain does not compensate them for the rise in rent.

E Counterfactuals Appendix

E.1 What is held fixed

Both experiments hold the estimated primitives at their baseline values. These are the preference parameters $(\sigma_b, \sigma_c, \alpha, \lambda, \kappa_\epsilon)$, the barcode qualities ξ_i , the category cost shifters $\bar{c}_n$, the neighborhood amenities $A_{\ell g}$, and the calibrated (τ, η, ρ_g) . Marginal cost is $c_{in} = c_0 \bar{c}_n \xi_i^\phi$ and carries no rent term, so rents reach the retail side only through entry, exit and commercial-floorspace demand and never through the pricing first-order condition.

Amenities are inverted once at the baseline allocation and then frozen. The amenity absorbs everything about a location the model does not name, so re-inverting it inside a counterfactual would absorb the treatment and return the baseline allocation. Freezing it attributes the reallo-

cation to the retail channel, and it is why the sorting responses are conditional on the amenity being orthogonal to the induced change in the grocery environment.

The two experiments of Section 8.1 differ only in the shock and solve the same equilibrium, so neither switches off a margin the other leaves on. Their responses are finite responses to different shocks, not derivatives at a common point, so they are compared and not multiplied.

E.2 The site and the treated area

The site is chosen from the tracts where a change in composition is plausible. A candidate tract has income between the 25th and 75th percentiles of tract income, a population at or above the 50th percentile of those tracts, reaches at least one store but no more than the 25th percentile of their reachable barcodes, and lies within 3.0 kilometers of a tract at or above the 75th percentile of tract income. Among the candidates, the site is the one with the highest store exposure, defined as the largest share of any reachable store's expenditure that the tract supplies. Exposure is the right criterion for both shocks, since it bounds how far a change in the tract can move the customers of any store it reaches.

The treated area grows outward from the site, nearest populated tract first, until it holds 3000 households or reaches 3.0 kilometers. In both experiments the radius binds first, and the area holds 2715 households in 12 tracts. The area is contiguous and is not restricted to candidate tracts: the screen chooses where the experiment takes place, and the treated area is the neighborhood around it.

The households placed in Section 8.3 come from the tracts with above-median income and, among those, the 287 in the top 33 percent of grocery access for the poorest quartile, drawn in proportion to each tract's top-quartile population.

E.3 The housing market

Writing $D_\ell = \sum_g \rho_g y_g N_{\ell g}$ for residential housing expenditure in a tract, the counterfactual rent solves the clearing condition of Lemma 2 against both claims on the stock,

$$\left(D_\ell^0 + H_\ell^{\text{firm},0} r_\ell^{\text{obs}} \right) x_\ell^{1+\eta} - H_\ell^{\text{firm}} r_\ell^{\text{obs}} x_\ell - D_\ell = 0, \quad r_\ell = r_\ell^{\text{obs}} x_\ell, \quad (50)$$

where $H_\ell^{\text{firm},0}$ is the floorspace of the stores operating in the observed city and H_ℓ^{firm} that of the stores operating in the counterfactual, entrants included and closed stores excluded. The first term is the supply of floorspace, calibrated so that the observed stock clears at observed rents; the second is the commercial claim on it. The positive root is unique and equals one when $D_\ell = D_\ell^0$ and $H_\ell^{\text{firm}} = H_\ell^{\text{firm},0}$, so the baseline reproduces observed rents exactly, and an entrant raises rent in its tract through its own claim as well as through the households it draws. The commercial claim is rent-inelastic and absorbs part of the supply response, so omitting it from (50) would

make rent more responsive to a residential demand shift than the model allows. With $\eta = 0.300$ at most 77 percent of a shift in demand passes into rent, attained where no store occupies the tract. Neither experiment adds housing units, so floorspace moves from the observed stock only along the supply elasticity as rents change.

E.4 The store margin

Which categories a store stocks is held at the observed set and the barcodes within them are predicted, so the category stocking cost enters through the cell threshold v_{fn} of (40), which is solved in each cell from its observed count. The remaining cost is the fixed cost of holding a site, which is not observed, so each store's exit margin is anchored to its own baseline position.

An incumbent closes when its operating profit falls below the lower of its own baseline coverage of rent and the thinnest surviving store's, by more than a relative tolerance of 10^{-5} , so no store closes in the absence of a shock. The 221 stores with a measured footprint whose operating profit is not positive at baseline are not rationalized by the model, which says they should not operate. A floor at their own baseline coverage would close them at the first shock of any size, which would reflect how their profit is measured and not the shock, so they are given no exit margin. Stores without a measured footprint have no rent bill to cover and no exit margin either. The participation condition (iii) of the equilibrium definition is therefore imposed on the stores the model rationalizes, the counterfactuals hold the others open, and the closures they report are those of rationalized stores only.

A candidate entrant is a clone of an observed store, planted dormant at a treated tract. At each tract the candidates are the 3 formats whose closed-form profit the change in the income mix of Section 8.3 raises most, with every format first pinned to break even there at baseline. The same candidates are planted in both experiments. Each candidate's overhead is set so that, switched on alone at the baseline and choosing its own prices and shelf, it exactly breaks even, which is the revealed-preference content of its absence from the observed city. It opens when the shock raises its profit by more than it raises its rent bill, and closes again if the rest of the city later moves against it. An entrant inherits its format's revealed stocking costs cell by cell and adapts that format's shelf to its own location and clientele by the rule incumbents follow. These are the most permissive overheads the data admit, so no planted candidate is kept out by its threshold; formats that are not planted are not considered. The store forced open in Section 8.2 is the treatment and keeps its format's shelf.

After the placement of Section 8.3, the same closed-form screen applied to every observed format at every treated tract finds 146 format-by-tract pairs above their thresholds, the largest by a margin of 0.0757 of its threshold. These are opportunities the planted candidates did not cover, and they indicate how much further entry the screen would allow.

E.5 The quality and price margins

Table 12 reports the quality and price terms of (12) net of each other. Separately, quality moves the treated area's index by -55.55 log points and price by 58.34 . Both move with the same shift in what households buy: they switch toward the entrant's barcodes, which are of higher quality and, with marginal cost rising in quality, of higher price. The incumbent shelf-quality rows of Tables 10 and 11 are unweighted means of $\log \xi$ over the incumbents' shelves and so move little, while the quality term weights each store by its share of the area's spending, of which the entrant takes 13 percent. The price term is a quality-weighted mean price of the shelf, so it moves with that composition even where no barcode is repriced, and the two terms are informative only together.

E.6 The shelf response

The shelf response uses two results of Appendix C.3. The stocking condition is, to first order, the condition for a shelf that no single addition or deletion improves (Lemma 1), and the observed shelf pins each barcode's cost draw to one side of its cell's threshold (Proposition 1). The counterfactual therefore never needs the level of the category fixed cost, only how the shock moves each barcode relative to the threshold its own store has revealed.

What neither result delivers is a counterfactual shelf that can be computed one barcode at a time. A barcode's contribution depends on the rest of its shelf, and adding one lowers every shelfmate's: at the estimated parameters one added barcode lowers its cell's expected count by more than one, so the expected count evaluated at the shelf it predicts has no fixed point in whole barcodes in many cells. Each barcode's stocking condition is therefore evaluated against its store's observed shelf, as the estimator evaluates it, while prices, the income mix of the store's customers, their spending and the population all take their counterfactual values. An entrant's reference shelf is the one it chose on opening. Shelves still enter prices, the location price index and sorting through the equilibrium itself. What is given up is the interaction among barcodes that move at the same time. Barcodes added together cannibalize one another, which the rule ignores and which overstates additions; barcodes dropped together each raise what their shelfmates contribute, which the rule also ignores and which overstates drops. With both margins active the error in the net variety response is not signed.

The count follows the model's expectation. Each cell's threshold is solved at baseline so that its expected number of barcodes carried equals the observed one. The threshold absorbs the level of the cell's expenditure, so holding it fixed while expenditure moves requires adding the log change in the cell's expenditure to each barcode's index. The cell then carries the observed count plus the change in that expectation, rounded to whole barcodes with a uniform drawn once per cell and held fixed. Which barcodes move follows the data. By Proposition 1, an uncarried barcode joins with the probability, conditional on its being uncarried at baseline, that

the shock lifts it over its threshold, and a carried one leaves with the corresponding probability that the shock takes it below. The barcodes with the highest such probability are the ones that move. The observed shelf is reproduced exactly at baseline, and a cell whose threshold does not solve is held at its observed shelf.

Entry of a barcode a store has never carried is bounded. Materializing every uncarried barcode for every store is not feasible, so each cell admits the 5 uncarried barcodes whose contribution sits closest to its threshold and the 5 highest-quality barcodes it does not carry. The estimated index rates the cheapest barcodes as the most profitable additions, so a pool drawn from the margin alone would let a shelf grow but never upgrade. With both ends offered, the shock decides which barcodes join. A cell that exhausts its pool has had its count set by the bound and not by the model, and 1 cells do in Section 8.3.

E.7 The frozen-retail comparison

Table 13 solves the placement of Section 8.3 a second time with the retail sector held at the observed equilibrium: every store keeps its observed prices and shelf, no candidate is planted and no incumbent may close. The location values then do not move, and households respond only to the placement and to the rents it produces. Three objects follow from the two solutions. The mean income of the treated area under the placement before anyone re-sorts is the mechanical part of the change; the frozen solution less that allocation is the part rents produce through sorting; and the full solution less the frozen one is the part the retail response adds.

E.8 Solution and equilibrium selection

Each experiment alternates four blocks from the observed allocation until none of them moves. Prices are re-solved for every store whose customers' income mix or spending has changed, by a predictor-corrector homotopy from the observed prices on the first pass, holding each barcode-store pricing residual at its observed value (Appendix D.7). Shelves are re-solved as above. Population and rents are solved to the fixed point of the sorting map given the location values. Then candidates may open and incumbents may close. From the second pass on, the population moves a share 0.50 of the way from its previous value toward the sorted allocation. A fixed point of the unrelaxed alternation is a fixed point of the relaxed one, and the relaxation removes a period-two oscillation that the unrelaxed alternation otherwise sustains in Section 8.2. Where the retail side has several equilibria, the relaxation rate is part of the selection rule, since it can change which of them the alternation reaches. The alternation stops when the largest change in any tract's population is below 10^{-5} of the largest tract or has stopped falling, the barcodes added and dropped are below 10^{-4} of those carried, and no store enters or exits. Prices are solved to 10^{-7} within each pass for every store whose customers' income mix or spending moved by more than 10^{-4} , so on the final pass, where no store's customers move by that much, prices are settled

as well. Both experiments meet these conditions within 40 passes.

A dormant candidate has no customers and must be switched on to be evaluated. Each is first screened at the prices and shelf it chose at baseline, which by the envelope theorem gives the change in its profit to first order, allowing a gain from re-optimizing those choices as large as the first-order change itself. A candidate that passes is switched on alone, re-solves its prices, chooses its shelf and re-solves its prices again, and opens if its profit then clears its threshold. Candidates are admitted one at a time, in order of their baseline profit relative to rent, and the rest are re-evaluated against the network each entrant creates.

Any multiplicity the solver faces comes from the retail side. With the network, the shelves and the prices held fixed, rents and the allocation of households are unique for every parameter value, with the placed households held in place (Proposition 5), so the population step of each pass and the frozen-retail solution are well defined. Once prices, shelves and entry respond, uniqueness is not available from the primitives (Appendix C.4). The reported equilibrium is the one the alternation reaches from the observed allocation, and the two experiments are comparable since each walks the same alternation from the same observed equilibrium.

The solver refuses to report a result unless four checks pass. The planted candidates, while dormant, leave the baseline unmoved; the latent barcode rows leave it unmoved; the stocking moment rebuilt from the counterfactual's own cell state reproduces the estimation's to within $1.0e - 11$; and a null shock moves no price, barcode, store or household.