Supply chain concentration in traditional retail*

Diego Santa María
Harvard University

November 16, 2026
[Click here for latest version]

Abstract

Corner stores selling consumer packaged goods are ubiquitous in developing countries, suggesting a competitive environment. Their supply chains, however, are highly concentrated, and suppliers' decisions about which stores to serve can limit their access to products. This paper asks whether greater upstream competition expands the variety stores offer, and how this affects their growth and consumer welfare. I randomize a new wholesaler's entry across 364 neighborhoods in Lima, offering stores an alternative source for a broad range of products. About ten months later, I survey 1,188 stores and collect their wholesale receipts over three weeks, covering 43,500 product-level transactions. Small, space-constrained stores in treated neighborhoods carry 19% more distinct products and spend 20% more on wholesale purchases, with no detectable gains among larger stores. Most of this growth comes from incumbent suppliers, who respond to competition by improving non-price terms. To rationalize these findings, I develop a model in which suppliers save delivery costs by concentrating sales among fewer stores. A wholesaler allows excluded stores to compete for those sales, making them more attractive for incumbents to serve. I estimate that consumer welfare gains from greater variety at nearby stores are equivalent to 10% of baseline travel costs.

*I am grateful to my advisors Emily Breza, Edward Glaeser, and Gabriel Kreindler for their generous advice and support. I also thank Pol Antràs, Jie Bai, Sachet Bangia, Augustin Bergeron, Tilman Graff, Deivy Houeix, Torsten Jaccard, Myrto Kalouptsidi, Dean Karlan, Robin Lee, Oluchi Mbonu, Pamela Medina, Santiago Medina, Marc Melitz, Oswaldo Molina, Vishan Nigam, Konstantin Poensgen, Lucina Rodriguez, Ben Roth, Emanuel Schertz, Meredith Startz, Edward Wiles, and participants at Harvard's Development, Trade, and IO workshops, NEUDC, the *Seminario de Microeconomía Aplicada en Perú*, the TRIOS workshop at UBC, and the IHS Entrepreneurship workshop for valuable feedback and discussions. Antony Ore provided excellent research assistance at every stage of the project. I am also grateful to Leonardo Gonzales, who sadly passed away before the project was completed, and to Paulo Carrasco for outstanding fieldwork coordination, as well as to all the enumerators and data-entry personnel who participated in this project. I am especially indebted to Eduardo Casado, whose trust made this study possible. I acknowledge financial support from the Weiss Fund, the Institute for Humane Studies, the Weatherhead Center for International Affairs, the Institute for Quantitative Social Science, SurveyCTO, and the Harvard Economics Department's Griffin, Garon Ross, LEAP, and Warburg Funds. The study is registered in the AEA RCT Registry (AEARCTR-0016388) and was approved by the Harvard University IRB (IRB23-1307). Author contact: dsantamaria@g.harvard.edu.

1 Introduction

In developing countries, consumer packaged goods such as beverages, pantry staples, and personal care products are mainly sold at small, independently owned corner stores. Although the spatial density of this sector suggests a competitive environment, upstream distribution is typically concentrated: a few large suppliers, each specializing in a narrow range of products, make regular deliveries through their own distribution networks (Kremer et al., 2013b). When logistics costs are high, suppliers with market power may not find it worthwhile to serve some stores, leaving them without access to products that are hard to source elsewhere. These limits on store variety may have substantial implications for entrepreneurs' incomes and consumer welfare, especially in low- and middle-income settings where preferences for proximity have been shown to be strong (Lagakos, 2016; Ramos-Menchelli and Sverdlin-Lisker, 2026).

In this paper, I ask whether greater upstream competition allows corner stores to offer more variety, and how this affects their growth and consumer welfare downstream. I study the entry of a new wholesaler that buys in bulk from suppliers and resells their products to corner stores in Lima, Peru.[1] This constitutes a shock to stores' outside options: they can source a broad range of products without buying directly from each product's main supplier. An outside option could particularly benefit relatively small stores, which are more likely to be excluded from supplier distribution networks, and allow them to expand the variety they offer. The equilibrium effects of a wholesaler's entry depend, however, on whether and how incumbents respond to direct competition for sales of their products, and how much consumers value greater variety at nearby stores. Answering these questions requires exogenous variation in supply chain structure across many comparable markets and detailed data on store operations, which are rarely available in developing countries.

I overcome these challenges by partnering with the wholesaler to randomize its entry across 364 low- and middle-income neighborhoods, covering over one quarter of Lima's population. Sales agents visit treated neighborhoods weekly, generating market-level variation in exposure to the entrant. Nine to thirteen months after the initial rollout, I survey 1,188 corner stores and collect a wholesale diary of all their wholesale orders over a three-week window, with receipts for most supplier orders. These data allow me to quantify the impacts of upstream competition on product variety, wholesale expenditure, and supplier relationships. To guide the empirical analysis, I develop a model in which the wholesaler's entry allows previously excluded stores to source more products and compete more effectively for consumers. The model predicts that incumbent suppliers may defend their market share by serving more stores or improving the terms they offer.

The experimental increase in upstream competition has large effects on the organization of traditional retail downstream that are consistent with the predictions of the model. I present three main findings: (1) exposure to the entrant leads to variety gains concentrated among small stores; (2) small stores spend more on wholesale purchases, and most of that growth comes from incumbent suppliers; and (3) incumbents respond through non-price

[1] This business model is becoming increasingly common in developing countries. Examples include Warung Pintar in Indonesia, MaxAB in Egypt, TradeDepot in Nigeria, and MarketForce in Kenya (Nieuwoudt, 2023).

terms rather than wholesale prices. Using a complementary consumer survey and a multiproduct extension of the model, I estimate that increased variety at nearby stores yields consumer welfare gains equivalent to 10% of baseline travel costs.

Corner stores account for a large share of the aggregate economy across Latin America, Africa, and Asia: roughly 10–20% of businesses and, in most countries with data, over a third of grocery spending (Appendix C). In northern Lima, the setting of the experiment, there is one corner store per 120 inhabitants, and consumers visit a store at least once a day on average. Using the store survey, I document that the supply chains serving these corner stores are highly concentrated: only ten suppliers account for 81% of supplier purchases, with minimal overlap in their product portfolios. I also find heterogeneity in product variety across stores: *large* stores above the control-group median floor area order from 35% more suppliers and purchase 64% more distinct products than *small* stores with less space. Floor area measures available space—a key determinant of corner store productivity[2]—and may also reflect other drivers of business performance that influence which stores suppliers choose to visit, such as access to capital and entrepreneurial skills. Product access may therefore vary across households, even in neighborhoods with many stores.

The entrant is part of a modern retail conglomerate that negotiates bulk purchases for all its formats, so suppliers cannot prevent it from selling to corner stores. It offers a broad range of products (accounting for over 60% of corner stores' purchases), competitive wholesale prices, and flexible ordering and delivery through regular visits and a digital app. To structure its operations, the wholesaler divided the study area in northern Lima into 119 similarly sized territories, called *cuadrantes*, each visited weekly by a sales agent to take and deliver orders. Using these territories as randomization strata, I divided each into three contiguous clusters with about 60 stores each, and randomly assigned one to the control group.[3] Sales agents were free to choose which stores to visit within treated clusters, but no active outreach was allowed in control clusters until after data collection.

The survey took place nine to thirteen months after the wholesaler's entry, giving it enough time to consolidate its presence and allowing stores and incumbent suppliers to adjust. Surveyed stores reported all their wholesale purchases over a three-week window, and the field team collected and digitized receipts for 71% of supplier orders. The resulting dataset covers more than 15,000 orders and 43,500 product-level purchases, with harmonized wholesale prices and quantities. It includes purchases and other relationship terms across all of a store's suppliers, rather than just the entrant. I use these data to construct the main outcomes for the analysis: product variety, wholesale expenditure, and incumbent supplier terms, including non-price terms such as trade credit, free goods, and equipment transfers. I also collected self-reported measures of financial performance, along with audit data on availability and retail prices for six popular products.

Demand for the new wholesaler and its clients' purchasing patterns show that it is an attractive sourcing option for corner stores. By the time of the survey, the entrant was a regular source of supply for roughly 30% of stores in treated neighborhoods—about three

[2]A physically smaller store faces a higher opportunity cost of stocking a product, as it takes up space that cannot be used for another one.

[3]Markets and their surrounding areas were excluded from the randomization. In a few strata without markets, I drew four rather than three clusters.

times the share in control neighborhoods[4]—and accounted for about 20% of total wholesale expenditure among those clients. Orders from the entrant are nearly twice as large in value and span 50% more product categories than orders placed through incumbents, highlighting its appeal across a broad product range. Within-product comparisons confirm the entrant's competitive terms: stores pay 5% lower wholesale prices when buying from the entrant.

I next present the equilibrium impacts of the wholesaler's entry on access to products, store scale, and incumbent supplier responses.

First, I find variety gains concentrated among small, space-constrained corner stores. Small stores in treated neighborhoods purchase 19% more distinct products than those in control neighborhoods ($p = 0.037$), while the effect among large stores is indistinguishable from zero (though still positive). Across all stores, the estimate is 9% and imprecise ($p = 0.229$), but significant in logs (14%; $p = 0.041$). Variety gains among small stores are largely driven by an increase of 13% in categories purchased ($p = 0.050$), reflecting a broader range of products rather than additional close substitutes for products they already carry. Small stores also shift their purchases toward products with more concentrated distribution channels, raising the average supplier concentration (HHI) of the products they buy by 3.6% ($p = 0.043$). Together, these patterns are consistent with incumbent suppliers excluding small, less efficient stores from their distribution networks before the wholesaler's entry.

Consumer gains are not limited to increased variety at small stores: the retail audit reveals lower stockout probabilities and, at large stores, lower retail prices. Across six popular products, exposure to the entrant reduces the probability that at least one is unavailable by 5.8 percentage points, from a control mean of 33.7% ($p = 0.029$). Average retail prices for these products do not change detectably across all stores, but fall by 0.6% at large stores ($p = 0.019$). I discuss possible explanations for the latter result below.

Second, variety gains among small stores are accompanied by a 20% increase in their total wholesale purchases ($p = 0.001$), flattening the relationship between store size and wholesale expenditure in treated neighborhoods. Small stores' expansion occurs without detectable reductions in their retail prices, consistent with consumers responding to variety gains rather than pass-through from improved wholesale terms. Estimated effects on aggregate wholesale expenditure are positive but small and imprecise in levels, so I cannot rule out full demand reallocation across stores.[5] To provide a direct test for spillovers, I construct a measure of exposure to treated stores that exploits proximity to neighborhood borders, which I instrument using the recentering procedure in Borusyak and Hull (2023). The results provide suggestive evidence that part of small-store growth comes at the expense of nearby stores.

Crucially, small-store growth is not driven only by purchases from the entrant: incumbent suppliers account for over three quarters of these additional purchases. Large stores, by contrast, appear to substitute purchases toward the entrant rather than expand their total wholesale purchases. These patterns hold when restricted to the 71% of supplier orders

[4]Take-up is not zero in control neighborhoods due to navigation mistakes by sales agents and self-enrollment in the app.

[5]Effects on self-reported measures of wholesale purchases and sales provide suggestive evidence of reallocation across stores but are very noisy, which I attribute to measurement error. Restricting the sample to stores whose reported purchases agree more closely with the wholesale diary yields larger effects on sales for small stores, though not enough to reach conventional significance levels (Appendix G).

backed by a receipt. Expenditure on visits to wholesale markets or modern stores—an alternative outside option, of limited importance for most stores—does not change significantly for either group.

Third, incumbents respond to upstream competition by improving their service and non-price terms for small stores, while also increasing equipment transfers to their larger clients. Increased small-store purchases from incumbent suppliers reflect more supplier visits (7%, p = 0.079) as well as more orders per visit (5%, p = 0.062) and, less precisely, higher expenditure per order (9%, p = 0.104). Small stores also become 8.5 percentage points more likely to receive free goods, from a control mean of 48.4% (p = 0.038), and 5.3 percentage points more likely to receive trade credit, from 21.0% (p = 0.086). At large stores, I find that incumbents increase the provision of highly valuable storage equipment such as shelves and coolers—roughly one additional unit of each for every three large stores (p = 0.035 and p = 0.004, respectively). These results are consistent with incumbents' efforts to retain their largest clients and strengthen their position amid growing competition from nearby stores. More efficient storage might partly explain why retail prices fall at large stores, and may also contribute to the lower stockout probabilities.

In contrast to observed effects on non-price terms, I find no evidence of incumbents adjusting wholesale prices in treated neighborhoods: product-level analyses reveal very precise null effects, and I can rule out reductions larger than 0.3% at the 95% confidence level. When including purchases from the entrant, the estimated effect on wholesale prices is -0.1%, consistent in magnitude with its wholesale price advantage and its share of transactions but not statistically significant (p = 0.287). Overall, the evidence points to incumbents responding through non-price terms to capture growing demand from small stores and defend their positions at large stores. I conjecture that the rigidity of incumbent suppliers' pricing policies in the experiment's setting makes non-price terms the more suitable margin to respond to localized competition.[6]

To rationalize the empirical findings, I develop a theoretical model that shows how logistics and upstream market power shape the organization of traditional retail. The model explains why incumbent suppliers optimally exclude some stores from their distribution networks, and how a wholesaler's entry can reverse that decision even when its wholesale prices are not lower. The basis is a single-product Hotelling environment where stores differ in location and handling costs (large stores are more efficient).[7] A monopolist supplier chooses a common wholesale price and which stores to serve, paying a fixed delivery cost per store. Stores that are served then decide which retail price to charge, and consumers choose whether and where to buy the product by trading off retail prices against travel costs. The key insight is that the supplier has an incentive to serve more stores: this brings the product closer to consumers, creating additional surplus that it can extract by raising wholesale prices.[8]

[6] Anecdotally, the entrant's sales agents report instances of incumbent suppliers lowering wholesale prices globally to match the entrant's more competitively priced products, but these responses happen at a broader geography than the randomization clusters.

[7] I abstract from interactions between products to isolate a supplier's incentives to exclude stores and its response to direct competition. Allowing for many suppliers selling different products would lose tractability but deliver similar intuition, modulated by the degree of substitutability between products.

[8] This is closely related to the kinked-demand region in Salop (1979), where additional retailers can raise

However, the additional profit from serving less efficient stores may not cover the associated delivery costs, so the supplier excludes those stores from its distribution network even though they would buy at the prevailing wholesale price.

I then consider an entrant offering the same product at the wholesale price the incumbent charges when it is the only supplier available. Previously excluded stores can now carry the product regardless of whether the incumbent serves them, and compete with the incumbent's original clients for consumers. The model generates predictions for the incumbent supplier's response. If additional delivery costs are not too high, the supplier may now find it worthwhile to serve previously excluded small stores, capturing purchases that would instead go to the entrant. Otherwise, it may prefer to lower the wholesale price it charges its existing clients (the large stores) to defend their market share. These inclusion and retention incentives help interpret the experimental results: incumbents increase service to small stores and improve non-price terms for large stores. In the model, retention operates through a lower wholesale price, whereas in the experiment it is consistent with greater provision of supplier-owned equipment. The entrant's presence also prevents the wholesale price hikes that network expansion would allow under monopoly. Aggregate consumer surplus can, but is not guaranteed to, increase through lower travel costs.

In the final part of the paper, I conduct a complementary consumer survey and combine it with a demand model to quantify the consumer welfare gains from the experiment's increase in variety. I randomly selected 30 control neighborhoods and one block within each, and surveyed consumers living in and around those blocks about their recent shopping behavior. Enumerators first mapped nearby corner stores, providing measures of variety, retail prices, and locations for the stores in a consumer's choice set. Consumers shop locally and value variety: they walk on average 3 minutes round trip, and stores in the top half of brand counts receive twice as many weekly visits as other stores.

I then extend the model's consumer demand from a single-product to a multiproduct setting. On each shopping occasion, consumers draw preferences for products and choose a store by trading off travel costs and retail prices against the utility of their best product match. They buy the product they like best at the chosen store. More variety therefore improves the match between consumers and available products, while idiosyncratic tastes for stores allow destination choices to differ even at similar distances and variety levels. I estimate this model by simulated maximum likelihood on weekly visit counts, using brand counts to measure product variety. The estimates reflect strong preferences for both proximity and variety: when one store offers twice as many brands as another, the extra variety compensates for additional walking equal to 36% of the average round trip. I simulate the experiment's increase in variety by applying the estimated size-dependent treatment effects on brand counts to the stores in the complementary survey's sample. The resulting welfare gain is equivalent to 10% of baseline weekly travel costs (90% CI: 3.1 to 20.5%)—roughly the travel cost of two-thirds of an average shopping trip per week—or a 5% reduction in retail prices (90% CI: 0.7 to 19.8%). Extrapolated to the approximately 2.5 million people living

retail prices as distances become shorter. In the model developed here, the supplier selects this region by raising the wholesale price until the marginal consumer is indifferent between buying and not buying. The distribution decision also relates to the choice of outlet density in Dixit (1983) and Mathewson and Winter (1983), who weigh the value of proximity against the fixed cost of an outlet.

in the study area, these gains are equivalent to 4.8 million walking minutes per week.

This paper contributes to several strands of literature. First, it builds on a long literature in development economics on the constraints that keep small firms from growing.[9] Given the size of the traditional retail sector (see Appendix C), small retailers make up a large share or the entirety of the sample in most of these studies. This paper relates particularly to work on how search costs, limited information, and weak contract enforcement can prevent firms from forming valuable trading relationships (Cai and Szeidl, 2018; Cai et al., 2024; Falcao Bergquist et al., 2024; Startz, 2025; Wiles and Houeix, 2026; see Szeidl, 2025 for a review). I show that limited access to suppliers can also arise from suppliers' own choices, and these choices depend on the competition they face. In Killeen (2026), for example, retailers are more willing to stock a new product when assured that its supplier will keep serving them.

My findings also complement recent evidence that competition among small firms in developing countries shapes the effects of growth interventions. Studies that offer credit, training, or subsidies to some firms in a market find that their gains can come largely at the expense of competitors (Cai and Szeidl, 2024) or spur growth and value creation in the market as a whole (Busso and Galiani, 2019; McKenzie and Puerto, 2021; Banerjee et al., 2026; Karlan et al., 2026). The intervention I study changes the sourcing options of every store in a treated neighborhood, and I find suggestive evidence that part of the gains of small stores come at the expense of nearby stores. These results are consistent with consumers' strong preference for proximity, a mechanism highlighted in previous work studying the prevalence of small firms in developing country retail (Lagakos, 2016; Ramos-Menchelli and Sverdlin-Lisker, 2026).

Second, this paper relates to an extensive literature in industrial organization on how suppliers with market power shape the retail networks that carry their products (see Rey and Tirole, 2007; Rey and Vergé, 2008; Lee et al., 2021, for reviews). Theoretical work shows that suppliers weigh the value of bringing products closer to consumers against the fixed cost of operating an outlet (Dixit, 1983; Mathewson and Winter, 1983), and that a supplier may limit the number of retailers it serves, so that competition among them does not erode its market power (Hart and Tirole, 1990; McAfee and Schwartz, 1994; Segal, 1999; Segal and Whinston, 2003). I contribute by inducing rare exogenous variation in upstream competition across markets, and studying the resulting changes in suppliers' distribution decisions. The results show that incumbent suppliers respond to upstream competition by serving more stores, consistent with the marginal benefit of doing so increasing once excluding a store no longer keeps it out of the market.

A related line of work argues that incumbent suppliers can use exclusive dealing[10] and

[9] This literature, reviewed by Quinn and Woodruff (2019) and Verhoogen (2023), has studied the role of access to capital (de Mel et al., 2008; McKenzie and Woodruff, 2008) and labor (de Mel et al., 2019; Hardy and McCasland, 2023), business practices (Karlan and Valdivia, 2011; Brooks et al., 2018; Dalton et al., 2021; Samaniego de la Parra and Shenoy, 2024; see McKenzie, 2021 for a review), and behavioral biases and preferences (Kremer et al., 2013a; Beaman et al., 2014; Bai et al., 2025; Killeen, 2026).

[10] For theoretical analyses of the conditions under which exclusive dealing can foreclose rivals, see Aghion and Bolton (1987); Mathewson and Winter (1987); Rasmusen et al. (1991); Bernheim and Whinston (1998); Segal and Whinston (2000); Fumagalli and Motta (2006).

other terms of their vertical relationships to prevent retailers from accommodating a competing supplier (Asker and Bar-Isaac, 2014). Empirical studies find that these arrangements can keep rival products out of stores, potentially limiting the variety available to consumers (see Nurski and Verboven, 2016; Conlon and Mortimer, 2021; Hristakeva, 2022; Asker (2016) finds no evidence of foreclosure from exclusive dealing in beer distribution). In a developing country setting, Ghani and Reed (2022) show that incumbent ice retailers in Sierra Leone respond to increased upstream competition by extending more credit to their loyal buyers, in an attempt to retain them.[11] I also find improvements in the non-price terms that incumbents offer their large clients, consistent with efforts to defend sales amid competitive pressure from an entrant.[12]

Third, it adds to an emerging literature on wholesale intermediation. Research in international trade emphasizes the role of wholesalers in spreading fixed sourcing costs across trading relationships, which allows buyers to access products that would otherwise be unprofitable to source directly (Åkerman, 2018; Ganapati, 2024; Tsai, 2025). Grant and Startz (2026) bridge this literature with development economics by building and estimating a model that shows how longer intermediation chains can bring products to remote markets and benefit consumers, despite possibly higher retail prices.[13] The entrant I study plays a similar role for corner stores: by delivering many suppliers' products together, it can serve stores that each specialized supplier finds too small to serve.

Rather than studying the emergence of intermediaries as an endogenous outcome, however, I exogenously vary their entry across markets to study how it shapes market structure and consumer welfare. Two previous papers provide experimental evidence on the entry of intermediaries in retail settings. In a randomized controlled trial in rural China, Couture et al. (2021) find that access to e-commerce lowered the cost of living for a minority of consumers, mainly by overcoming logistical barriers, with little effect on the earnings of rural businesses. Closest to this work, Iacovone and McKenzie (2022) randomize access to an order-aggregation and delivery service across produce retailers in Bogotá and show that the service allows them to bypass wholesale markets, reducing travel time and wholesale prices.[14] I build on their findings by showing that an intermediary's entry may also affect the behavior of incumbent suppliers, and that these effects are essential to understanding equilibrium outcomes.

The rest of the paper proceeds as follows. Section 2 describes the context. Section 3

[11] In contrast, Macchiavello and Morjaria (2021) and Leone et al. (2025) find that when contracting environments are imperfect, more competitive market structures in agricultural supply chains can harm rather than help farmers by reducing the services that intermediaries provide to them.

[12] Suppliers facing more competition also extend more trade credit in Fisman and Raturi (2004) and Fabbri and Klapper (2016). Firms often set prices uniformly across locations despite local differences in demand and competition (DellaVigna and Gentzkow, 2019), which may explain why responses to the wholesaler's entry operate through non-price terms.

[13] A related literature shows that intermediary market power along supply chains in developing countries lowers the share of surplus that reaches smaller farmers and consumers (Atkin and Donaldson, 2015; Mitra et al., 2018; Falcao Bergquist and Dinerstein, 2020; Zavala, 2022; Chatterjee, 2023; Dhingra and Tenreyro, 2025; Domínguez-Iino, 2025). Casaburi and Reed (2022) find a competitive intermediary market.

[14] Bartkus et al. (2022) show that joint investments by fishing communities in the Amazon allow them to reduce the costs of accessing markets directly, bypassing middlemen and increasing their earnings.

presents the entrant and the experimental design, and Section 4 describes the data. Section 5 outlines the conceptual framework and the empirical strategy, Section 6 presents the results, Section 7 develops the model, Section 8 quantifies consumer welfare gains, and Section 9 concludes.

2 Context

2.1 Corner stores in Lima

Across Latin America, Africa, and Asia, corner stores account for a large share of the aggregate economy: roughly 10–20% of businesses and, in most countries with data, over a third of grocery spending (Appendix Tables C1 and C2). In Peru, corner stores account for 36% of grocery spending and approximately 9% of nonfarm micro and small businesses.[15] They specialize in packaged goods such as beverages, pantry staples, snacks, and cleaning and personal-care products, while markets (and increasingly, modern stores) are typically used for fresh products and larger purchases. Fresh products account for only 4% of corner stores' reported sales.[16]

In northern Lima, the setting of the experiment, there is one corner store per 120 residents, and 98% of stores have another store within 100 meters (Appendix Figure B1). Consumers use corner stores for frequent, small purchases, making proximity particularly valuable. In the complementary survey described in Section 8.1.1, consumers make 7.4 visits per week, with an average round-trip walk of 3 minutes, and spend about 8 USD PPP per visit. Retail prices also vary little across stores, consistent with supplier efforts to coordinate them: across six popular products covered in the retail audit (described in Section 4.2), an average of 64% of control stores carrying a product charge its modal retail price across the study area. The variety available at nearby stores may therefore matter for consumer welfare even in neighborhoods with many stores to choose from.

Corner stores are predominantly home-based businesses run by middle-aged women, providing a large source of self-employment in low-income neighborhoods. Among control stores in the survey, 73% of owners are women, their average age is 50, and 19% have higher education. 57% of owners keep business records and 65% report paying taxes. These businesses have operated for 9.5 years on average, and 91% operate from the owner's home. Their median floor area is 20 square meters, and only 5% have paid employees (Appendix Figure B2 illustrates stores in the study area). Despite their small scale, they remain open for long hours, averaging 12.7 hours per day, with 84% opening every day.

The average control store buys from 7.3 suppliers over the three-week wholesale diary and places four supplier orders per week. The wholesale diary records 323 USD PPP in weekly wholesale expenditure on consumer packaged goods for the average store, against

[15] Appendix C gives the sources, reference periods, and construction of the business and spending shares.

[16] Among control stores, the largest categories by share of wholesale expenditure are beer (23%), soda (22%), snacks and candy (11%), pantry staples (11%), cleaning and personal-care products (12%), dairy (7%), and bottled water (4%). Pantry staples include canned fish, oil, rice, and pasta. Cleaning and personal-care products include detergent, toilet paper, and hair care.

self-reported sales of 915 USD PPP.[17]

2.2 Larger stores carry more product variety

Despite their close proximity, stores differ substantially in operating scale and the product variety they offer. I use floor area to characterize these differences because available space is a key determinant of corner store productivity: anecdotally, owners emphasize the importance of product rotation, since slow-moving products occupy space that cannot be used for another product. Consistent with this, owners report using 83% of their storage space on average. Floor area may also reflect other drivers of business performance that influence which stores suppliers choose to visit, such as access to capital and entrepreneurial skills. Smaller stores may therefore be more likely to face supply constraints because they are less attractive clients for suppliers, even when they have space to carry additional products.

Larger stores' higher wholesale expenditure largely reflects purchases of a wider range of products. Large control stores spend 74% more on wholesale purchases than small stores (Appendix Figure B3). In the receipt-verified sample, more suppliers account for 57% of the gap in mean log expenditure, and more products per supplier account for a further 19%, with only about one quarter coming from expenditure per product (see Figure 1 and Panel B of Appendix Figure B3).[18] The complementary survey, described in detail in Section 8.1.1, supports interpreting these purchase differences as differences in supplier service and the product variety available to consumers: within neighborhoods, stores in the larger half of the floor-area distribution receive visits from 40% more suppliers and carry 31% more brands (Appendix Figure B5).

2.3 Supply chains are concentrated

Although corner stores are numerous, the supply chains serving them are highly concentrated. The four largest suppliers account for 60% of incumbent-supplier purchases, and the ten largest account for 81% (Appendix Table A1). Most specialize in a narrow range of products, with their top three categories accounting for 78% of purchases from them on average. Supplier portfolios also overlap little: on average, 80% of spending on a supplier's products goes to that supplier. Access to each product's main supplier therefore determines much of the variety a store can offer.[19]

Each supplier operates a separate distribution network, with sales agents visiting stores on fixed routes, typically weekly, to take orders that trucks deliver over the following days. Stores that are not served can buy from wholesale markets or modern stores, but these alternatives are often far away and their wholesale prices are not always lower, as Kremer et al. (2013b) also document in Kenyan retail. As a result, purchases on these trips account for only 15% of recorded wholesale expenditure among control stores. Stores therefore have lim-

[17] The implied margin of 65% is implausibly high, considering anecdotal reports of markups around 25–30%. I return to issues in self-reporting in Section 6.3.2.

[18] Consistent with this, conditional on buying the same product, large stores buy 14% more packages per transaction ($p = 0.003$, Appendix Figure B4).

[19] Similar supply chain structures are common in other countries. For illustration, Ramos-Menchelli and Sverdlin-Lisker (2026) use records from a single large supplier covering the universe of corner stores in Mexico.

ited alternatives when a product's main supplier chooses not to serve them, leaving suppliers considerable discretion over which stores can carry their products.

Wholesale prices vary little across stores, consistent with suppliers setting common prices across their distribution networks rather than adjusting them to local conditions. Across 13,204 control transactions in 993 products, 61% are priced within half a percent of the product's median (Appendix Figure B6). Suppliers have more scope to adjust non-price terms across stores, even within neighborhoods. Minimum order values and discrete package sizes affect how much a store needs to buy at once, with 73% of control supplier orders subject to a positive minimum order value. Suppliers also provide free goods in 22% of orders and trade credit in 4% of observed store–supplier relationships. Equipment loans can substantially increase a store's capacity to carry products, with the average store holding 1.7 supplier-provided coolers and 1.8 supplier-provided shelves. Commercial coolers in Peru are worth roughly 1,000–2,200 USD PPP, and suppliers often require a commitment from stores receiving them. These non-price terms therefore give suppliers a way to support stores' operations and strengthen client relationships without changing their wholesale prices.

3 Intervention and experimental design

3.1 The entrant

To study the effects of increasing upstream competition in traditional retail, I randomize the entry of a wholesaler to corner stores across low- and middle-income neighborhoods in Lima.

The wholesaler is part of a modern retail conglomerate that also operates supermarket chains. It sells the same brands that incumbent suppliers distribute. Because incumbent suppliers price discriminate between the modern and traditional retail sectors, the group can negotiate bulk purchases at terms unavailable to corner stores and resell to them profitably. Its offer combines three features that distinguish it from incumbent suppliers.

Product mix. The entrant carries a broad portfolio spanning most of the categories and brands that incumbent suppliers distribute through separate networks. The products it sold during the study account for over 60% of the value of control stores' receipt-verified purchases from incumbent suppliers. Orders must meet a minimum total value—slightly above the minimums that incumbent suppliers typically impose—but can combine products freely across categories.

Wholesale prices. The group's purchasing terms allow the entrant to offer wholesale prices competitive with those of incumbent suppliers. Section 6.1.2 documents that realized entrant purchases carry lower wholesale prices than incumbent purchases of the same products.

Ordering convenience. Stores can order during weekly visits from the entrant's sales agents, whose incentive schemes resemble those of incumbent sales forces, or at any time through an app with on-demand delivery. Incumbent ordering is instead almost entirely face to face: only 6% of supplier orders placed by control stores are made without a salesperson at the store, compared with 36% of entrant orders among treated stores. Within an order, products can be bought in any quantity, without product-level package requirements that

are the norm in this sector.

In combination, these features offer stores consolidated sourcing across categories through a single relationship, and they place the entrant in direct competition with incumbent suppliers for the same purchases.

3.2 Experimental design

3.2.1 Entrant's coverage strategy

The study takes place in an area of northern Lima comprising eight districts and home to about 2.5 million people, over one quarter of the city's population. Income levels range widely across the area: about 38% of the study-area population lives in poor or very poor blocks, 39% in lower-middle-class blocks, and 23% in upper-middle-class blocks (as classified by Peru's National Statistics Institute—INEI). Panel A of Figure B7 shows the study area, with blocks colored by socioeconomic status (SES).

The entrant reaches corner stores through a sales force that visits them in person regularly. Using a 2018–2019 corner store census conducted by INEI as a reference, it divided the city into territories of about 200 stores each—contiguous blocks small enough for one sales agent to cover with weekly visits. These territories were drawn idiosyncratically, often following main roads. The study area comprises 119 territories in total. An additional 24, where the entrant also operates, were excluded from the study, either because their wealthier neighborhoods have different retail dynamics or because they were unsafe for enumerators (Appendix D.1). Panel B of Figure B7 shows the division, along with all corner stores in the area, from the INEI census. By combining this with a validation exercise, I estimate that the total number of stores was roughly 20,700 during the study period.[20]

3.2.2 Randomization and treatment

These territories served as the randomization strata. Within each stratum, I manually drew the clusters that form the units of randomization. I refer to these clusters informally as neighborhoods. First, all urban markets and their immediate surroundings—typically dense commercial blocks—were set aside for the entrant to cover freely from the initial rollout. These areas host a different mix of businesses and face different competitive conditions from ordinary residential areas, and therefore fall outside the scope of this study. The remainder of each stratum was split into three clusters, and one cluster per stratum was then randomly assigned to the control group.[21]

Sales agents visit treated clusters weekly from the rollout onward. They were not allowed to visit control clusters until data collection ended, so stores there received no active outreach,

[20] Enumerators conducted a full store listing exercise in three territories in a different part of the city, corresponding to the pilot area. The listing counted about 19% fewer stores than the census recorded, and a non-trivial share of corner stores found by enumerators did not match a point in the census data. The gap appears to stem mainly from churn—especially around the COVID pandemic, which began shortly after the census—and from imprecise location markers.

[21] All but seven strata had at least one market. In the seven strata that did not, four clusters were formed instead and three assigned to treatment. This was agreed with the entrant to leave sales agents a minimum number of potential clients in all strata. Market locations were provided by the entrant.

though they could still self-enroll through the app. Altogether, the design comprises 364 clusters across 119 strata, with one cluster per stratum in control.

Panels A and B of Figure 2 show the study clusters and treatment assignment, respectively. The experimental design drew on lessons from a pilot conducted earlier in a different part of the city. Clusters were made as compact and as similar in size as possible, and were shaped so that the treatment area would remain contiguous under any possible assignment—letting a sales agent navigate it without crossing control areas. Cluster boundaries were kept off streets where possible, so that treated and control stores would not end up directly across from one another. This was meant to limit direct exposure between the two arms and reduce both fairness concerns and the risk of navigation errors.

Balance. Tables A2 and A3 compare treated and control units on cluster characteristics—store counts, geography, population, and socioeconomic composition—and on predetermined store and owner characteristics obtained from the survey. Differences are small throughout. Joint orthogonality tests do not reject balance at either level ($p = 0.256$ for cluster characteristics; $p = 0.908$ for local characteristics and $p = 0.407$ for store and owner characteristics).

3.3 Implementation

The entrant agreed to hold out part of each stratum in order to concentrate early on smaller areas, ensuring coverage and giving it time to refine its sales strategy. The corner store survey took place nine to thirteen months after the rollout (Section 4).

During the survey months, the entrant's administrative records show 6,788 distinct client stores placing 76,284 orders in the study area.[22] Take-up grew through the first months of the rollout and leveled off around the start of the survey window: Figure 3 shows entrant clients with any purchase in each three-week period, and Appendix Figure B8 reports the corresponding cumulative shares. Many stores never tried the entrant, but use among those that did was persistent: of the 7,617 stores that ever placed an order, 89% remained active during the survey months, placing about two orders per month. I explore take-up patterns in detail, including the experimental first stage, in Section 6.1.1.

4 Data and outcomes

The main data source is an original in-person survey of corner stores that combines a three-week wholesale diary with modules on store characteristics, supplier relationships, retail availability and prices, and store performance.

4.1 Corner store survey

The corner store survey was fielded between March and August 2025, nine to thirteen months after the rollout. Survey capacity could not cover all strata at once, so they were surveyed

[22] The entrant's administrative records cover all client orders throughout the rollout, including those from stores outside the survey sample. They record order dates and delivery locations, which I use to assign clients to study clusters.

in monthly rounds, randomly assigned so that survey timing is uncorrelated with stratum characteristics. 77% of stores were surveyed by the end of May, in the first three rounds (Appendix Table D1).

4.1.1 Sampling and field protocol

The target sample was ten stores in each of the 119 strata in the study area, totaling 1,190: four from the control cluster and six from the treatment clusters, thus oversampling control areas relative to their share of clusters. The sampling frame was the same corner store census introduced in Section 3.2.1. The protocol used census locations to start searches, not to select stores: enumerators approached the nearest eligible store and surveyed it if open and its owner agreed to participate. Failed attempts triggered a move to the next-nearest store until a survey was completed or the search was abandoned. A handful of unresolved searches left a realized sample of 1,188 stores, just short of the target. Of these, five lie outside the study area and are dropped, leaving 1,183 stores for the analysis. The full sampling, eligibility, and search protocol is in Appendix D.

4.1.2 Visit structure and the wholesale diary

Each store was surveyed over two visits, two weeks apart. The first visit collected owner demographics, store characteristics and operations, self-reported performance, and wholesale purchases made during the preceding week. At the end, enumerators asked respondents to keep their receipts for a return visit exactly two weeks later. The second visit recorded purchases made during those two weeks, repeated the financial questions described below, and administered the retail audit. In some stores, the owner could not be found or declined the second visit, so 2.5% of stores have only one visit. Outcomes that aggregate across visits use their single response. Appendix D provides further detail on respondent payments and receipt collection.

The survey's central component is the resulting wholesale diary: every wholesale purchase made by each store over a three-week window, spanning the week before the first visit and the two weeks between visits. Purchases are recorded as orders and transactions. An order is a single purchase occasion from one source—a supplier sale, a market trip, or a purchase at a modern store. A transaction is a product within an order. The diary covers more than 14,000 supplier orders and 43,500 transactions with wholesale prices and quantities, from both the entrant and incumbent suppliers. Every reported order records its source and total expenditure. 71% of supplier orders are verified against receipts, which provide product-level detail and are digitized line by line. Appendix Table D2 describes receipt coverage, and Appendix E documents digitization and product harmonization.

4.2 Outcomes

The outcomes combine the wholesale diary with the store and supplier-relationship modules and the retail audit. Flow outcomes—expenditure, orders, and visits—are expressed as weekly averages over each store's diary window. Counts of distinct products, categories, brands, and suppliers instead cover the full window.

Store characteristics. The survey records predetermined store and owner characteristics: floor area, store age, and owner demographics. These variables enter the balance tests of Section 3.2.2, and floor area defines the store-size split introduced in Section 5.2.

Wholesale diary outcomes. From reported orders, I construct entrant take-up, order counts, and weekly wholesale expenditure in total and separately for the entrant, incumbent suppliers, and purchases at wholesale markets and modern stores. I also report receipt-verified counterparts. Digitized transactions provide product variety and within-product wholesale prices and quantities. A product is a brand–size–flavor combination, such as a personal-size bottle of Coca-Cola Light.[23] The catalog is also linked to control-group measures of median package price, store share, and supplier concentration used in the product-composition analyses. Appendix E provides the construction details.

Supplier relationships. Orders come from the wholesale diary, while visits without a purchase are reported at the store level. Together they provide total visits, orders per visit, and expenditure per order. Trade credit is reported for each supplier and collapsed to whether any incumbent extends trade credit to the store. Free goods are identified from free-product transactions and collapsed to a store indicator, while supplier-provided shelves and coolers are reported directly at the store level.

Retail audit. The second visit records availability and retail prices for six popular products, each the leading product in its category. The retail price index is the equal-weight mean across available products of log retail price minus that product's mean log price among control stores. Complementary outcomes record the shares priced below and above the product's modal price among control stores. Appendix D describes the implementation of this audit.

Store operations and performance. Self-reported total weekly purchases and sales of the packaged products covered by the study, fresh-product sales, household consumption of store products, other sources of store income (such as banking agents), and waste were reported at both visits and are averaged across them. All income adds fresh-product sales, household consumption, and other income to packaged-product sales. Customer counts, opening hours, monthly operating costs, and store-owned assets were measured at the first visit. The self-reported spending measure is distinct from wholesale purchases calculated order by order in the diary, allowing the two to be compared directly.

5 Conceptual framework and empirical strategy

To guide the empirical analysis, I first outline a conceptual framework that summarizes the main elements and predictions of the theoretical model developed in Section 7. I then present the empirical specifications used throughout the analysis.

[23] Pooling flavors within brand and size leaves the results unchanged. The coarser definition provides more observations per product but can combine flavors sold at different wholesale prices and requires judgment about which flavors should be grouped.

5.1 Conceptual framework

The model considers a neighborhood in which corner stores sell a single product and differ in their location and in their handling costs. Store size serves as a proxy for these costs: large stores are more efficient, as they are less constrained by space (Section 2.2). Consistent with the concentrated supply chains and uniform wholesale prices described in Section 2.3, a monopolist supplier chooses a wholesale price and which stores to serve, paying a fixed delivery cost per store. Stores that are served then decide which retail price to charge, and consumers choose whether and where to buy the product by trading off retail prices against travel costs.

Prediction 1: Exclusion of less efficient stores. *Incumbent suppliers may choose to exclude relatively inefficient stores from their distribution networks to save on logistics costs, while still covering the market.*

Travel costs create an incentive for the supplier to serve more stores: this brings the product closer to consumers, creating additional surplus that it can extract by raising wholesale prices.[24] However, additional gross profits from serving the less efficient stores may not cover the associated delivery costs. The supplier may therefore prefer to exclude those stores from its distribution network, even if they would be willing to buy the product at the prevailing wholesale price. In the model, excluded stores have no other way to source the product before the wholesaler's entry. Large stores can still serve the entire neighborhood, so the supplier does not lose any sales by excluding small stores. This prediction is consistent with the patterns documented in Section 2.2, where small stores buy from fewer suppliers and purchase fewer distinct products than large stores. It also implies that an outside option should particularly benefit small stores, allowing them to expand their product variety.

Prediction 2: Incumbent responses to the wholesaler's entry. *After the wholesaler's entry, incumbent suppliers may respond by serving more stores or by improving the terms they offer their existing clients.*

To isolate the effect of a new outside option for excluded stores, I assume that an entrant offers the product at the same wholesale price as the incumbent. After the wholesaler's entry, small stores can carry the product regardless of whether the incumbent serves them, and thus compete with large stores for consumers. The model predicts that, if additional delivery costs are not too high, the incumbent may now find it worthwhile to serve small stores, capturing purchases that would instead go to the entrant. Otherwise, it keeps excluding them and may lower the wholesale price it charges large stores to defend their market share (or exit if delivery costs are high enough). Either way, the incumbent earns less than as a monopolist: it incurs an additional delivery cost in the first case, and gives up volume or margin in the second.

Because suppliers in this setting rarely price discriminate across corner stores and instead manage client relationships through non-price terms (Section 2.3), wholesale prices in the model may reflect both price and non-price relationship terms. In a single-product model,

[24] This is similar to the trade-off an integrated manufacturer faces in spatial models of distribution (Dixit, 1983; Mathewson and Winter, 1983). Here, the fixed cost is the cost of supplying an existing store rather than opening a retail outlet.

at most one response arises for given parameters. In practice, differences in store locations and costs, supplier cost structures, and product valuations may generate both responses simultaneously, within and across suppliers.

Once small stores carry the product, whether from the incumbent or the entrant, demand reallocates toward them and consumers walk less on average, unless better terms allow large stores to keep the entire market. In the model, all consumers already buy the product before the wholesaler's entry, so small stores grow at the expense of large stores, leaving total sales in the neighborhood unchanged.

5.2 Empirical strategy

All experimental estimates are intent-to-treat effects of the cluster-level rollout. For store-level outcomes, the main specification is

$$Y_{ics} = \alpha + \beta\, T_c + \gamma_s + \varepsilon_{ics}, \tag{1}$$

where Y_{ics} is the outcome of store i in cluster c and stratum s, T_c indicates treated clusters, and γ_s are stratum fixed effects. Standard errors are clustered at the stratum level rather than the cluster level, following de Chaisemartin and Ramirez-Cuellar (2024), to account for correlation in treatment assignment within strata.

Results are also reported separately by store size throughout the analysis. Store size is a good proxy for space constraints and operating efficiency. It is also time-invariant, and can therefore be used safely for heterogeneity analysis in empirical specifications. I classify stores as large when their floor area is above the control-group median and as small otherwise. Table A3 shows that treated and control stores do not differ in this dimension, as expected. The heterogeneity specification interacts both treatment and the stratum fixed effects with an indicator L_i for large stores, as follows:

$$Y_{ics} = \beta_1\, T_c + \beta_2\, T_c \times L_i + \gamma_{s\times L} + \varepsilon_{ics}, \tag{2}$$

where $\gamma_{s\times L}$ are stratum-by-size fixed effects, so that β_1 is the effect among small stores and $\beta_1 + \beta_2$ the effect among large stores. Tables report the p-value of the test that the two are equal.

Some analyses use the wholesale diary's transactions as the unit of observation, where a transaction is an order–product combination. These specifications add product fixed effects to the regressors in equations (1) and (2):

$$Y_{pics} = \beta\, T_c + \gamma_s + \delta_p + \varepsilon_{pics}, \tag{3}$$

where δ_p are product fixed effects, so that outcomes such as wholesale prices and quantities are compared within product. In the heterogeneity version, treatment, the stratum fixed effects, and the product fixed effects are all interacted with L_i. The exceptions are specifications that quantify changes in product composition (Appendix Table A4): their outcomes are fixed product characteristics that product fixed effects would absorb, so they include only the fixed effects of equations (1) and (2).

Finally, Section 6.1 analyzes how stores use the entrant by comparing treated-area stores' orders and transactions with the entrant against those with incumbent suppliers. The specification is

$$Y_j = \alpha + \beta\, E_j + \phi_{g(j)} + \varepsilon_j, \tag{4}$$

where j indexes an order or transaction, E_j indicates purchases from the entrant, and β is the entrant–incumbent difference. The fixed-effect group $g(j)$ is the stratum in order-level versions. Transaction-level versions add product fixed effects, so wholesale prices and quantities are again compared within product. In heterogeneity versions, the entrant indicator and the fixed-effect groups are interacted with L_i as in equation (2). Because stores choose whether and what to buy from the entrant, these comparisons are descriptive rather than experimental.

6 Results

This section presents the experimental results. Section 6.1 shows that the rollout raises demand for the entrant substantially and describes how stores use it. The next three subsections present the main findings: variety gains concentrated among small stores, together with better product availability for consumers (Section 6.2); higher wholesale purchases among small stores, most of them from incumbent suppliers (Section 6.3); and incumbent responses through non-price terms rather than wholesale prices (Section 6.4).

6.1 Entrant take-up and use patterns

I first examine whether the entrant's outreach increases take-up. I then describe how stores use the entrant and compare its terms with those of incumbent suppliers.

6.1.1 The randomized rollout increases entrant take-up

Table 1 reports treatment effects on entrant take-up over the three-week wholesale diary. Stores in treated areas are 10.8 percentage points more likely to place at least one order with the entrant ($p < 0.01$), roughly doubling the control-group share of 11.7%. The first stages for weekly entrant expenditure and orders are similarly large (both $p < 0.01$). Table A5 reports similar receipt-verified counterparts, with a less precise expenditure estimate ($p < 0.10$). On the intensive margin, usage is high: entrant clients in treated areas source roughly 20% of their wholesale expenditure through it.

The effect is stronger in the entrant's administrative client records, matched by cluster and date to the survey periods (columns (4)–(6)). In these records, 8.0% of stores in control areas have a delivered entrant order. Treatment raises this share by 16.7 percentage points ($p < 0.01$), to roughly three times the control-group share. Administrative expenditure and order counts also show large increases (both $p < 0.01$). Not every regular client places an order within a given three-week window. Extrapolating from clients' observed order frequencies under a Poisson benchmark, the share of stores with an ongoing entrant relationship is above 30% in treated areas and below 10% in control areas.[25] A direct comparison suggests

[25] A client ordering at the arm's average rate per three-week window places at least one order in a given

that the survey overstates control-group purchasing from the entrant, so I conclude that the survey-based first stage is conservative. Take-up had also leveled off around the start of the survey (Figure 3, Section 3.3), so the effects below capture a settled rather than transitory response.

The rollout increases take-up among both small and large stores. The effects are 12.6 percentage points for small stores ($p < 0.01$) and 7.9 for large stores ($p < 0.10$), with no detectable difference between them ($p = 0.355$). The proportional increase is about three times as large among small stores because their control-group take-up is lower, at 8.1% compared with 16.3%. More generally, I find little evidence that the response to outreach varies systematically with observed store characteristics or local conditions. Table A6 shows that the treatment interactions with these characteristics are not jointly significant ($p = 0.280$ and 0.238, respectively).

6.1.2 The entrant is an attractive sourcing option across many categories

I now compare how treated-area stores use the entrant and incumbent suppliers. These comparisons are descriptive rather than experimental, since stores freely choose whether and what to buy from the entrant, but they show what the new channel offers. Table A7 shows that entrant orders are nearly twice as large in value as incumbent orders (89%, $p < 0.01$), and the gap is larger over receipt-verified purchases. They also contain 23% more products ($p = 0.029$) and span 50% more categories ($p < 0.01$). Stores therefore use the entrant to buy across categories rather than to deepen the range within each, and the patterns are similar for small and large stores.

Within-product comparisons show that the entrant is also competitive on wholesale prices. Table A8 shows that entrant transactions carry 5% lower wholesale prices than incumbent transactions for the same product ($p < 0.01$), with no detectable difference between small and large stores and no difference in the number of packages per transaction. This gap partly reflects selection: stores buy from the entrant precisely where it is attractive, so the comparison describes realized savings rather than a full-catalog comparison of wholesale prices. Clients also take advantage of flexible ordering, with 36% of entrant orders placed remotely, compared with 6% of incumbent-supplier orders in control areas.

Overall, the rollout raises entrant take-up substantially across store sizes, and stores find in the entrant an attractive alternative to incumbents across many categories, at competitive wholesale prices. The wholesaler's entry is therefore a competition shock to incumbent suppliers. Sections 6.2–6.4 evaluate its effects on what stores stock and buy, and on how incumbents respond.

6.2 Product variety and consumer access to products

First, I ask whether access to the entrant changes the products stores carry, and what this implies for the products consumers find nearby.

window with probability one minus the exponential of minus that rate. Dividing the share of stores with a delivered order in the administrative records by this probability gives the implied share of stores with an ongoing relationship.

6.2.1 Variety gains are concentrated among small stores

The wholesaler's entry expands product variety primarily among small stores. Table 2 shows that small stores in treated neighborhoods purchase 5.5 more distinct products over the diary window than those in control neighborhoods, an increase of 19% ($p = 0.037$). The large-store estimate is also positive but indistinguishable from zero. Across all stores, the effect is 9% and imprecise in levels ($p = 0.229$), but significant in the log specification in Table A9 (14%, $p = 0.041$). Figure B9 shows this pattern without a median split: treated stores purchase more products throughout the lower half of the size distribution, with a treatment–control gap indistinguishable from zero above the median.[26]

The expansion comes from breadth rather than depth. Small stores buy from 1.3 more categories, 13% of the control mean ($p = 0.050$), and 7% more brands per category ($p = 0.031$), with no change in products per brand. In the log specification in Table A9, where the three margins add up to the total, categories account for 19 of the 23 log points of the small-store effect and brands per category for 7, with products per brand slightly negative. Small stores are therefore not simply adding close substitutes of products they already carry. They broaden the range of products available locally.

The composition of purchases helps explain why the gains concentrate among small stores. Small-store purchases shift toward products distributed through more concentrated incumbent channels: the average supplier concentration (HHI) of the products they buy rises 3.6% ($p = 0.043$; Table A4). When a product's distribution is concentrated among one or few suppliers, carrying it is more likely to require a separate relationship and its associated fixed costs. Together, these patterns are consistent with incumbent suppliers excluding small, less efficient stores from their distribution networks before the wholesaler's entry.

6.2.2 Product availability improves, with limited retail-price responses

The wholesaler's entry also improves the availability of products that stores commonly carry. The retail audit covers six popular products, each the top product in its category and stocked by most control stores.[27] These measures capture availability at the time of the visit, complementing the expansion in product variety documented above. Table 3 shows that treatment reduces the probability that a store is missing at least one of these products by 5.8 percentage points, from a control mean of 33.7% ($p = 0.029$). The estimated reduction is 9.1 percentage points among large stores ($p = 0.016$) and 5.0 among small stores ($p = 0.215$), although the two effects do not differ significantly ($p = 0.449$).

Turning to retail prices, Table 3 shows no detectable change in the retail price index in the full sample or among small stores ($p = 0.165$ for the latter). Among large stores, however, the index declines by 0.6% ($p = 0.019$), with a significant difference between the two size groups ($p = 0.008$). Columns (4) and (5) also show a shift from above-modal to below-modal retail prices among large stores (both $p < 0.10$), while among small stores the share priced below the mode falls ($p < 0.05$) without a detectable change in the index. These movements are notable because many of these products are subject to resale price

[26] The continuous size interactions in Tables 2 and A9 also indicate declining effects with floor area, with $p = 0.119$ in levels and $p = 0.022$ in logs.

[27] Rice was also surveyed but is excluded because almost no stores carry the reference brand.

maintenance by incumbent suppliers: 64% of control-store retail prices equal the product's modal price across the study area. The reduction at large stores could reflect competitive pressure from better-stocked small stores, or lower storage costs from the equipment that suppliers provide to large stores in treated areas (Section 6.4.1). This equipment may also contribute to their lower stockout probabilities.

Appendix Tables A10 and A11 provide product-specific estimates of availability and retail prices, which are noisy but suggest that the results are not driven by any particular product.

6.3 Wholesale expenditure and store scale

Second, I ask whether broader variety is accompanied by additional wholesale purchases, and whether those purchases come at the expense of nearby stores.

6.3.1 Small stores expand wholesale purchases, mostly from incumbents

Small stores expand their wholesale purchases alongside their gains in product variety. Table 4 shows that their weekly expenditure rises by 46 USD PPP, or 20% of the control mean ($p = 0.001$), while the large-store estimate is negative and imprecise. Across all stores, the effect is positive but imprecise in levels, at 4% of the control mean ($p = 0.531$). The log estimates in Table A12 show a significant increase of 10% across all stores ($p = 0.041$) and 22% among small stores ($p = 0.002$). The concentration of gains among small stores also extends beyond the median split: the estimated treatment effect declines by 72 USD PPP per standard deviation of floor area ($p = 0.020$). Figure 4 shows that treated stores spend more than control stores throughout the lower half of the size distribution, while the treatment–control gap is indistinguishable from zero above it, flattening the relationship between store size and wholesale expenditure.

Incumbent suppliers account for over three quarters of small stores' additional wholesale expenditure. In columns (2) and (3) of Table 4, purchases from the entrant rise 8.7 USD PPP per week ($p < 0.01$), while purchases from incumbent suppliers rise 37 USD PPP ($p < 0.01$). The same holds over receipt-verified orders (columns (5) and (6); $p < 0.01$ and $p < 0.10$). Access to the entrant therefore crowds in purchases from incumbent suppliers rather than merely replacing their sales. Among large stores, the entrant estimate is positive and the incumbent estimate negative, a pattern consistent with some substitution toward the entrant given the take-up results, but both are imprecise. There is no detectable change in expenditure at wholesale markets and modern stores for either group (column (4)). These channels account for a relatively small fraction of all purchases.

Small stores' expansion occurs without detectable reductions in their retail prices (Section 6.2.2), consistent with consumers responding to variety gains rather than to pass-through of improved wholesale terms.

6.3.2 Effects on self-reported measures of store performance

Tables A13 and A14 show imprecise effects on reported wholesale purchases, sales, and customer counts, despite the substantial increase in small stores' purchases documented in Section 6.3.1. Appendix G examines whether reporting errors can explain the lack of

statistically significant effects on these outcomes. I take advantage of the survey's two parallel measures of wholesale purchases: the order-by-order wholesale diary and owners' direct reports. A direct comparison finds large reporting discrepancies across stores. While observable characteristics are associated with these discrepancies, they explain a small share of their variation (Table G3). Restricting the sample to stores whose self-reports agree more closely with the diary produces larger small-store sales effects (though still not significant at conventional levels).[28]

Appendix G also presents power calculations suggesting that sales estimates are too noisy to reliably detect the gain implied by the diary, along with other evidence of substantial reporting inaccuracies. Temporary inventory accumulation is also unlikely to explain the purchase response: the survey took place nine to thirteen months after rollout, once take-up had leveled off, and there is no significant increase in capacity utilization or waste (Table G1).

Other outcomes in Table A13 provide some evidence of increased business activity. Stores remain open 0.2 additional hours per day on average ($p = 0.065$), with similar estimates for small and large stores. Monthly operating costs—which consist largely of utility bills—rise among large stores by 45 USD PPP ($p = 0.043$), with small-store estimates similar in size but imprecise. While assets appear to decrease among large stores, this effect does not account for the value of shelves loaned by suppliers, whose provision increases in treated areas as discussed in Section 6.4.1.

6.3.3 Spillovers to nearby stores and the size of the market

As discussed in Section 6.3.1, the positive but imprecise level effect on wholesale expenditure across all stores leaves open whether small-store growth comes at the expense of other stores. Such interactions between stores may also extend across cluster boundaries, since treated and control clusters are geographically contiguous. While the literature on causal inference with experiments has generally treated such interference as a threat to identification, studying spillovers can also reveal the market interactions underlying treatment–control differences (Casaburi and Reed, 2022). To disentangle direct treatment effects from potential spillovers, I use variation in the treatment of nearby clusters to estimate the effects of exposure to treated stores, following the recentered-instrument approach proposed by Borusyak and Hull (2023).

I assume that spillovers dissipate with distance and measure exposure as the number of treated stores within 450 meters of each focal store, including those in its own cluster (varying this threshold gives qualitatively similar results, though less precise). Because stores in the same cluster share treatment status, their own access to the entrant changes together with the treatment of their neighbors. To separate these effects, I assume that exposure to treated neighbors has the same effect on a store's outcomes whether those neighbors are inside or outside its cluster. This allows me to infer within-cluster spillovers from variation in the treatment of neighboring clusters.[29] I leverage the randomization to

[28] These comparisons are suggestive because treatment may affect inclusion in the restricted samples. However, treatment effects on the main diary-based measures of expenditure and variety remain qualitatively similar, alleviating this concern.

[29] This is a strong assumption. Within strata, I drew similarly sized clusters and kept their boundaries off streets where possible to reduce direct exposure and compliance problems. Strata, on the other hand, were drawn by the wholesaler's operations team, often following main roads or district boundaries that could limit

calculate expected exposure for each focal store, conditional on its own treatment status to account for correlated treatment assignment within strata. Deviations from this expectation provide an instrument for exposure. Exposure effects for control stores are identified only from neighboring strata, because all other clusters in their own stratum are treated by design. Appendix F provides details of the procedure and additional results.

Panel B of Table A15 separates the effect of own treatment, holding exposure fixed, from spillovers due to treatment of nearby stores. For comparison with the cluster-level estimates reported elsewhere, the table also combines direct and spillover effects into an implied cluster effect: the average change at a store when its whole cluster is treated, holding exposure outside the cluster fixed. For small stores, own treatment raises weekly wholesale expenditure by 82 USD PPP ($p = 0.002$) and the number of products purchased by 13.1 ($p = 0.011$). Greater exposure works in the opposite direction: an increase equal to one standard deviation of the instrument (12.7 treated stores) reduces expenditure by 14 USD PPP ($p = 0.046$) and the number of products purchased by three ($p = 0.058$). Exposure effects for large stores are imprecisely estimated, as are the differences between the two size groups, so the estimates do not establish that large stores are unaffected.

A comparison of Panels A and B indicates that accounting for exposure reduces the estimated effects of cluster treatment on small stores' weekly wholesale expenditure from 46 to 33 USD PPP ($p = 0.019$) and on product variety from 5.5 to 2.4 products ($p = 0.436$). The expenditure effect remains statistically significant, while the product-variety effect remains positive but is no longer significant at conventional levels. This suggests that the treatment–control differences estimated without accounting for exposure partly reflect losses at control stores, which have more treated neighbors outside their own cluster by design.[30]

Overall, the analysis provides suggestive evidence of spillovers between nearby stores. With this design, I cannot distinguish whether these arise from direct competition for consumers or from suppliers reallocating their service across stores. Yet both mechanisms are consistent with the market interactions captured in the conceptual framework (Section 5.1). Accounting for spillovers also leaves positive own-cluster effects for small stores, so their gains are not explained entirely by losses at nearby control stores. However, the full-sample level estimates remain too imprecise to determine whether the wholesaler's entry expands total market activity or simply reallocates it across stores.

6.4 Incumbent supplier responses

Third, I ask how incumbent suppliers respond to the wholesaler's entry.

interactions between stores.

[30] Panel C estimates the own-cluster effect directly by counting only treated neighbors outside the cluster in exposure. Its treatment coefficient combines own access and within-cluster spillovers without requiring these spillovers to equal those across cluster boundaries. Both specifications use the same variation in treatment of outside neighbors and estimate comparable own-cluster effects, so their results should be closely related. The estimates are indeed similar: 34 USD PPP ($p = 0.013$) and 2.8 products ($p = 0.336$).

6.4.1 Incumbents extend more visits and non-price benefits to small stores and more equipment to large ones

Small stores' additional purchases from incumbent suppliers reflect both more frequent visits and more purchasing during those visits. Table 5 shows that supplier visits to small stores rise by 7% (p = 0.079).[31] Small stores also place orders on a larger share of those visits, with orders per visit rising by 5% (p = 0.062). Expenditure per order also rises by 9%, although this estimate is imprecise (p = 0.104). In the log specification in Table A16, the three margins account for roughly equal shares of the increase in incumbent expenditure. Together, these patterns are consistent with incumbents accommodating growing demand from small stores by visiting them more frequently. Among large stores, I find no detectable changes in these margins.

Incumbents also extend volume-contingent benefits to small stores. Table 6 shows that small stores in treated areas are 8.5 percentage points more likely to receive free goods, from a control mean of 48.4% (p = 0.038), and 5.3 percentage points more likely to receive trade credit, from 21.0% (p = 0.086), terms typically reserved for high-volume clients. At the store–supplier level, Appendix Table A17 shows that the share of purchasing relationships that carry free goods or trade credit rises for small stores as well. Trade credit to large stores falls by 6.5 percentage points (p = 0.060), but this effect is not robust to the store–supplier specification in Appendix Table A17.

Among large stores, incumbents respond by increasing equipment transfers instead. Suppliers lend 0.29 additional display shelves (p = 0.035) and 0.36 additional coolers (p = 0.004) to large stores in treated areas, about one added unit of each for every three large stores, and the shelves effect rises with floor area (p = 0.016). Shelves and coolers expand a store's storage capacity and lower its storage costs, and suppliers may also use them to restrict the products a client carries. These effects are therefore consistent with suppliers making costly efforts to retain their most valuable clients, and to help those clients defend their market share amid growing competition from nearby stores.

6.4.2 No evidence of changes in incumbent wholesale prices or quantities

In contrast to the non-price margins, Table 7 shows no detectable change in incumbent wholesale prices in treated neighborhoods. The effect is a precise null: the coefficient is -0.1% with a standard error of 0.1%, so the 95% confidence interval rules out reductions larger than 0.3% and increases larger than 0.1%. Including entrant purchases in column (3) leaves the estimate indistinguishable from zero (-0.1%, p = 0.287). This magnitude is close to the mechanical effect implied by the entrant's 5% wholesale-price advantage and its higher share of transactions in treated areas if incumbent wholesale prices are unchanged.[32]

[31] The sample excludes stores without an incumbent order. Total visits combine orders recorded in the wholesale diary with retrospectively reported visits that did not result in a purchase. The latter may be subject to recall bias and other measurement error.

[32] The number of packages purchased per transaction shows no detectable change (-1.1%, p = 0.648), and most purchases are already of one package. Neither the wholesale-price nor quantity effects differ significantly between small and large stores. The continuous size interaction for wholesale prices is statistically significant but small, at 0.2% per standard deviation of floor area (p = 0.004).

A likely explanation is that suppliers set wholesale prices centrally for geographic areas broader than the experiment's neighborhoods and, as discussed in Section 2.3, rarely price discriminate across corner stores.[33] I conjecture that the rigidity of incumbent suppliers' pricing policies in this setting makes non-price terms the more suitable margin for responding to localized competition. Section 7 rationalizes these inclusion and retention responses in a model of supplier distribution networks.

7 Theoretical model

The empirical results show that increased competitive pressure from the entry of a wholesaler can induce incumbent suppliers to (1) sell to small stores that they might otherwise not serve, and (2) improve the terms they offer large stores, specifically by transferring storage equipment that arguably lowers their costs. To formalize the mechanisms outlined in Section 5.1, I present a theoretical model of how a monopolist supplier chooses which stores to sell its product to, and what this means for consumer welfare.

The basis is a Hotelling environment, where stores are horizontally differentiated by location and differ in handling costs (one is small and one is large). A monopolist supplier chooses a wholesale price and which stores to serve, paying a fixed delivery cost. I first characterize its distribution network and show that, if the fixed cost is sufficiently high, it may prefer to concentrate sales among lower-cost stores.[34] I then examine how the supplier responds to an increase in upstream competition: an entrant offering the same product to any store at the incumbent's pre-entry wholesale price (throughout this section, pre-entry and post-entry refer to the wholesaler's entry).

After the wholesaler's entry, consumers benefit from lower travel costs, but retail prices may increase where more stores carry the product. I show that when travel costs are sufficiently high, the former effect dominates and upstream competition increases aggregate consumer welfare.

7.1 Setup

Two stores, S (small) and L (large), are located on opposite ends of a unit interval (a neighborhood), with a unit mass of consumers uniformly distributed along the line. A consumer at $x \in [0, 1]$ obtains utility

$$U_L(x) = \nu - p_L - \tau x,$$
$$U_S(x) = \nu - p_S - \tau(1 - x),$$

[33] Anecdotally, the entrant's sales agents report instances of incumbent suppliers lowering wholesale prices globally to match the entrant's more competitively priced products.

[34] Fixed delivery costs can justify excluding a store even under vertical integration: adding the store is profitable only if the increase in total producer surplus (from higher retail prices) covers the additional delivery cost. Conversely, models without fixed costs can also generate exclusion when suppliers cannot price-discriminate and the marginal store is inefficient enough. Intuitively, inefficient stores require lower wholesale prices to be viable, reducing the supplier's margin on all purchases.

where p_L and p_S are the retail prices, ν is the value of the product, and $\tau > 0$ is the cost per unit of distance. Consumers buy at most one unit and can choose not to purchase, obtaining utility zero.

A monopolist supplier has constant marginal cost $c \geq 0$ and charges the same wholesale price w to every store it supplies. It pays a fixed delivery cost $F \geq 0$ for each store that places positive orders. Its profits are therefore

$$\Pi = (w - c)(q_S + q_L) - F\mathbf{1}\{q_S > 0\} - F\mathbf{1}\{q_L > 0\},$$

where q_j is the quantity sold by each store $j \in \{S, L\}$.

Store j has marginal handling cost h_j, so its profits are

$$\pi_j = (p_j - w - h_j)q_j.$$

The timing is as follows:

1. The supplier chooses which stores to offer the product to, and sets a uniform wholesale price w.
2. Stores that receive an offer choose whether to purchase and simultaneously choose retail prices. Quantities are determined by consumer demand.
3. Consumers observe which stores carry the product and their retail prices and choose whether and where to purchase.

I study pure-strategy subgame-perfect equilibria: stores price optimally after observing the supplier's offer, and the supplier anticipates their responses. Appendix H provides detailed derivations of all results.

Assumption 1: Store heterogeneity. $0 \leq h_L < h_S < h_L + 2\tau$.

Large stores are more efficient at handling the product, but small stores remain efficient enough to attract some consumers if both stores carry it.

Assumption 2: Consumer valuation. $\nu \geq c + h_L + 4\tau$.

Consumer valuation is high enough that, when relying only on the large store, the supplier prefers to serve the full market rather than increase its wholesale margin and leave some consumers unserved.

7.2 Before the wholesaler's entry

The supplier's choice consists in comparing its best outcome with one active store and with two, net of delivery costs. I characterize each and derive a condition for exclusion of the small store to be preferred. Superscripts L and LS denote the outcomes with one and two stores, respectively.

One active store. Because $h_L < h_S$, L generates at least as much supplier profit as S for any given w, so in the relevant one-store outcome the supplier serves only the large store.

With only L active, the supplier sets the highest wholesale price consistent with full market coverage: below it, profits rise one for one with w; above it, L would leave the most distant consumers unserved, which Assumption 2 makes unprofitable for the supplier.

The optimal wholesale price is therefore

$$w^L = \nu - h_L - 2\tau.$$

At this wholesale price,

$$p_L^L = \nu - \tau, \qquad q_L^L = 1,$$

so supplier profit is

$$\Pi^L = w^L - c - F.$$

At w^L, the small store would profitably purchase the product if offered, so its absence from the supplier's network reflects active exclusion.

Two active stores. If both stores are supplied, the supplier again sets the highest wholesale price consistent with full market coverage. Because consumers are now closer to a store, this price is higher than w^L:

$$w^{LS} = \nu - \tau - \frac{h_L + h_S}{2}.$$

At this wholesale price,

$$q_L^{LS} = \frac{1}{2} + \frac{h_S - h_L}{4\tau}, \qquad q_S^{LS} = \frac{1}{2} - \frac{h_S - h_L}{4\tau},$$

so the two stores jointly cover the full market. Supplier profit is

$$\Pi^{LS} = w^{LS} - c - 2F.$$

Adding S reduces the distance consumers must travel and therefore allows the supplier to sustain a higher wholesale price while continuing to sell one unit.

Proposition 1: Equilibrium before the wholesaler's entry. *Define*

$$F^{pre} = w^{LS} - w^L = \tau - \frac{h_S - h_L}{2} > 0.$$

Under Assumptions 1 and 2, the supplier's distribution network depends on the delivery cost as follows:

(i) if $F < F^{pre}$, it serves both stores at w^{LS}, and the two stores cover the full market with quantities q_L^{LS} and q_S^{LS};

(ii) if $F^{pre} < F < w^L - c$, it serves only L at w^L, and L covers the full market, $q_L^L = 1$;

(iii) if $F > w^L - c$, it serves no store.

Assumption 2 implies $w^L - c \geq 2\tau > F^{pre}$, so the exclusion region (ii) is nonempty.

Because total quantity is one in both candidate outcomes (L and LS), F^{pre} equals the additional gross wholesale profit from serving both stores. It is therefore the largest delivery cost the supplier is willing to incur to serve the small store before the wholesaler's entry: above it, the additional surplus from serving S does not cover the extra delivery, and S is excluded even though it would buy at w^L. Above $w^L - c$, not even L is worth a delivery. Appendix H.2 proves Proposition 1.

Throughout, I assume that $F^{pre} < F < w^L - c$, so the supplier optimally excludes S in the pre-entry equilibrium.

Consumer welfare. Consumer surplus with only L and with both stores active is

$$CS^L = \frac{\tau}{2}, \qquad CS^{LS} = \frac{\tau}{4} + \frac{(h_S - h_L)^2}{16\tau}.$$

Assumption 1 implies $CS^L > CS^{LS}$, so exclusion of the small store benefits consumers. Intuitively, although adding S reduces consumers' travel distances, it also allows the supplier to charge a higher wholesale price while maintaining full market coverage. The resulting increase in retail prices more than offsets the travel savings in aggregate, though some consumers sufficiently close to S nevertheless benefit from the additional outlet.

The main takeaway is therefore that delivery costs limit the supplier's ability to extract consumer surplus by adding more retail outlets.

7.3 After the wholesaler's entry

Suppose an entrant offers the same product at the pre-entry wholesale price $w_E = w^L$. Stores purchase from the supplier offering the lowest wholesale price and break ties in favor of the incumbent.

Fixing the entrant's wholesale price at the pre-entry level isolates the effect of a new outside option for excluded stores, abstracting away from possibly lower wholesale prices. Even if the incumbent refuses to sell to S, it can no longer prevent S from carrying the product. Charging more than w_E loses both stores' purchases to the entrant, so w_E is effectively a ceiling on the incumbent's wholesale price.

Assumption 3: Retail competition. $h_S - h_L \leq \tau$.

The small store's handling-cost disadvantage is at most the travel cost τ.[35] *When both stores pay the wholesale price w^L, the full market is covered and the retail equilibrium is unique.*

The incumbent compares its profits from serving both stores with its best alternative when continuing to serve only L to determine its distribution network after the wholesaler's

[35]For $\tau < h_S - h_L < 2\tau$, the pre-entry results are unchanged. After the wholesaler's entry, one can show that there remains a nonempty range of delivery costs for which the incumbent serves both stores at w^L, the entrant receives no orders, and the full consumer market is covered. The retail pricing game can, however, admit multiple just-covered equilibria, so retail prices, market shares, and the exact boundary of the inclusion region depend on equilibrium selection.

entry. Below I derive the optimal wholesale price under each network. Superscripts L, E and LS, E denote outcomes conditional on the incumbent serving only L or both stores after the wholesaler's entry.

Both stores. At or below w_E, both stores buy from the incumbent. Since total sales are already one, lowering the wholesale price cannot increase quantity. The incumbent therefore sets $w^{LS,E} = w_E = w^L$ and earns

$$\Pi^{LS,E} = w^L - c - 2F.$$

Only L. If the incumbent continues to exclude S, the latter can buy from the entrant at w^L and compete for consumers. At the original wholesale price, L loses market share and the incumbent cedes the corresponding wholesale sales to the entrant. The incumbent can defend these sales by lowering its wholesale price to L, but doing so reduces its margin on all units sold.

Lemma 1: Wholesale price when the incumbent continues to exclude S. *Suppose that after the wholesaler's entry the incumbent serves only L, so that S can buy from the entrant at w^L. Under Assumptions 1–3, the incumbent's optimal wholesale price depends on its pre-entry wholesale margin $w^L - c$:*

(i) if $w^L - c \leq 3\tau + (h_S - h_L)$, it keeps w^L and cedes part of L's market share to S;

(ii) if $3\tau + (h_S - h_L) < w^L - c < 9\tau - (h_S - h_L)$, it sets $w^{L,E} = c + \frac{3\tau + (w^L - c) + (h_S - h_L)}{2} < w^L$, and L recovers part of its market share;

(iii) if $w^L - c \geq 9\tau - (h_S - h_L)$, it sets $w^{L,E} = w^L + (h_S - h_L) - 3\tau$, the highest wholesale price at which L serves the entire market.

Its profit in this case is $\Pi^{L,E}$.

The incumbent's wholesale margin determines the marginal profits from expanding the large store's market share. At low margins each regained unit is worth little, so the incumbent accepts the loss of market share. At higher margins the regained sales are worth more, and the incumbent lowers the wholesale price until the sacrificed margin balances the regained sales, or until L serves the entire market. Because $w^L - c = \nu - c - h_L - 2\tau$ rises with the product's value, the incumbent tends to cede market share on low-value products and defend it on high-value ones.

Appendix H.4 derives the optimal wholesale price and quantities. In every case, the incumbent's gross profits are lower than before the wholesaler's entry: L sells less than the full market at the same margin, or at most the full market at a lower margin.

Proposition 2: Equilibrium after the wholesaler's entry. *Let F^{post} denote the incumbent's gain in gross profits from serving both stores at w^L rather than only L at the wholesale price determined by Lemma 1, so that*

$$\Pi^{LS,E} - \Pi^{L,E} = F^{post} - F.$$

Under Assumptions 1–3, $F^{post} > F^{pre}$. Then:

(i) if $F < F^{post}$, the incumbent serves both stores at $w^{LS,E} = w^L$, the entrant receives no orders, and the full market is covered with $q_L^{LS,E} = \frac{1}{2} + \frac{h_S - h_L}{6\tau}$ and $q_S^{LS,E} = \frac{1}{2} - \frac{h_S - h_L}{6\tau}$;

(ii) if $F > F^{post}$, the incumbent continues to exclude S, which can buy from the entrant, and sets its wholesale price to L as in Lemma 1 (it exits once F exceeds its gross profits from serving L alone).

In particular, for $F^{pre} < F < F^{post}$ the incumbent excludes S before the wholesaler's entry but serves both stores after it.

F^{post} exceeds F^{pre} because excluding S no longer keeps it out of the retail market. Before the wholesaler's entry, the gain from serving S is only the increase in the wholesale margin that shorter travel distances allow. After the wholesaler's entry, the gain is the margin on the sales S would take from L, or the gross profits the incumbent would sacrifice to defend them (Lemma 1). Appendix H.4 proves Proposition 2.

The wholesaler's entry lowers the incumbent's profits whichever network it chooses. If it serves both stores, profits are $\Pi^{LS,E} = w^L - c - 2F$, exactly F below the pre-entry level: total sales and the wholesale margin are unchanged, but an additional delivery is required. If it continues to exclude S, it either cedes sales to the entrant or sacrifices margin to defend L's market share (Lemma 1).

These two responses have empirical counterparts in the results. Inclusion of small stores mirrors the observed effect on incumbent visits and purchases among small stores. Lowering the wholesale price to defend L mirrors the equipment transfers incumbents offer large stores, under the broader interpretation that wholesale prices in the model may reflect both price and non-price relationship terms. In a single-product model, at most one response arises for given parameters. In practice, different configurations of store locations and costs, as well as different cost structures across suppliers and valuations of their products, may lead to both responses occurring simultaneously as the data suggest.

Consumer welfare. Consumer welfare after the wholesaler's entry depends only on the wholesale prices the two stores pay, not on whether S buys from the incumbent or the entrant. I first consider the inclusion equilibrium of Proposition 2, in which both stores pay w^L. The same retail outcome arises when the incumbent excludes S but keeps its wholesale price (case (i) of Lemma 1), or exits, in which case both stores buy from the entrant at w^L.

The wholesaler's entry then leaves total product sales and the wholesale price unchanged, but gives some consumers access to a closer store. The covered-market retail prices are

$$p_L^{LS,E} = \nu - \tau + \frac{h_S - h_L}{3}, \qquad p_S^{LS,E} = \nu - \tau + \frac{2(h_S - h_L)}{3}.$$

Both retail prices exceed the pre-entry retail price at L, $p_L^L = \nu - \tau$.[36] Before the wholesaler's entry, L must keep its retail price low enough to serve even the most distant

[36] This increase reflects the assumption that the entrant charges exactly w^L. If it undercut the incumbent (as the data suggest) and the incumbent matched it, retail prices could be lower. Empirically, effects on retail prices are positive but imprecise at small stores (p = 0.165), while large-store retail prices fall 0.6% (p = 0.019), consistent with the better terms incumbents offer them.

consumer. Once S becomes active, L no longer needs to attract these distant consumers and raises its retail price.

Since all consumers purchase before and after the wholesaler's entry, the change in consumer surplus equals travel-cost savings minus the increase in retail payments:

$$CS^{LS,E} - CS^{L} = \underbrace{\left[\frac{\tau}{4} - \frac{(h_S - h_L)^2}{36\tau}\right]}_{\text{Travel-cost savings}} - \underbrace{\left[\frac{h_S - h_L}{2} - \frac{(h_S - h_L)^2}{18\tau}\right]}_{\text{Increase in retail payments}}.$$

A larger handling-cost gap raises retail prices and reduces the small store's market share, so fewer consumers switch to the closer store and travel savings decrease. The wholesaler's entry raises aggregate consumer surplus if and only if

$$h_S - h_L < (9 - 6\sqrt{2})\tau.$$

Consumers who continue purchasing from L lose from its higher retail price, while consumers sufficiently close to S gain from the shorter trip. When the incumbent instead lowers the wholesale price to defend L (cases (ii) and (iii) of Lemma 1), L's retail price, and S's when it remains active, are lower than if the incumbent served both stores. Every consumer is then at least as well off as in the inclusion equilibrium, so the condition above is sufficient for the wholesaler's entry to raise aggregate consumer surplus in every post-entry outcome.

The wholesaler's entry therefore separates improved access from the supplier's ability to extract its value through a higher wholesale price. Before the wholesaler's entry, activating S raises the wholesale price from w^L to w^{LS}. Afterward, S becomes active at w^L or L's wholesale price falls.

8 Quantifying consumer welfare

The expansion in product variety documented in Section 6 brings more products within reach of consumers, particularly through small stores. How much do consumers value this improved access? To answer this, I combine the experimental estimates with a complementary survey of neighborhood stores and their customers. The survey links consumers' shopping destinations to the variety, retail prices, and location of the stores available to them. I first describe the data and the patterns of store choice. I then develop and estimate a multiproduct extension of the consumer demand model presented in Section 7. Finally, I simulate an expansion in product variety like the one induced by the wholesaler's entry and quantify the associated consumer welfare gains.

8.1 Data

8.1.1 Complementary survey

The complementary survey was fielded in July and August 2026 in 30 control clusters, randomly sampled from the main experiment's clusters (stratified by district). In each cluster, I randomly selected a residential block away from the cluster boundary. Enumerators

first mapped all corner stores in and around the sampled block to build a plausible choice set for its residents. As part of this exercise, they also conducted a short survey of all store owners that consented (79% of listed stores). Enumerators then surveyed adults living in or near the sampled block about their shopping behavior. The resulting sample contains 576 consumers and 518 listed stores (17 stores per neighborhood on average). Appendix I describes the field protocol and reports neighborhood, consumer, and store characteristics.

In the store survey, owners reported the number of brands carried in nine categories selected to cover a broad range of leading products typically sold at corner stores. The survey also records selling area and a five-point variety score before requesting consent, so both are available for stores that declined the owner survey.[37] The survey records availability and retail prices for the same six popular products as the main experiment, from which I construct the same retail price index. Consumer and store GPS locations provide walking times between each household and the surrounding stores.

Consumers reported the two neighborhood corner stores they had visited most often during the preceding week. For each, they reported the number of visits, average expenditure per visit, and products bought on their most recent visit. They also reported purchases at other corner stores, markets, and supermarkets. When interviewing consumers, enumerators had a physical map of all neighborhood stores (based on the data they had collected) and were familiar with their names, owners, and appearance. They were therefore able to match shopping reports to actual stores, allowing me to link consumers to the characteristics of the destinations they chose and the alternatives they passed over.

8.1.2 Descriptives

The survey confirms the importance of corner stores described in Section 2. Corner stores account for 39% of consumers' reported grocery expenditure, and consumers make an average of 7.4 visits per week. Shopping trips are short, but consumers perceive them to be longer: the average round trip takes 3 minutes (calculated using GPS locations), yet consumers' reported walking times are 1.74 times as long on average (Appendix I.3 documents the comparison). The two named destinations account for 92% of their corner-store visits and 92% of corner-store expenditure. Thus, for each consumer, a small number of stores accounts for nearly all corner-store shopping, despite many nearby alternatives. Appendix Figure I2 illustrates this setting with two examples.

Figure 5 shows how corner store proximity and variety shape consumers' purchasing choices. High-variety stores (those with brand counts above the median) receive 100% more weekly visits and 113% more expenditure on average. Consumers are less likely to visit stores farther away: regardless of variety, the share of consumers that visit a given store is about 80% lower at three to four minutes of round-trip walking than within one minute. Yet proximity and variety leave considerable room for individual preferences: only 24% of consumers visit their closest low-variety store and 44% visit their closest high-variety store. Perhaps surprisingly, average expenditure and the number of products per visit do not differ significantly between low- and high-variety stores.

[37] I use these scores to impute missing brand counts. Their rank correlation with verified counts is 0.60, and Appendix I.2 describes the procedure.

8.2 Estimating consumer preferences

8.2.1 Multiproduct demand model

The consumer demand model in Section 7 considers the purchase of a single product: consumers trade off its retail price against the cost of reaching a store that carries it. Here I develop an extension of this model in which stores can carry many products. To make the problem tractable, I assume products are symmetric ex ante, but consumers have idiosyncratic tastes for particular products that vary across shopping occasions. Consumers buy one unit of one product per shopping occasion. The model therefore abstracts from differences in basket size, consistent with the lack of a significant difference in expenditure per visit between low- and high-variety stores documented above. For simplicity, I take the number of corner-store trips as given and omit an outside option, consistent with full coverage in equilibrium in Section 7 and the limited expenditure response in the experiment.

Variety attracts consumers by improving the match between their tastes and available products. Before choosing a store on each shopping occasion, consumers draw and observe product tastes η_{ik}, where i indexes consumers and k indexes products, which are independent Type-I extreme-value draws with scale $\sigma_\eta > 0$. Consumers buy the product they like best among those the chosen store carries. I assume a common ordering of availability, so a store carrying n_j products offers products $1, \ldots, n_j$. Stores carrying more products therefore offer all the products available at stores carrying fewer products, plus additional ones.[38] The same product has the same taste shock wherever it is carried. Let $M_i(n_j) = \max_{k \le n_j} \eta_{ik}$ denote consumer i's best product match at store j.

Consumers weigh their product match against travel costs and other store characteristics. Consumers have a general perception of each store's retail price level, summarized by the price index $\tilde{p}_j$. Idiosyncratic tastes for stores also allow destination choices to differ even at similar distances and variety levels, consistent with Figure 5. I add store tastes ϵ_{ij}, independently distributed Type-I extreme value, and normalize their scale to one without loss of generality. These tastes are independent of product tastes and redrawn on each shopping occasion. Let d_{ij} be round-trip walking minutes and X_j a vector of additional store controls with coefficients γ. Utility from store j is $U_{ij} = v_{ij}(\eta) + \epsilon_{ij}$, where

$$v_{ij}(\eta) = M_i(n_j) - \tau d_{ij} - \alpha \tilde{p}_j + X_j' \gamma. \tag{5}$$

Here τ is the cost per minute of round-trip walking, corresponding to the travel-cost parameter in Section 7, and α measures sensitivity to retail prices. The price index averages log relative retail prices and enters utility directly.

Under this taste distribution, expected product-match utility increases with log variety. Taking the maximum over n_j independent Type-I extreme-value draws gives

$$\mathbb{E}[M_i(n_j)] = \sigma_\eta(\log n_j + \gamma_E),$$

[38] This ordering can arise from supplier logistics that make it profitable to distribute some products more broadly than others. Anecdotally, corner stores exhibit little specialization: small stores tend to offer similar products to one another and subsets of the products found at larger stores.

where γ_E is Euler's constant. The coefficient on log variety in expected product-match utility is therefore σ_η, the dispersion of product tastes. Greater dispersion makes additional variety more valuable by creating more opportunities to find a product consumers particularly like.

Consumers observe their product tastes before choosing, but I do not observe these tastes in the survey. Conditional on product tastes, store choice is logit. Averaging these conditional probabilities over possible product tastes gives

$$P_{ij} = \mathbb{E}_\eta \left[\frac{\exp(v_{ij}(\eta))}{\sum_{\ell \in \mathcal{J}_i} \exp(v_{i\ell}(\eta))} \right], \tag{6}$$

where $\mathcal{J}_i$ is the set of mapped stores in consumer i's neighborhood.

8.2.2 Estimation

I estimate preferences by fitting predicted store-choice probabilities to consumers' weekly visit counts. I predict brand counts across all categories from the complementary survey's nine-category counts, using a log-log relationship estimated among experimental controls. Brands are counted separately within each category, so a brand appearing in two categories counts twice in both surveys. These predicted counts proxy for n_j and determine the joint distribution of product matches under the common ordering of availability. For each simulated draw, I generate these matches jointly across stores and calculate the conditional choice probabilities. Averaging across draws gives the probabilities in equation (6). I choose the preference parameters that make observed visits most likely by simulated maximum likelihood (Train, 2009), conditional on each consumer's total visits to the two named stores.[39] All mapped stores retained in the demand sample enter as alternatives, including those with zero reported visits.

I control for floor area in levels, as size may attract consumers through channels other than product variety. I set missing retail price indices to zero for the 18% of stores in the demand sample that declined the survey and include an indicator for consent to retain these observations.[40] For robustness, I allow stores to differ in persistent appeal beyond their observed characteristics. This specification adds a normally distributed store random effect, shared by all consumers and fixed across shopping occasions. Standard errors are clustered by neighborhood and account for uncertainty in the mapping between variety measures. Appendix J.1 gives the derivation, likelihood, and computational details.

The estimates confirm that consumers value variety and prefer nearby stores (Table 8). In the main specification, the coefficients on log brand counts and round-trip walking minutes are 1.23 and -0.63, respectively (both $p < 0.01$).[41] Retail prices enter negatively ($p = 0.045$),

[39] An alternative in which consumers learn product tastes only after arriving at the store yields a standard logit that can be estimated via PPML with consumer fixed effects. The resulting distance–variety trade-off and welfare gains are similar to those obtained with the preferred model. The linked consumer–store data allow direct likelihood estimation, without the aggregate market-share inversion used by Berry et al. (1995).

[40] The indicator may also capture other unobserved characteristics of the store, such as the probability that it is open and the owner's personality.

[41] A specification with store fixed effects exploits differences across consumers in their distance to the same stores, holding all store attributes fixed. It yields a distance coefficient of -0.71, close to the main estimate of -0.63 (column (3) of Table 8). Store fixed effects absorb variety, so this specification cannot identify the variety coefficient needed for the welfare exercise.

while floor area enters positively ($p = 0.006$), even after accounting for variety. The store-consent indicator is also positive ($p = 0.001$). Allowing for store random effects yields a similar variety coefficient and a larger response to distance. These preferences imply a substantial trade-off between variety and proximity. At equal distance and holding other store characteristics fixed, the estimated model predicts that high-variety stores receive 72% more weekly visits than low-variety stores. In a choice between two otherwise identical stores, a store offering twice as many brands can be an additional 1.1 minutes of round-trip walking away and still attract half the visits.

Concerns about endogeneity of variety are limited in this setting. First, I control for retail prices and floor area, which may attract consumers independently of variety. Second, corner stores are fairly homogeneous in other respects, with little credit to consumers or specialization beyond adding product categories. Third, a separate discrete-choice experiment that holds other store characteristics fixed shows a similarly strong preference for variety.[42]

8.3 Counterfactual: simulating variety gains from the experiment

I simulate the experiment in the complementary survey's store sample by applying the estimated treatment effect on log brand counts across all categories and its interaction with floor area. Specifically, I construct each store's counterfactual brand count as $n_j^{\text{cf}} = n_j \exp(a+bS_j)$, where S_j is floor area and a and b are the estimated coefficients on treatment and its interaction with floor area.[43] Locations, retail prices, floor area, and the number of consumer trips remain fixed, isolating the welfare effect of expanded variety.

As discussed in Section 2, wholesale purchase records from the main survey provide a useful proxy for the product variety available to consumers. Suppliers typically replenish stores weekly, so the three-week wholesale diary covers several replenishment cycles. The complementary survey supports this connection: restricting both surveys to the same nine categories, an additional square meter of floor area is associated with 1% more brands purchased among experimental controls and 1.6% more brands on the shelves of complementary-survey stores (Appendix Table J2). I use experimental effects on brands rather than products to match the variety measure in the demand estimation. Using all categories captures the broader expansion in product variety and preserves the full experimental sample. Appendix J.2 compares the experimental responses across category coverage and samples, and reports a version that uses only the nine survey categories in both demand estimation and the counterfactual.

For each consumer, I compare expected maximum utility under the baseline and counterfactual product offerings. Given product tastes, expected maximum utility is $\log \sum_j \exp(v_{ij}(\eta))$, up to a constant that cancels between the baseline and counterfactual.

[42] Respondents allocated hypothetical weekly visits between images of low- and high-variety stores, assuming that retail prices, their relationship with the owner, and other store characteristics were the same. At equal distance, they allocated 75% of visits to the high-variety store. Appendix I.4 describes the exercise and its limitations.

[43] The surveys have broadly similar floor-area distributions. I also transfer the experimental response using each store's size percentile within its survey sample to assess sensitivity to the remaining differences (Appendix J.2).

Averaging the difference over product tastes gives the per-trip gain

$$
\begin{aligned}
G_i &= \mathbb{E}_{\eta,\epsilon}\left[\max_{j\in\mathcal{J}_i} U_{ij}^{\text{cf}} - \max_{j\in\mathcal{J}_i} U_{ij}\right] \\
&= \mathbb{E}_{\eta}\left[\log\sum_{j\in\mathcal{J}_i}\exp(v_{ij}^{\text{cf}}(\eta)) - \log\sum_{j\in\mathcal{J}_i}\exp(v_{ij}(\eta))\right].
\end{aligned} \tag{7}
$$

Multiplying by the consumer's weekly visits to the named stores, T_i, and dividing by τ expresses this gain in equivalent walking minutes. I also report the uniform proportional reduction in retail prices that produces the same average utility gain. The 90% confidence intervals account for uncertainty in the estimated demand parameters, the mapping between variety measures, and the experimental effect on variety. Appendix J.3 describes their construction.

The welfare gain from increased variety is equivalent to 10% of baseline weekly travel costs (90% CI: [3.1, 20.5]). This corresponds to 1.9 minutes of round-trip walking per week, out of a predicted baseline of 19 minutes. Equivalently, it is worth a 5% reduction in retail prices (90% CI: [0.7, 19.8]). Using consumers' perceptions of walking time gives an equivalent of 3.3 minutes per week. Extrapolated to the 2,491,942 residents living in the study area, these gains are equivalent to 4,763,337 walking minutes per week.

Appendix Table J3 reports the results across demand specifications and alternative variety counterfactuals. The model with store random effects gives a similar gain, worth 8.7% of baseline travel costs. Together, these results indicate that the gains from increased variety at the stores already close to consumers were substantial.

9 Conclusion

Traditional retail in developing countries is fragmented downstream but often concentrated upstream. Large suppliers distribute largely non-overlapping product portfolios through separate networks, so the variety a store can offer depends on which suppliers are willing to serve it. I study how an increase in upstream competition changes stores' access to products, and its implications for consumers downstream. I randomize the entry of a multibrand wholesaler across low- and middle-income neighborhoods in Lima and use receipt-verified wholesale diaries to trace the responses of corner stores and incumbent suppliers.

I have three main results. First, access to an alternative sourcing option leads to variety gains concentrated among small, space-constrained stores, with no detectable changes among larger ones. Second, these gains are accompanied by an increase in small stores' wholesale purchases, and incumbent suppliers account for over three quarters of that growth. Third, incumbents respond to competitive pressure by increasing service and improving non-price terms for small stores, while also increasing equipment transfers to large stores. These adjustments occur without detectable changes in incumbent wholesale prices. Consumers also face fewer stockouts and, at large stores, lower retail prices.

A theoretical model explains these responses through the interaction of logistics costs and upstream market power. A monopolist supplier may exclude small stores when the additional

profits from serving them do not cover delivery costs, even though these stores would purchase at its prevailing wholesale price. A competitor's entry changes this trade-off: once excluded stores can source the product elsewhere, the incumbent may prefer to serve them rather than lose these sales to the entrant. The model also identifies an incentive to improve terms for established large-store clients to defend their market share. The effects of the wholesaler's entry therefore extend beyond purchases from the entrant itself, as incumbents adjust their service and the terms they offer to competing stores. These findings show why interventions aimed at alleviating sourcing constraints must account for incumbent suppliers' responses, especially in markets where they have market power.

Greater upstream competition can also reshape retail markets by allowing small stores to compete through product variety. In treated neighborhoods, the relationship between floor area and wholesale expenditure becomes flatter, consistent with consumers reallocating demand toward small stores. Such reallocation can benefit consumers even without an increase in aggregate expenditure, by making a wider range of products available close to home. Combining the experimental variety gains with a model of consumer store choice, I estimate a welfare gain equivalent to 10% of baseline weekly travel costs—roughly the travel cost of two-thirds of an average shopping trip per week.

The findings also offer a different perspective on the modernization of retail in developing countries. Previous research has shown that the entry of modern retailers benefits consumers through lower retail prices and better amenities, but these gains may come at the expense of independent stores (Atkin et al., 2018; Talamas Marcos, 2025). Modern wholesale intermediaries, by contrast, can strengthen independent stores by expanding their access to products, allowing consumers to benefit from greater variety without sacrificing proximity.

Tables

Table 1. Effect of entrant access on take-up

	Store survey			Administrative records		
	Any purchase	Expenditure	Orders (#)	Any purchase	Expenditure	Orders (#)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full sample						
Treatment	0.108***	7.56***	0.06***	0.167***	13.18***	0.11***
	(0.023)	(2.33)	(0.01)	(0.013)	(1.38)	(0.01)
Control mean	0.117	7.88	0.06	0.080	7.75	0.06
Observations	1,183	1,183	1,183	364	364	364
Panel B: Heterogeneity by store size						
Treatment × small	0.126***	8.67***	0.07***	—	—	—
	(0.030)	(2.22)	(0.02)	—	—	—
Treatment × large	0.079*	7.05	0.05	—	—	—
	(0.040)	(5.79)	(0.03)	—	—	—
Control mean (small)	0.081	4.93	0.04	—	—	—
Control mean (large)	0.163	11.80	0.08	—	—	—
p-value, equal effects	0.355	0.794	0.620	—	—	—
Size (cont.) interaction: coef	-0.018	-0.69	-0.00	—	—	—
Size (cont.) interaction: p-val	0.317	0.705	0.703	—	—	—
Observations	1,150	1,150	1,150	—	—	—

Notes: In columns (1)–(3) the unit of observation is the store. Column (1) equals one if the store placed at least one order with the entrant over the wholesale diary; column (2) is weekly expenditure with the entrant; column (3) is weekly orders placed with it. In columns (4)–(6) the unit of observation is the randomization cluster and the outcomes come from the entrant's administrative client records over the three-week window matched to the cluster's survey dates: column (4) is the number of clients with a delivered order divided by the census-adjusted number of stores in the cluster; columns (5) and (6) are those clients' expenditure and orders, per store and week. Monetary values are in USD PPP. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Table A5 reports columns (1)–(3) over receipt-verified purchases. Back to Section 6.1.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2. Effect of entrant access on product variety

	Products (#) (1)	Categories (#) (2)	Brands/ category (3)	Products/ brand (4)
Panel A: Full sample				
Treatment	3.16	0.85*	0.05	-0.02
	(2.61)	(0.51)	(0.04)	(0.03)
Control mean	36.96	11.48	1.69	1.69
Observations	996	996	996	996
Panel B: Heterogeneity by store size				
Treatment × small	5.50**	1.28**	0.12**	-0.05
	(2.61)	(0.64)	(0.05)	(0.04)
Treatment × large	0.71	0.35	0.01	0.00
	(4.77)	(0.87)	(0.05)	(0.05)
Control mean (small)	28.36	9.79	1.61	1.69
Control mean (large)	45.34	13.15	1.76	1.70
p-value, equal effects	0.382	0.405	0.150	0.363
Size (cont.) interaction: coef	-6.50	-0.74	-0.06	-0.01
Size (cont.) interaction: p-val	0.119	0.213	0.053	0.832
Observations	963	963	963	963

Notes: The unit of observation is the store. The sample is the stores with at least one receipt-verified supplier purchase. Products are counted over receipt-verified supplier transactions, so purchases at wholesale markets and modern stores are excluded. Column (1) counts total distinct products, where a product is a brand–size–flavor combination; column (2) counts distinct categories; column (3) is distinct category-by-brand combinations per category; column (4) is products per category-by-brand combination. The outcome in column (1) is the product of the outcomes in columns (2), (3) and (4). Free goods are counted, and the 2.3% of transactions without a matched category or brand are excluded. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Table A9 reports the same estimates in logs. Back to Section 6.2.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3. Effect of entrant access on product availability and retail prices

	Availability		Retail prices		
	Any missing (1)	Share missing (2)	Price index (log) (3)	Share below mode (4)	Share above mode (5)
Panel A: Full sample					
Treatment	-0.058**	-0.012*	-0.001	-0.000	-0.008
	(0.026)	(0.007)	(0.002)	(0.007)	(0.011)
Control mean	0.337	0.070	0.000	0.137	0.232
Observations	880	880	880	880	880
Panel B: Heterogeneity by store size					
Treatment × small	-0.050	-0.013	0.004	-0.022**	0.003
	(0.040)	(0.011)	(0.003)	(0.009)	(0.015)
Treatment × large	-0.091**	-0.013*	-0.006**	0.019*	-0.026*
	(0.037)	(0.007)	(0.002)	(0.011)	(0.015)
Control mean (small)	0.429	0.095	-0.003	0.188	0.255
Control mean (large)	0.223	0.041	0.003	0.070	0.194
p-value, equal effects	0.449	0.996	0.008	0.003	0.144
Size (cont.) interaction: coef	0.036	0.014	-0.003	0.014	-0.008
Size (cont.) interaction: p-val	0.227	0.157	0.237	0.067	0.509
Observations	857	857	857	857	857

Notes: The unit of observation is the store. The sample is the stores covered by the retail audit. Six popular products are tracked: canned tuna (Florida 140g), Coca-Cola (600ml), bottled water (Cielo 625ml), beer (Pilsen 630ml), toilet paper (Noble) and evaporated milk (Gloria 390g). Column (1) equals one if the store is missing any of the six; column (2) is the share of the six that are missing. Column (3) is the retail price index, the equal-weight mean across available products of log retail price minus that product's mean log price among control stores. Columns (4) and (5) are the shares of a store's priced products lying below and above the product's modal price among control stores. A product counts as missing when the store could not quote a retail price for it at the second visit. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Tables A10 and A11 report availability and retail prices product by product. Back to Section 6.2.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4. Effect of entrant access on wholesale expenditure

	All purchases				Receipt-verified	
	Total exp.	Entrant	Inc. suppliers	Markets/ modern stores	Entrant	Inc. suppliers
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full sample						
Treatment	12.18	7.56***	5.49	-0.87	5.06*	9.05
	(19.38)	(2.33)	(18.37)	(6.80)	(2.80)	(17.89)
Control mean	323.07	7.88	266.98	48.20	5.62	183.66
Observations	1,183	1,183	1,183	1,183	997	997
Panel B: Heterogeneity by store size						
Treatment × small	46.22***	8.67***	37.03***	0.53	8.46***	25.49*
	(13.73)	(2.22)	(13.60)	(5.78)	(2.23)	(13.55)
Treatment × large	-11.28	7.05	-28.25	9.91	4.23	-18.11
	(40.36)	(5.79)	(36.82)	(9.17)	(5.81)	(32.58)
Control mean (small)	233.03	4.93	196.42	31.68	3.28	120.96
Control mean (large)	419.68	11.80	339.88	68.00	8.38	245.41
p-value, equal effects	0.179	0.794	0.098	0.385	0.495	0.213
Size (cont.) interaction: coef	-71.86	-0.69	-70.91	-0.27	-0.70	-51.24
Size (cont.) interaction: p-val	0.020	0.705	0.030	0.970	0.644	0.070
Observations	1,150	1,150	1,150	1,150	964	964

Notes: The unit of observation is the store. In columns (1)–(4) the outcome is weekly expenditure over all orders, whether or not a receipt was produced. Column (1) covers purchases from every source, including wholesale markets and modern stores; column (2) is expenditure with the entrant, column (3) expenditure with incumbent suppliers, and column (4) purchases at wholesale markets and modern stores. In columns (5) and (6) the sample is stores with at least one receipt-verified supplier order and the outcome is weekly expenditure on the products digitized from those receipts: column (5) is expenditure with the entrant and column (6) expenditure with incumbent suppliers. Monetary values are in USD PPP. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Table A12 reports wholesale expenditure in logs. Back to Section 6.3.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5. Effect of entrant access on incumbent expenditure components: supplier visits, orders per visit, and order size

	All incumbent purchases				Receipt-verified		
	Total exp.	Total visits (#)	Orders/ visit	Exp./ order	**Total exp.**	Orders (#)	Exp./ order
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Full sample							
Treatment	5.26	0.11	0.017*	-1.74	9.97	0.29	-10.58
	(18.54)	(0.19)	(0.010)	(3.52)	(18.15)	(0.21)	(11.21)
Control mean	270.43	5.88	0.648	69.28	184.60	3.25	62.69
Observations	1,169	1,169	1,169	1,169	988	988	988
Panel B: Heterogeneity by store size							
Treatment × small	36.25***	0.36*	0.029*	5.85	27.59**	0.58**	1.24
	(13.66)	(0.20)	(0.016)	(3.57)	(13.82)	(0.23)	(4.62)
Treatment × large	-28.25	-0.26	0.010	-9.70	-20.01	-0.05	-22.97
	(36.82)	(0.34)	(0.012)	(6.71)	(32.76)	(0.35)	(21.52)
Control mean (small)	201.29	5.21	0.603	63.20	121.59	2.47	50.47
Control mean (large)	339.88	6.62	0.693	75.46	246.72	4.04	74.43
p-value, equal effects	0.102	0.122	0.337	0.044	0.178	0.136	0.275
Size (cont.) interaction: coef	-71.13	-0.42	-0.010	-14.28	-51.32	-0.42	-8.00
Size (cont.) interaction: p-val	0.030	0.142	0.241	0.014	0.071	0.173	0.061
Observations	1,137	1,137	1,137	1,137	954	954	954

Notes: The unit of observation is the store. In columns (1)–(4) the sample is stores with at least one incumbent-supplier order and the outcomes cover all incumbent orders: column (1) is weekly expenditure with incumbent suppliers; column (2) is total weekly supplier visits; column (3) is orders per visit; column (4) is expenditure per order. An order counts as a successful visit. Orders are from the wholesale diary and visits without purchase, excluding the entrant's, are from the store survey. In columns (5)–(7) the sample is stores with at least one receipt-verified incumbent-supplier order and the outcomes cover orders with a digitized receipt: column (5) is weekly expenditure with incumbent suppliers, column (6) the weekly number of orders and column (7) expenditure per order. The outcome in column (1) is the product of the outcomes in columns (2)–(4), and the outcome in column (5) the product of the outcomes in columns (6) and (7). Monetary values are in USD PPP. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Table A16 reports columns (1)–(4) in logs. Back to Section 6.4.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6. Effect of entrant access on non-price supplier benefits

	Supplier shelves (#) (1)	Supplier coolers (#) (2)	Received free goods (3)	Received trade credit (4)
Panel A: Full sample				
Treatment	0.08	0.15*	0.034	-0.011
	(0.10)	(0.09)	(0.028)	(0.022)
Control mean	1.81	1.66	0.570	0.249
Observations	1,183	1,183	1,183	1,183
Panel B: Heterogeneity by store size				
Treatment × small	-0.02	0.03	0.085**	0.053*
	(0.14)	(0.12)	(0.041)	(0.031)
Treatment × large	0.29**	0.36***	0.009	-0.065*
	(0.14)	(0.12)	(0.039)	(0.034)
Control mean (small)	2.10	1.57	0.484	0.210
Control mean (large)	1.46	1.72	0.660	0.273
p-value, equal effects	0.103	0.050	0.169	0.013
Size (cont.) interaction: coef	0.24	0.03	-0.010	-0.014
Size (cont.) interaction: p-val	0.016	0.768	0.705	0.551
Observations	1,150	1,150	1,150	1,150

Notes: The unit of observation is the store. Column (1) counts display shelves supplied and column (2) coolers; column (3) equals one if any digitized receipt records free goods (extra units delivered at no charge alongside a paid order) and column (4) if any incumbent supplier extended trade credit in the last month. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Table A17 reports free goods and trade credit at the store–supplier level. Back to Section 6.4.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7. Effect of entrant access on wholesale prices and quantities

	Incumbent purchases		All purchases	
	Price/package (log)	Packages (log)	Price/package (log)	Packages (log)
	(1)	(2)	(3)	(4)
Panel A: Full sample				
Treatment	-0.001	-0.011	-0.001	-0.015
	(0.001)	(0.024)	(0.001)	(0.024)
Control mean	2.060	0.113	2.064	0.114
Observations	33,368	33,368	34,330	34,330
Panel B: Heterogeneity by store size				
Treatment × small	-0.002	-0.002	-0.003	0.001
	(0.002)	(0.033)	(0.002)	(0.032)
Treatment × large	-0.000	-0.017	-0.001	-0.025
	(0.001)	(0.034)	(0.002)	(0.034)
Control mean (small)	2.084	0.037	2.089	0.037
Control mean (large)	2.046	0.158	2.049	0.160
p-value, equal effects	0.545	0.748	0.445	0.585
Size (cont.) interaction: coef	0.002	-0.023	0.001	-0.023
Size (cont.) interaction: p-val	0.004	0.240	0.065	0.255
Observations	33,300	33,300	34,266	34,266

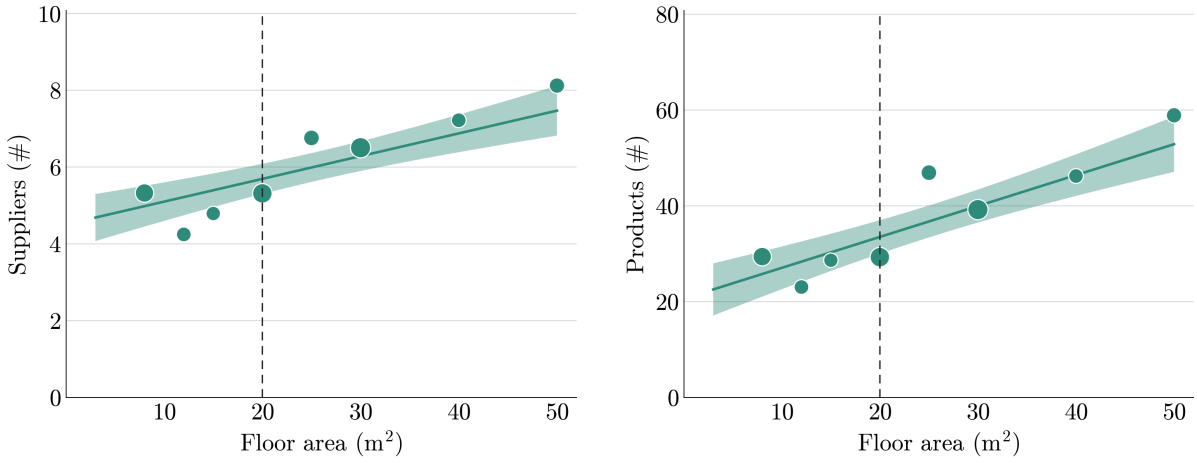
Notes: The unit of observation is the transaction, one product inside one order. Columns (1) and (3) are the log price of one package and columns (2) and (4) the log number of packages, a package being the modal quantity in which the product is bought. A product is a brand–size–flavor combination. Columns (1) and (2) exclude purchases from the entrant and columns (3) and (4) include them. Products without a matched category or brand, free goods and purchases at modern stores are excluded, as are products bought fewer than three times by control stores and prices outside the 1st and 99th percentiles of the ratio to the product's control-group median. Monetary values are in USD PPP. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum and product fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size and product-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Back to Section 6.4.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 8. Consumer preferences over neighborhood stores

	Main specification (1)	Store random effects (2)	Store fixed effects (3)
Log brand count (σ_η)	1.226	1.370	—
	(0.445)	(0.430)	
Round-trip walking minutes	-0.627	-0.753	-0.708
	(0.061)	(0.075)	(0.064)
Floor area (m^2)	0.018	0.028	—
	(0.007)	(0.014)	
Retail price index	-3.517	-4.136	—
	(1.757)	(3.134)	
Store consented to survey	0.693	0.917	—
	(0.205)	(0.366)	
Extra round-trip minutes for twice the brands	1.1	1.2	—
McFadden's pseudo R^2	0.225	0.386	0.293
Consumer–store pairs	9,427	9,427	5,103
Consumers	563	563	563
Stores	499	499	271
Neighborhoods	30	30	30

Notes: The unit of observation is the consumer–store pair in a mapped neighborhood. Columns (1) and (2) estimate the informed-choice model in equation (6), using weekly visit counts and all mapped alternatives meeting the demand-sample requirements. Column (2) adds an independent, normally distributed store random effect that is common to consumers and persists across shopping occasions. Column (3) estimates a destination logit with store fixed effects. Stores receiving no reported visits are excluded from column (3). The scale of occasion-specific store tastes is normalized to one. The log-brand-count row reports the product-taste scale σ_η. Brand counts sum distinct brands within categories. All-category counts are predicted from owner-reported, enumerator-verified nine-category counts, or counts imputed from pre-consent variety scores, using a log-log relationship estimated among experimental controls. The retail price index averages log relative prices of the same six popular products as the main experiment. Missing values are set to zero, and an indicator identifies stores that consented to the survey. McFadden's pseudo R^2 uses equal choice probabilities within the alternatives retained in each column. The row labeled Extra round-trip minutes for twice the brands compares two otherwise identical stores, one with double the other's brand count. It reports the additional round-trip walking minutes to the higher-variety store that make their predicted visit shares equal, integrating over product tastes and, in column (2), store random effects. Standard errors, in parentheses, are clustered at the neighborhood level. Columns (1) and (2) also incorporate mapping uncertainty, clustered by experimental stratum. Back to Section 8.2.2.

Figures

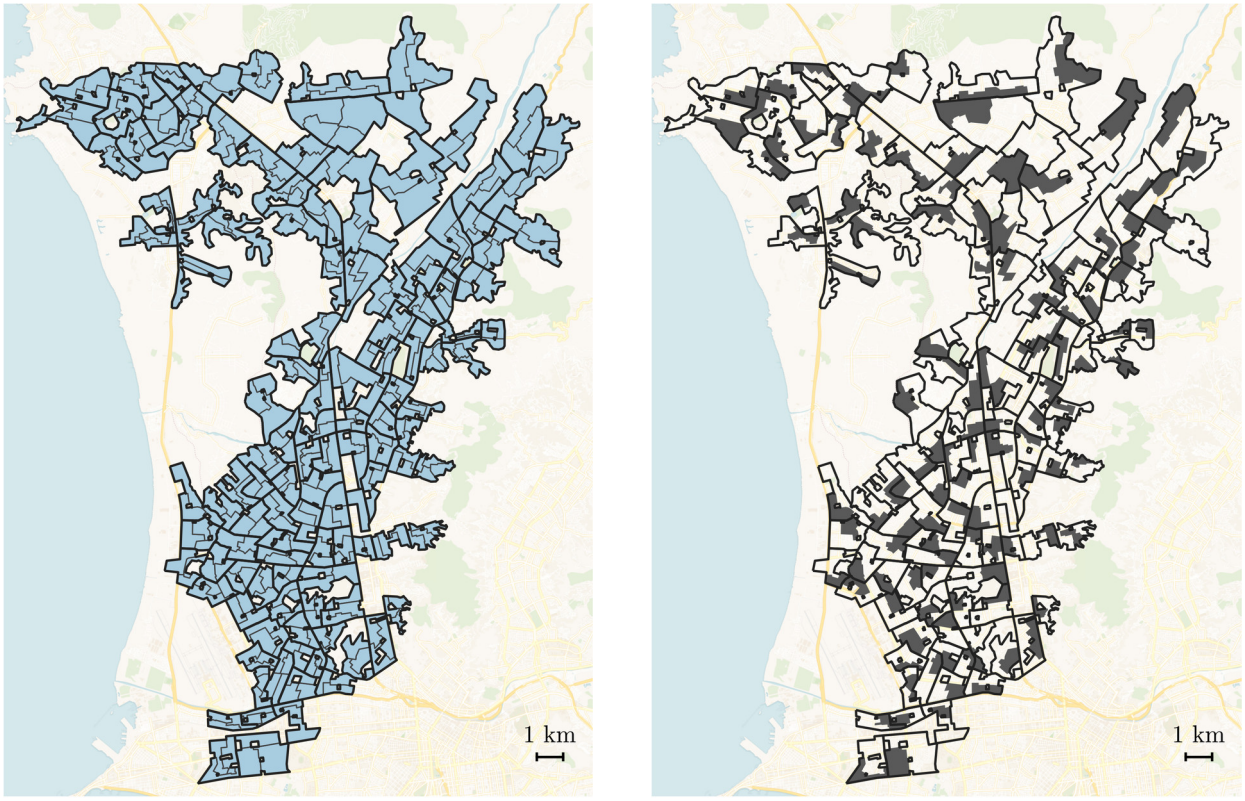


Panel A: Number of suppliers

Panel B: Number of products

Figure 1. Supplier relationships, product variety, and floor area

Notes: Panel A plots the number of suppliers and Panel B the number of distinct products purchased against floor area. Both panels use receipt-verified supplier purchases over the three-week diary among 393 control stores with positive observed expenditure, supplier counts, and product counts. Supplier counts include named suppliers and count each order from an "other" supplier separately. Unidentified suppliers are excluded. Products are distinct brand–size–flavor combinations with an assigned category and brand. Markers (green) report outcome means within quantile bins, placed at the median floor area in each bin. Larger markers denote bins containing more stores. The solid lines are store-level linear fits, and the shaded areas are their pointwise 90% confidence intervals from conventional OLS standard errors. The fitted slope is 0.06 suppliers per m^2 ($R^2 = 0.05$) in Panel A and 0.64 products per m^2 ($R^2 = 0.07$) in Panel B. The dashed vertical line marks the median floor area among control stores. Floor area is top-coded at 50 m^2 for plotting only. Regressions use individual-store observations and uncapped floor area. Back to Section 2.2.



Panel A: Study clusters

Panel B: Treatment assignment

Figure 2. Study clusters and treatment assignment

Notes: The figure shows the 364 clusters of the study in northern Lima, drawn within the 119 randomization strata. Panel A shades all clusters (blue). Areas left blank within a stratum—urban markets and their surroundings, and empty land such as industrial zones, farmland, public spaces, and hillsides—fall outside the study. Panel B shades the 119 control clusters (one per stratum) in dark gray. The remaining 245 treated clusters are left unshaded. Stratum boundaries are outlined in black in both panels. Background map tiles are CartoDB (OpenStreetMap data). Back to Section 3.2.2.

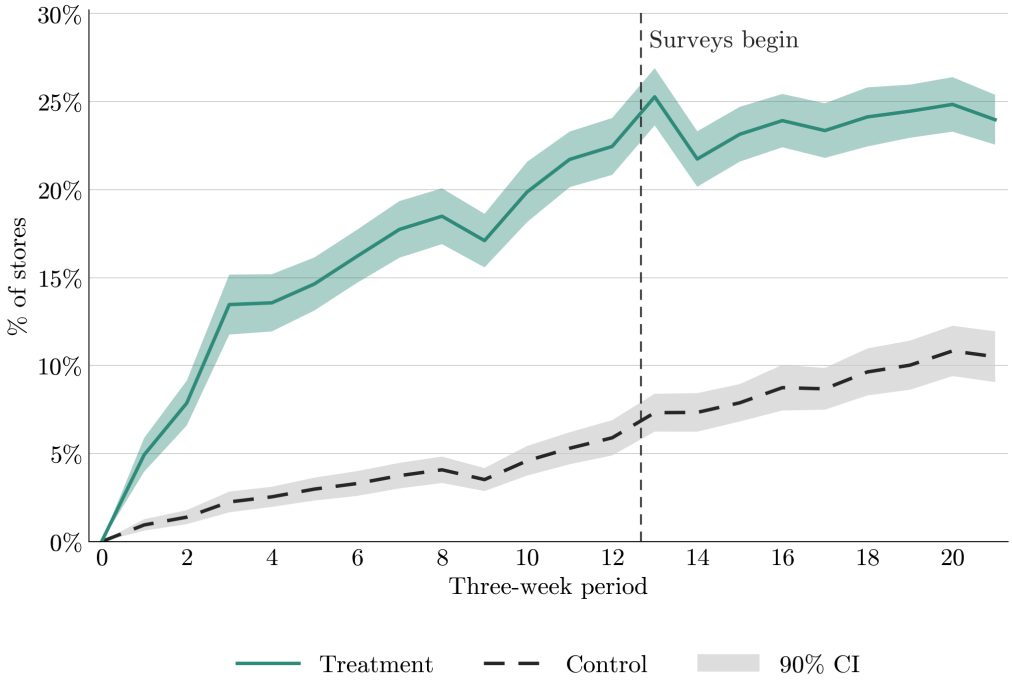


Figure 3. Entrant clients with any purchase in each three-week period

Notes: This figure shows the number of distinct clients with a delivered purchase in each three-week period, from the entrant's administrative client records, divided by the census-adjusted number of stores in the cluster. Period 1 begins July 1, 2024. The dashed vertical line marks the start of store surveys, and the last period shown is the last survey period. The solid green and dashed black lines are raw means across treated and control randomization clusters; shaded areas are pointwise 90% confidence intervals using standard errors clustered by stratum. Appendix Figure B8 reports the corresponding cumulative client shares. Back to Section 3.3.

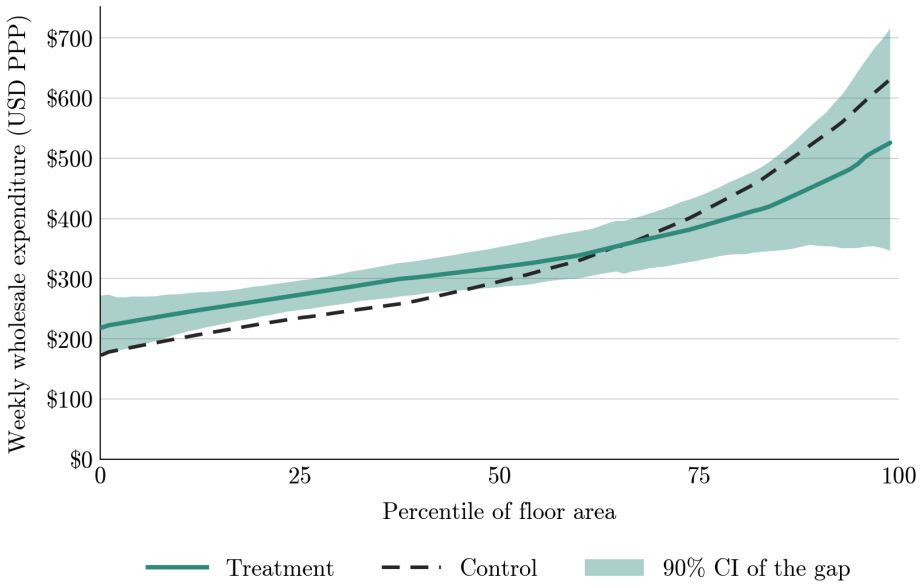


Figure 4. Wholesale expenditure by floor area, treatment and control

Notes: This figure shows weekly wholesale expenditure from every source (supplier orders plus purchases at wholesale markets and modern stores), in USD PPP, against store floor area, for control stores (dashed black line) and treated stores (solid green line). The horizontal axis is the percentile of floor area in the control-store distribution. The control line is a lowess smoother of the store-level outcome among control stores (bandwidth 0.8). The treatment line is the control line plus the smoothed treatment–control gap at that store size, and the shaded band is the pointwise 90% confidence interval of that gap from a bootstrap that resamples strata (500 replications). No fixed effects are absorbed. The figure covers 1,183 stores. Back to Section 6.3.1.

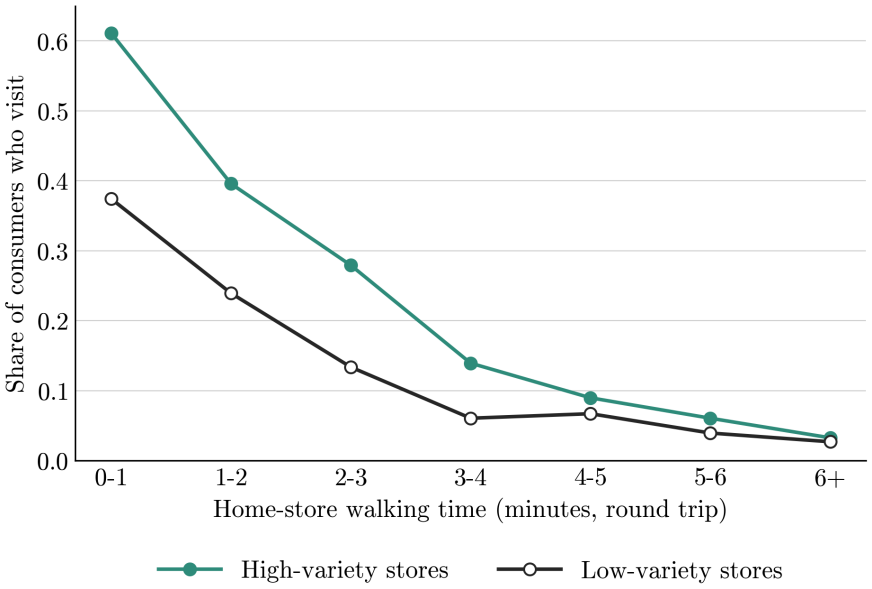


Figure 5. Distance, product variety, and store patronage

Notes: This figure shows the share of consumers who visited a store during the preceding week, over consumer–store pairs in the complementary survey, within bins of calculated round-trip walking time, for high-variety stores (green, filled markers) and low-variety stores (black, open markers). High variety means more than 27 brands, the median brand count. Brand counts are imputed from enumerators' variety scores at stores without complete brand counts. Back to Section 8.1.2.

References

Aghion, Philippe and Patrick Bolton. Contracts as a barrier to entry. *American Economic Review*, 77(3):388–401, 1987.

Åkerman, Anders. A theory on the role of wholesalers in international trade based on economies of scope. *Canadian Journal of Economics*, 51(1):156–185, 2018. doi: 10.1111/caje.12319.

Asker, John. Diagnosing foreclosure due to exclusive dealing. *Journal of Industrial Economics*, 64(3):375–410, 2016.

Asker, John and Heski Bar-Isaac. Raising retailers' profits: On vertical practices and the exclusion of rivals. *American Economic Review*, 104(2):672–686, 2014.

Atkin, David and Dave Donaldson. Who's getting globalized? The size and implications of intra-national trade costs. NBER Working Paper 21439, National Bureau of Economic Research, 2015.

Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro. Retail globalization and household welfare: Evidence from Mexico. *Journal of Political Economy*, 126(1):1–73, 2018. doi: 10.1086/695476.

Bai, Jie, David Sungho Park, and Ajay Shenoy. Misperceptions and product choice: Evidence from a randomised trial in Zambia. PEDL Working Paper, 2025.

Banerjee, Abhijit, Greg Fischer, Dean Karlan, Matt Lowe, and Benjamin N. Roth. Do microenterprises maximize profits? A nonrandomized vegetable-market experiment in India. CESifo Working Paper 12833, CESifo, 2026.

Bartkus, Viva Ona, Wyatt Brooks, Joseph P. Kaboski, and Carolyn Pelnik. Big fish in thin markets: Competing with the middlemen to increase market access in the Amazon. *Journal of Development Economics*, 155:102757, 2022.

Beaman, Lori, Jeremy Magruder, and Jonathan Robinson. Minding small change among small firms in Kenya. *Journal of Development Economics*, 108:69–86, 2014. doi: 10.1016/j.jdeveco.2013.12.010.

Bernheim, B. Douglas and Michael D. Whinston. Exclusive dealing. *Journal of Political Economy*, 106(1):64–103, 1998. doi: 10.1086/250003.

Berry, Steven, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

Borusyak, Kirill and Peter Hull. Nonrandom exposure to exogenous shocks. *Econometrica*, 91(6):2155–2185, 2023. doi: 10.3982/ECTA19367.

Brooks, Wyatt, Kevin Donovan, and Terence R. Johnson. Mentors or teachers? Microenterprise training in Kenya. *American Economic Journal: Applied Economics*, 10(4):196–221, 2018. doi: 10.1257/app.20170042.

Busso, Matias and Sebastian Galiani. The causal effect of competition on prices and quality: Evidence from a field experiment. *American Economic Journal: Applied Economics*, 11 (1):33–56, 2019. doi: 10.1257/app.20150310.

Cai, Jing and Adam Szeidl. Interfirm relationships and business performance. *Quarterly Journal of Economics*, 133(3):1229–1282, 2018. doi: 10.1093/qje/qjx049.

Cai, Jing and Adam Szeidl. Indirect effects of access to finance. *American Economic Review*, 114(8):2308–2351, 2024. doi: 10.1257/aer.20220711.

Cai, Jing, Wei Lin, and Adam Szeidl. Firm-to-firm referrals. NBER Working Paper 33082, National Bureau of Economic Research, 2024. URL https://www.nber.org/papers/w33082.

Casaburi, Lorenzo and Tristan Reed. Using individual-level randomized treatment to learn about market structure. *American Economic Journal: Applied Economics*, 14(4):58–90, 2022. doi: 10.1257/app.20200306.

Chatterjee, Shoumitro. Market power and spatial competition in rural India. *Quarterly Journal of Economics*, 138(3):1649–1711, 2023. doi: 10.1093/qje/qjad004.

Conlon, Christopher T. and Julie Holland Mortimer. Efficiency and foreclosure effects of vertical rebates: Empirical evidence. *Journal of Political Economy*, 129(12):3357–3404, 2021.

Couture, Victor, Benjamin Faber, Yizhen Gu, and Lizhi Liu. Connecting the countryside via e-commerce: Evidence from China. *American Economic Review: Insights*, 3(1):35–50, 2021. doi: 10.1257/aeri.20190382.

Dalton, Patricio S., Julius Rüschenpöhler, Burak Uras, and Bilal Zia. Curating local knowledge: Experimental evidence from small retailers in Indonesia. *Journal of the European Economic Association*, 19(5):2622–2657, 2021. doi: 10.1093/jeea/jvab007.

de Chaisemartin, Clément and Jaime Ramirez-Cuellar. At what level should one cluster standard errors in paired and small-strata experiments? *American Economic Journal: Applied Economics*, 16(1):193–212, 2024.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. Returns to capital in microenterprises: Evidence from a field experiment. *Quarterly Journal of Economics*, 123(4): 1329–1372, 2008. doi: 10.1162/qjec.2008.123.4.1329.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. Measuring microenterprise profits: Must we ask how the sausage is made? *Journal of Development Economics*, 88 (1):19–31, 2009. doi: 10.1016/j.jdeveco.2008.01.007.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. Labor drops: Experimental evidence on the return to additional labor in microenterprises. *American Economic Journal: Applied Economics*, 11(1):202–235, 2019. doi: 10.1257/app.20170497.

DellaVigna, Stefano and Matthew Gentzkow. Uniform pricing in U.S. retail chains. *Quarterly Journal of Economics*, 134(4):2011–2084, 2019. doi: 10.1093/qje/qjz019.

Dhingra, Swati and Silvana Tenreyro. The rise of agribusinesses and its distributional consequences. Working Paper, 2025.

Dixit, Avinash. Vertical integration in a monopolistically competitive industry. *International Journal of Industrial Organization*, 1(1):63–78, 1983. doi: 10.1016/0167-7187(83)90023-1.

Domínguez-Iino, Tomás. Efficiency and redistribution in environmental policy: An equilibrium analysis of agricultural supply chains. Working Paper, 2025.

Fabbri, Daniela and Leora F. Klapper. Bargaining power and trade credit. *Journal of Corporate Finance*, 41:66–80, 2016. doi: 10.1016/j.jcorpfin.2016.07.001.

Falcao Bergquist, Lauren and Michael Dinerstein. Competition and entry in agricultural markets: Experimental evidence from Kenya. *American Economic Review*, 110(12):3705–3747, 2020. doi: 10.1257/aer.20171397.

Falcao Bergquist, Lauren, Craig McIntosh, and Meredith Startz. Search costs, intermediation, and trade: Experimental evidence from Ugandan agricultural markets. NBER Working Paper 33221, National Bureau of Economic Research, 2024.

Fisman, Raymond and Mayank Raturi. Does competition encourage credit provision? Evidence from African trade credit relationships. *Review of Economics and Statistics*, 86(1): 345–352, 2004. doi: 10.1162/003465304323023859.

Fumagalli, Chiara and Massimo Motta. Exclusive dealing and entry, when buyers compete. *American Economic Review*, 96(3):785–795, 2006. doi: 10.1257/aer.96.3.785.

Ganapati, Sharat. The modern wholesaler: Global sourcing, domestic distribution, and scale economies. NBER Working Paper 32036, National Bureau of Economic Research, 2024.

Ghani, Tarek and Tristan Reed. Relationships on the rocks: Contract evolution in a market for ice. *American Economic Journal: Microeconomics*, 14(1):330–365, 2022. doi: 10.1257/mic.20190166.

Grant, Matthew and Meredith Startz. Cutting out the middleman: The structure of chains of intermediation. Working Paper, 2026.

Hardy, Morgan and Jamie McCasland. Are small firms labor constrained? Experimental evidence from Ghana. *American Economic Journal: Applied Economics*, 15(2):253–284, 2023. doi: 10.1257/app.20200503.

Hart, Oliver and Jean Tirole. Vertical integration and market foreclosure. *Brookings Papers on Economic Activity. Microeconomics*, 1990:205–286, 1990. doi: 10.2307/2534783.

Hristakeva, Sylvia. Vertical contracts with endogenous product selection: An empirical analysis of vendor-allowance contracts. *Journal of Political Economy*, 130(12):3202–3252, 2022. doi: 10.1086/720631.

Iacovone, Leonardo and David McKenzie. Shortening supply chains: Experimental evidence from fruit and vegetable vendors in Bogotá. *Economic Development and Cultural Change*, 71(1), 2022.

Karlan, Dean and Martin Valdivia. Teaching entrepreneurship: Impact of business training on microfinance clients and institutions. *Review of Economics and Statistics*, 93(2):510–527, 2011. doi: 10.1162/rest_a_00074.

Karlan, Dean, Natalia Rigol, and Benjamin N. Roth. When microenterprises grow, are consumers better off? Evidence from large loans to microenterprises in Chile. NBER Working Paper 35729, National Bureau of Economic Research, 2026.

Killeen, Grady. Risk aversion and barriers to firm growth: Experimental evidence from small retailers. Working Paper, 2026.

Kremer, Michael, Jean Lee, Jonathan Robinson, and Olga Rostapshova. Behavioral biases and firm behavior: Evidence from Kenyan retail shops. *American Economic Review*, 103 (3):362–368, 2013a. doi: 10.1257/aer.103.3.362.

Kremer, Michael, Jean Lee, Jonathan Robinson, and Olga Rostapshova. The return to capital for small retailers in Kenya: Evidence from inventories. Unpublished, 2013b.

Lagakos, David. Explaining cross-country productivity differences in retail trade. *Journal of Political Economy*, 124(2), 2016.

Lee, Robin S., Michael D. Whinston, and Ali Yurukoglu. Structural empirical analysis of contracting in vertical markets. In Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, editors, *Handbook of Industrial Organization*, volume 4, chapter 9, pages 673–742. Elsevier, 2021. doi: 10.1016/bs.hesind.2021.11.009.

Leone, Fabrizio, Rocco Macchiavello, Josepa Miquel-Florensa, and Nicola Pavanini. Market structure, vertical integration and farmers' welfare in the Costa Rica coffee industry. Working Paper, 2025.

Macchiavello, Rocco and Ameet Morjaria. Competition and relational contracts in the Rwanda coffee chain. *Quarterly Journal of Economics*, 136(2):1089–1143, 2021. doi: 10.1093/qje/qjaa048.

Mathewson, G. F. and R. A. Winter. Vertical integration by contractual restraints in spatial markets. *Journal of Business*, 56(4):497–517, 1983. doi: 10.1086/296214.

Mathewson, G. F. and R. A. Winter. The competitive effects of vertical agreements: Comment. *American Economic Review*, 77(5):1057–1062, 1987.

McAfee, R. Preston and Marius Schwartz. Opportunism in multilateral vertical contracting: Nondiscrimination, exclusivity, and uniformity. *American Economic Review*, 84(1):210–230, 1994.

McKenzie, David. Small business training to improve management practices in developing countries: Re-assessing the evidence for 'training doesn't work'. *Oxford Review of Economic Policy*, 37(2):276–301, 2021. doi: 10.1093/oxrep/grab002.

McKenzie, David and Susana Puerto. Growing markets through business training for female entrepreneurs: A market-level randomized experiment in Kenya. *American Economic Journal: Applied Economics*, 13(2):297–332, 2021. doi: 10.1257/app.20180340.

McKenzie, David and Christopher Woodruff. Experimental evidence on returns to capital and access to finance in Mexico. *World Bank Economic Review*, 22(3):457–482, 2008. doi: 10.1093/wber/lhn017.

Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria. Asymmetric information and middleman margins: An experiment with Indian potato farmers. *Review of Economics and Statistics*, 100(1):1–13, 2018. doi: 10.1162/rest_a_00699.

Nieuwoudt, Tielman. The digitisation of the traditional retailer in Africa. The Supply Chain Lab, March 2023. https://thesupplychainlab.blog/2023/03/03/the-digitalisation-of-the-traditional-retailer-in-africa/.

Nurski, Laura and Frank Verboven. Exclusive dealing as a barrier to entry? evidence from automobiles. *Review of Economic Studies*, 83(3):1156–1188, 2016. doi: 10.1093/restud/rdw002.

Quinn, Simon and Christopher Woodruff. Experiments and entrepreneurship in developing countries. *Annual Review of Economics*, 11:225–248, 2019. doi: 10.1146/annurev-economics-080218-030246.

Ramos-Menchelli, Daniel and Diana Sverdlin-Lisker. Fragmented markets and the proliferation of small firms: Evidence from mom-and-pop shops in Mexico. Working Paper, 2026.

Rasmusen, Eric B., J. Mark Ramseyer, and John S. Wiley, Jr. Naked exclusion. *American Economic Review*, 81(5):1137–1145, 1991.

Rey, Patrick and Jean Tirole. A primer on foreclosure. In Mark Armstrong and Robert Porter, editors, *Handbook of Industrial Organization*, volume 3, chapter 33, pages 2145–2220. Elsevier, 2007. doi: 10.1016/S1573-448X(06)03033-0.

Rey, Patrick and Thibaud Vergé. Economics of vertical restraints. In Paolo Buccirossi, editor, *Handbook of Antitrust Economics*, chapter 9, pages 353–390. MIT Press, Cambridge, MA, 2008.

Salop, Steven C. Monopolistic competition with outside goods. *The Bell Journal of Economics*, 10(1):141–156, 1979. doi: 10.2307/3003323.

Samaniego de la Parra, Brenda and Ajay Shenoy. Measuring and estimating retail productivity. IGC Working Paper ZMB-22099, International Growth Centre, 2024. URL https://www.theigc.org/publications/measuring-and-estimating-retail-productivity.

Segal, Ilya. Contracting with externalities. *Quarterly Journal of Economics*, 114(2):337–388, 1999. doi: 10.1162/003355399556016.

Segal, Ilya and Michael D. Whinston. Robust predictions for bilateral contracting with externalities. *Econometrica*, 71(3):757–791, 2003. doi: 10.1111/1468-0262.00425.

Segal, Ilya R. and Michael D. Whinston. Naked exclusion: Comment. *American Economic Review*, 90(1):296–309, 2000. doi: 10.1257/aer.90.1.296.

Startz, Meredith. The value of face-to-face: Search and contracting problems in Nigerian trade. Working Paper, 2025.

Szeidl, Adam. Firm-to-firm access and private sector development. CEPR Discussion Paper 20755, Centre for Economic Policy Research, 2025. URL https://cepr.org/publications/dp20755.

Talamas Marcos, Miguel Ángel. Surviving competition: Neighbourhood shops versus convenience chains. *The Review of Economic Studies*, 92(1):553–585, 2025. doi: 10.1093/restud/rdae023.

Train, Kenneth E. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2 edition, 2009. URL https://eml.berkeley.edu/books/choice2.html.

Tsai, Yiu Hing Barron. When cutting out the middleman backfires: Disintermediation, wholesale markups, and misallocation. Job Market Paper, 2025.

Verhoogen, Eric. Firm-level upgrading in developing countries. *Journal of Economic Literature*, 61(4):1410–1464, 2023. doi: 10.1257/jel.20221633.

Wiles, Edward and Deivy Houeix. Relational frictions along the supply chain: Evidence from Senegalese traders. Working Paper, 2026. URL https://edwardwiles.github.io/storage/papers/senegal_traders.pdf.

Zavala, Lucas. Unfair trade? Monopsony power in agricultural value chains. Working Paper, 2022.

Supplemental Appendix

Table of Contents

A Additional tables **56**

B Additional figures **73**

C International evidence on small retailers **82**

D Survey protocol **91**

E Digitization and product harmonization **96**

F Neighborhood exposure **97**

G Self-reported measures **105**

H Model derivations **115**

I Complementary survey **126**

J Demand estimation and welfare quantification **135**

A Additional tables

Table A1. Concentration and specialization of incumbent suppliers

Rank	Core categories	Exp. share	Store share	Within-product exp. share	Category concentration (top 3 share)
		(1)	(2)	(3)	(4)
1	Beer	0.277	0.496	0.916	0.970
2	Beverages	0.150	0.696	0.978	0.982
3	Food/household goods	0.096	0.779	0.677	0.250
4	Dairy	0.079	0.547	0.962	0.909
5	Food/household goods	0.054	0.430	0.508	0.331
6	Beverages	0.052	0.545	0.942	0.709
7	Beverages	0.034	0.368	0.909	0.940
8	Food/household goods	0.030	0.226	0.548	0.475
9	Beverages	0.021	0.394	0.842	0.940
10	Food/household goods	0.018	0.185	0.343	0.340

Notes: The unit of observation is the supplier. The sample is the 470 control stores, and the table lists the ten largest incumbent suppliers, ranked by column (1). Core categories describe each supplier's main product categories, with Food/household goods marking multicategory suppliers. Column (1) is each supplier's share of all incumbent-supplier expenditure; column (2) is the share of stores purchasing from the supplier; column (3) is the supplier-expenditure-weighted mean of its share of expenditure within each product it sells; column (4) is the share of the supplier's expenditure accounted for by its three largest categories. Columns (3) and (4) use 14,824 receipt-verified supplier transactions from 393 of these stores. Back to Section 2.3.

Table A2. Balance at the cluster level

	Control		Difference		
	Mean	SD	Mean	SE	Observations
	(1)	(2)	(3)	(4)	(5)
Corner stores (#)	69.47	[14.11]	1.53	(0.98)	364
Road density (km/km^2)	27.0	[4.5]	-0.2	(0.3)	364
Slope (degrees)	4.7	[3.4]	0.2	(0.2)	364
Altitude (m)	132.4	[88.4]	0.7	(2.3)	364
Area (km^2)	0.48	[0.29]	0.05**	(0.02)	364
Population	6,764.14	[2,134.37]	182.28	(125.18)	364
Mid SES	0.391	[0.400]	-0.011	(0.022)	364
High SES	0.175	[0.325]	0.007	(0.019)	364
Joint orthogonality (p-value)		0.256			364

Notes: The unit of observation is the randomization cluster. Characteristics are from the INEI corner store census and block-level population and socioeconomic status (SES), OpenStreetMap roads, and the SRTM elevation model. Columns (1) and (2) report the control-group mean and standard deviation (in brackets); columns (3) and (4) report the treatment-control difference and its standard error. Differences are estimated with stratum fixed effects. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). The joint orthogonality row reports the p-value from a test that the characteristics do not jointly predict treatment assignment under the same specification. Back to Section 3.2.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A3. Balance at the store level

	Control		Difference		
	Mean (1)	SD (2)	Mean (3)	SE (4)	Observations (5)
Panel A: Local characteristics					
Corner stores (#)	13.71	[5.73]	0.15	(0.34)	1,183
Road density (km/km^2)	27.4	[5.9]	0.2	(0.4)	1,183
Distance to major road (m)	108.9	[96.2]	4.8	(6.6)	1,183
Distance to market (m)	466.5	[536.9]	-18.5	(19.5)	1,183
Slope (degrees)	4.1	[3.1]	0.1	(0.2)	1,183
Altitude (m)	128.8	[86.4]	-0.8	(1.8)	1,183
Population	1,302.42	[592.95]	3.74	(25.95)	1,183
Mid SES	0.378	[0.440]	0.017	(0.024)	1,183
High SES	0.183	[0.365]	-0.012	(0.019)	1,183
Panel B: Store and owner characteristics					
Female	0.728	[0.446]	-0.002	(0.029)	1,183
Age (years)	49.9	[14.9]	1.1	(0.9)	1,183
Higher education	0.194	[0.396]	0.027	(0.026)	1,181
Has partner	0.681	[0.467]	0.002	(0.030)	1,177
Household size	4.46	[1.41]	0.11	(0.08)	1,164
Store at home	0.906	[0.292]	0.003	(0.015)	1,183
Years in business	9.5	[9.2]	-0.1	(0.5)	1,180
Floor area (m^2)	24.7	[17.0]	0.5	(0.8)	1,183
Joint orthogonality (p-value)					
Local characteristics			0.908		1,183
Store and owner characteristics			0.407		1,183
All characteristics			0.666		1,183

Notes: The unit of observation is the store. Panel A contains local characteristics measured within 150 meters of the store from the INEI corner store census and block-level population and socioeconomic status (SES), OpenStreetMap roads, and the SRTM elevation model. Panel B contains store and owner characteristics from the corner store survey. Columns (1) and (2) report the control-group mean and standard deviation (in brackets); columns (3) and (4) report the treatment-control difference and its standard error. Differences are estimated with stratum fixed effects. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). The joint orthogonality rows regress treatment assignment on the indicated characteristics under the same specification. For these tests only, missing Panel B characteristics are set to zero and missing-value indicators are added as controls. The p-values test the characteristics jointly, excluding the indicators. Back to Section 3.2.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4. Effect of entrant access on wholesale product composition

	Entrant product	Median package price (log)	Store share	Supplier HHI
	(1)	(2)	(3)	(4)
Panel A: Full sample				
Treatment	0.026**	0.022	0.003*	0.013*
	(0.012)	(0.016)	(0.002)	(0.007)
Control mean	0.400	2.073	0.078	0.659
Observations	34,330	34,330	34,330	34,201
Panel B: Heterogeneity by store size				
Treatment × small	0.017	0.026	0.001	0.024**
	(0.017)	(0.025)	(0.003)	(0.012)
Treatment × large	0.027*	0.014	0.003	0.015
	(0.016)	(0.024)	(0.002)	(0.012)
Control mean (small)	0.414	2.096	0.083	0.655
Control mean (large)	0.391	2.059	0.075	0.662
p-value, equal effects	0.620	0.721	0.416	0.610
Size (cont.) interaction: coef	-0.000	-0.010	0.001	-0.007
Size (cont.) interaction: p-val	0.991	0.479	0.643	0.205
Observations	34,330	34,330	34,330	34,201

Notes: The unit of observation is the transaction, one product inside one supplier order, where a product is a brand–size–flavor combination. Purchases from the entrant are included. Column (1) indicates whether the product appears in any paid entrant supplier transaction observed in either experimental arm. Column (2) is the log median control-group price of one modal package of the product. Column (3) is the share of control stores purchasing the product. Column (4) is the expenditure-share Herfindahl index across incumbent suppliers distributing the brand in the control group. Products without a matched category or brand, free goods and purchases at modern stores are excluded, as are products bought fewer than three times by control stores and prices outside the 1st and 99th percentiles of the ratio to the product's control-group median. Monetary values are in USD PPP. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Back to Section 6.2.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A5. Effect of entrant access on receipt-verified take-up

	Any purchase (1)	Expenditure (2)	Orders (#) (3)
Panel A: Full sample			
Treatment	0.094***	5.06*	0.05***
	(0.020)	(2.80)	(0.01)
Control mean	0.066	5.62	0.04
Observations	997	997	997
Panel B: Heterogeneity by store size			
Treatment × small	0.135***	8.46***	0.07***
	(0.031)	(2.23)	(0.02)
Treatment × large	0.061*	4.23	0.05
	(0.032)	(5.81)	(0.03)
Control mean (small)	0.052	3.28	0.04
Control mean (large)	0.085	8.38	0.05
p-value, equal effects	0.086	0.495	0.542
Size (cont.) interaction: coef	-0.024	-0.70	-0.00
Size (cont.) interaction: p-val	0.094	0.644	0.787
Observations	964	964	964

Notes: The unit of observation is the store. The sample excludes stores with no receipt-verified supplier order from any supplier. Column (1) equals one if the store placed at least one order with the entrant over the wholesale diary; column (2) is weekly expenditure with the entrant; column (3) is weekly orders placed with it. All three outcomes count only entrant orders with a digitized receipt. Monetary values are in USD PPP. Panels A and B, the Size (cont.) interaction rows and the standard errors are as in Table 1. Back to Section 6.1.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6. Correlates of entrant take-up

	All stores		Always-takers		Compliers		
	Mean	SD	Coef.	SE	Coef.	SE	Observations
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Local characteristics							
Corner stores (#)	13.77	[5.79]	0.018	(0.021)	-0.039	(0.023)	1,183
Road density (km/km^2)	27.5	[6.0]	-0.016	(0.019)	0.001	(0.022)	1,183
Distance to major road (m)	112.1	[102.3]	0.019	(0.020)	-0.050**	(0.023)	1,183
Distance to market (m)	456.8	[501.7]	-0.001	(0.023)	-0.026	(0.026)	1,183
Slope (degrees)	4.2	[3.3]	0.045	(0.032)	-0.037	(0.026)	1,183
Altitude (m)	129.2	[86.0]	-0.041	(0.055)	-0.022	(0.022)	1,183
Population	1,298.83	[599.16]	0.017	(0.029)	-0.017	(0.025)	1,183
Mid SES	0.391	[0.441]	-0.002	(0.025)	0.004	(0.024)	1,183
High SES	0.175	[0.356]	-0.011	(0.024)	0.018	(0.022)	1,183
Panel B: Store and owner characteristics							
Female	0.727	[0.446]	0.023	(0.015)	-0.028	(0.023)	1,183
Age (years)	50.6	[14.9]	-0.042***	(0.015)	0.004	(0.021)	1,183
Higher education	0.211	[0.408]	-0.023*	(0.014)	0.035*	(0.019)	1,181
Has partner	0.684	[0.465]	0.012	(0.016)	-0.030	(0.022)	1,177
Household size	4.53	[1.49]	0.003	(0.017)	0.010	(0.027)	1,164
Store at home	0.907	[0.291]	0.008	(0.014)	-0.034	(0.021)	1,183
Years in business	9.4	[9.4]	-0.021	(0.015)	-0.022	(0.021)	1,180
Large store	0.461	[0.499]	0.028	(0.017)	-0.027	(0.025)	1,183
Joint test of the interactions (p-value)							
Local characteristics				0.238			1,183
Store and owner characteristics				0.280			1,157
All characteristics				0.131			1,157

Notes: The unit of observation is the store, and take-up equals one if the store placed at least one order with the entrant over the wholesale diary. Columns (1) and (2) are the mean and standard deviation (in brackets) of each characteristic in natural units. Each row is a separate regression of take-up on cluster-level treatment assignment, that row's characteristic, and their interaction, with stratum fixed effects: column (3) reports the coefficient on the characteristic and column (5) the coefficient on the interaction, with standard errors in columns (4) and (6) and the regression's sample size in column (7). Each characteristic is centered and divided by its own standard deviation, so every coefficient is per standard deviation. Large store equals one if the store's floor area is above the control-group median. In control areas take-up equals always-taker status, so column (3) is the change in the share of always-takers and column (5) the change in the share of compliers. Panel A characteristics are measured within 150 meters of the store, from the INEI corner store census, OpenStreetMap roads, the SRTM elevation model and block-level population and socioeconomic status (SES), with distances to the nearest feature of each type. Panel B characteristics are from the corner store survey. The joint tests regress take-up on every characteristic in the block and its interaction at once, and report the p-value of the test that those interactions are all zero. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Back to Section 6.1.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A7. Entrant versus incumbent orders: expenditure and product variety

	Expenditure per order		Products per order			
	All purchases (1)	Receipt-verified (2)	**Products (#)** (3)	Categories (#) (4)	Brands/ category (5)	Products/ brand (6)
Panel A: Full sample						
Entrant	60.08***	65.53***	1.00**	1.05***	-0.16***	-0.28***
	(7.27)	(10.35)	(0.45)	(0.16)	(0.05)	(0.04)
Inc. mean	67.35	59.03	4.42	2.08	1.47	1.59
Observations	8,540	5,917	5,917	5,917	5,917	5,917
Panel B: Heterogeneity by store size						
Entrant × small	58.10***	61.51***	0.92**	1.13***	-0.20***	-0.23***
	(7.29)	(8.21)	(0.42)	(0.20)	(0.07)	(0.06)
Entrant × large	58.72***	73.36***	1.19	1.01***	-0.10	-0.36***
	(13.63)	(20.68)	(0.83)	(0.27)	(0.07)	(0.06)
Inc. mean (small)	65.45	53.55	4.31	2.07	1.47	1.57
Inc. mean (large)	69.01	63.35	4.51	2.09	1.48	1.61
p-value, equal effects	0.967	0.597	0.764	0.738	0.315	0.129
Size (cont.) interaction: coef	-2.29	-0.18	-0.13	-0.11	0.00	-0.03
Size (cont.) interaction: p-val	0.667	0.988	0.776	0.543	0.953	0.370
Observations	8,540	5,914	5,914	5,914	5,914	5,914

Notes: The unit of observation is the order, and the sample is supplier orders placed by stores in treated areas. Purchases at wholesale markets and modern stores are excluded. Column (1) is expenditure per order as reported by the store, over all supplier orders; column (2) is the value of the products digitized from the receipt. Column (3) counts the distinct products in the order, column (4) the distinct categories, column (5) the distinct category-by-brand combinations per category, and column (6) the products per category-by-brand combination. A product is a brand–size–flavor combination, and it enters columns (3)–(6) only if it carries both a category and a brand. The outcome in column (3) is the product of the outcomes in columns (4), (5) and (6). Columns (2)–(6) are restricted to orders with a digitized receipt. The rows labeled Inc. mean report the mean among incumbent orders. Monetary values are in USD PPP. Panel A reports the coefficient on the entrant indicator from a regression with stratum fixed effects. Panel B interacts the entrant indicator with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the entrant–incumbent difference per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with the entrant indicator. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Back to Section 6.1.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A8. Entrant versus incumbent transactions: wholesale prices and quantities

	Price/package (log) (1)	Packages (log) (2)
Panel A: Full sample		
Entrant	-0.047***	0.012
	(0.003)	(0.039)
Inc. mean	2.082	0.092
Observations	20,850	20,850
Panel B: Heterogeneity by store size		
Entrant × small	-0.050***	0.041
	(0.006)	(0.042)
Entrant × large	-0.048***	-0.026
	(0.005)	(0.061)
Inc. mean (small)	2.105	0.033
Inc. mean (large)	2.068	0.135
p-value, equal effects	0.794	0.358
Size (cont.) interaction: coef	-0.000	0.030
Size (cont.) interaction: p-val	0.973	0.313
Observations	20,692	20,692

Notes: The unit of observation is the transaction, one product inside one order, and the sample is stores in treated areas. Column (1) is the log price of one package and column (2) the log number of packages, a package being the modal quantity in which the product is bought. A product is a brand–size–flavor combination. The rows labeled Inc. mean report the mean among incumbent transactions. Products without a matched category or brand, free goods and purchases at modern stores are excluded, as are products bought fewer than three times by control stores and prices outside the 1st and 99th percentiles of the ratio to the product's control-group median. Monetary values are in USD PPP. Panel A reports the coefficient on the entrant indicator from a regression with stratum and product fixed effects. Panel B interacts the entrant indicator with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size and product-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the entrant–incumbent difference per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with the entrant indicator. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Back to Section 6.1.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A9. Effect of entrant access on product variety (logs)

	Products (1)	Categories (2)	Brands/ category (3)	Products/ brand (4)
Panel A: Full sample				
Treatment	0.136**	0.106**	0.036*	-0.006
	(0.066)	(0.052)	(0.020)	(0.016)
Control mean	3.058	2.094	0.472	0.492
Observations	996	996	996	996
Panel B: Heterogeneity by store size				
Treatment × small	0.227**	0.187**	0.068**	-0.028
	(0.088)	(0.074)	(0.028)	(0.022)
Treatment × large	0.061	0.040	0.018	0.003
	(0.102)	(0.076)	(0.028)	(0.028)
Control mean (small)	2.835	1.917	0.428	0.491
Control mean (large)	3.266	2.260	0.512	0.494
p-value, equal effects	0.226	0.177	0.196	0.383
Size (cont.) interaction: coef	-0.127	-0.086	-0.034	-0.007
Size (cont.) interaction: p-val	0.022	0.045	0.040	0.649
Observations	963	963	963	963

Notes: The unit of observation is the store. The sample is the stores with at least one receipt-verified supplier purchase. Each outcome is in logs. Products are counted over receipt-verified supplier transactions, so purchases at wholesale markets and modern stores are excluded. Column (1) counts total distinct products, where a product is a brand–size–flavor combination; column (2) counts distinct categories; column (3) is distinct category-by-brand combinations per category; column (4) is products per category-by-brand combination. The coefficients in columns (2), (3) and (4) sum to the coefficient in column (1). Free goods are counted, and the 2.3% of transactions without a matched category or brand are excluded. Panels A and B, the Size (cont.) interaction rows and the standard errors are as in Table 2. Back to Section 6.2.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A10. Effect of entrant access on product availability, by product

	Product missing at retail					
	Tuna	Coca-Cola	Water	Beer	Toilet paper	Milk
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full sample						
Treatment	0.010	-0.006	0.005	-0.035	-0.033	-0.011
	(0.016)	(0.007)	(0.010)	(0.023)	(0.022)	(0.009)
Control mean	0.063	0.011	0.014	0.180	0.137	0.017
Observations	880	880	880	880	880	880
Panel B: Heterogeneity by store size						
Treatment × small	0.039	-0.011	0.006	-0.062*	-0.029	-0.020
	(0.024)	(0.009)	(0.015)	(0.036)	(0.039)	(0.017)
Treatment × large	-0.025	0.004	0.008	-0.026	-0.038*	0.000
	(0.020)	(0.010)	(0.012)	(0.035)	(0.019)	(0.000)
Control mean (small)	0.073	0.016	0.021	0.236	0.199	0.026
Control mean (large)	0.054	0.007	0.007	0.122	0.047	0.007
p-value, equal effects	0.042	0.294	0.933	0.477	0.851	0.239
Size (cont.) interaction: coef	-0.024	0.011	0.009	0.055	0.021	0.012
Size (cont.) interaction: p-val	0.201	0.304	0.488	0.026	0.357	0.215
Observations	857	857	857	857	857	857

Notes: The unit of observation is the store. The sample is the stores covered by the retail audit. Each outcome equals one if that product is missing and zero otherwise. A product counts as missing when the store could not quote a retail price for it at the second visit. The six popular products are those of Table 3. Panels A and B, the Size (cont.) interaction rows and the standard errors are as in Table 3. Back to Section 6.2.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A11. Effect of entrant access on retail prices, by product

	Retail price (log)					
	Tuna	Coca-Cola	Water	Beer	Toilet paper	Milk
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full sample						
Treatment	0.003	-0.007*	-0.008	0.002	0.004	-0.002
	(0.004)	(0.004)	(0.008)	(0.003)	(0.003)	(0.002)
Control mean	1.128	0.536	-0.427	1.309	0.085	0.813
Observations	819	873	865	739	780	871
Panel B: Heterogeneity by store size						
Treatment × small	-0.002	0.000	0.008	0.006	0.006	-0.001
	(0.005)	(0.005)	(0.010)	(0.004)	(0.004)	(0.002)
Treatment × large	0.008	-0.015***	-0.022**	-0.004	0.004	-0.002
	(0.005)	(0.004)	(0.011)	(0.004)	(0.004)	(0.003)
Control mean (small)	1.117	0.536	-0.429	1.315	0.075	0.814
Control mean (large)	1.142	0.535	-0.425	1.302	0.097	0.812
p-value, equal effects	0.154	0.018	0.051	0.097	0.730	0.693
Size (cont.) interaction: coef	0.000	-0.005	-0.004	-0.002	0.002	0.000
Size (cont.) interaction: p-val	0.985	0.143	0.610	0.527	0.486	0.962
Observations	797	850	842	711	761	848

Notes: The unit of observation is the store. The sample is the stores covered by the retail audit. Each outcome is the log retail price. Stores that could not quote a price for a product are excluded from that column. The six popular products are those of Table 3. Monetary values are in USD PPP. Panels A and B, the Size (cont.) interaction rows and the standard errors are as in Table 3. Back to Section 6.2.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A12. Effect of entrant access on wholesale expenditure (logs)

	All purchases			Receipt-verified	
	Total exp.	Supplier exp.	Supplier exp. excl. entrant	Supplier exp.	Supplier exp. excl. entrant
	(1)	(2)	(3)	(4)	(5)
Panel A: Full sample					
Treatment	0.100**	0.107*	0.083	0.136*	0.112
	(0.048)	(0.055)	(0.055)	(0.076)	(0.077)
Control mean	5.414	5.239	5.206	4.520	4.485
Observations	1,178	1,170	1,169	993	988
Panel B: Heterogeneity by store size					
Treatment × small	0.219***	0.237***	0.211***	0.323***	0.258***
	(0.068)	(0.074)	(0.076)	(0.096)	(0.098)
Treatment × large	-0.004	-0.020	-0.038	-0.027	-0.019
	(0.078)	(0.083)	(0.085)	(0.121)	(0.124)
Control mean (small)	5.133	4.986	4.958	4.209	4.188
Control mean (large)	5.702	5.493	5.454	4.811	4.760
p-value, equal effects	0.037	0.025	0.034	0.028	0.087
Size (cont.) interaction: coef	-0.151	-0.181	-0.180	-0.207	-0.193
Size (cont.) interaction: p-val	0.001	0.000	0.001	0.001	0.004
Observations	1,146	1,138	1,137	959	954

Notes: The unit of observation is the store. Each outcome is in logs, and stores with a zero outcome are excluded from that column. In columns (1)–(3) the outcome is weekly expenditure over all orders, whether or not a receipt was produced. Column (1) covers purchases from every source, including wholesale markets and modern stores; column (2) restricts to supplier orders, which excludes those purchases; column (3) further excludes the entrant. In columns (4) and (5) the sample is stores with at least one receipt-verified supplier order and the outcome is weekly expenditure on the products digitized from those receipts. Monetary values are in USD PPP. Panels A and B, the Size (cont.) interaction rows and the standard errors are as in Table 4. Back to Section 6.3.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A13. Effect of entrant access on self-reported store performance

	Prod. purchases (1)	Prod. sales (2)	All income (3)	Customers (#) (4)	Hours (5)	Op. costs (6)	Asset value (7)
Panel A: Full sample							
Treatment	5.71	-18.83	-15.11	-1.25	0.2*	23.27	-146.13
	(20.25)	(58.34)	(60.17)	(1.87)	(0.1)	(17.10)	(136.79)
Control mean	353.11	914.52	1,024.83	34.11	12.7	158.78	1,903.16
Observations	1,160	1,150	1,150	1,181	1,173	1,183	1,160
Panel B: Heterogeneity by store size							
Treatment × small	20.56	30.07	27.84	0.30	0.3*	41.98	-3.68
	(20.25)	(34.78)	(36.37)	(1.01)	(0.1)	(40.77)	(145.71)
Treatment × large	2.11	-53.28	-46.36	-3.72	0.2	44.84**	-277.88
	(42.36)	(118.28)	(121.69)	(4.01)	(0.2)	(21.97)	(280.40)
Control mean (small)	282.58	695.93	790.24	28.60	12.9	159.96	1,739.24
Control mean (large)	432.10	1,148.55	1,276.94	40.07	12.6	150.40	2,081.18
p-value, equal effects	0.703	0.503	0.563	0.337	0.958	0.950	0.397
Size (cont.) interaction: coef	-23.73	-99.02	-93.92	-5.06	0.1	13.87	-397.47
Size (cont.) interaction: p-val	0.561	0.493	0.523	0.283	0.187	0.442	0.134
Observations	1,127	1,117	1,117	1,147	1,140	1,150	1,127

Notes: The unit of observation is the store. Columns (1) and (2) are weekly spending on and weekly sales of the study-covered packaged products, respectively. Column (3) adds sales from fresh categories, self-consumption of any products, and other income to packaged-product sales. Column (4) is customers per day; column (5) is hours open per day; column (6) is monthly operating costs; and column (7) is the total value of store-owned assets. Monetary values are in USD PPP. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Table A14 reports the same estimates in logs. Back to Section 6.3.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A14. Effect of entrant access on self-reported store performance (logs)

	Prod. purchases (1)	Prod. sales (2)	All income (3)	Customers (4)	Hours (5)	Op. costs (6)	Asset value (7)
Panel A: Full sample							
Treatment	0.018	-0.009	-0.005	-0.011	0.016*	0.013	-0.108*
	(0.037)	(0.034)	(0.032)	(0.027)	(0.009)	(0.058)	(0.057)
Control mean	5.624	6.571	6.710	3.360	2.528	4.609	7.106
Observations	1,159	1,150	1,150	1,180	1,173	1,139	1,159
Panel B: Heterogeneity by store size							
Treatment × small	0.033	-0.010	-0.010	0.003	0.018	-0.037	-0.039
	(0.049)	(0.040)	(0.038)	(0.032)	(0.012)	(0.077)	(0.073)
Treatment × large	0.033	0.026	0.028	-0.027	0.019	0.188**	-0.137
	(0.065)	(0.059)	(0.056)	(0.048)	(0.015)	(0.090)	(0.112)
Control mean (small)	5.456	6.370	6.517	3.222	2.536	4.639	7.053
Control mean (large)	5.805	6.777	6.908	3.504	2.517	4.520	7.150
p-value, equal effects	0.999	0.628	0.591	0.610	0.971	0.061	0.478
Size (cont.) interaction: coef	-0.039	-0.041	-0.037	-0.080	0.012	0.027	-0.153
Size (cont.) interaction: p-val	0.466	0.379	0.415	0.066	0.141	0.682	0.036
Observations	1,126	1,117	1,117	1,146	1,140	1,106	1,126

Notes: The unit of observation is the store. Each outcome is in logs, and stores with a zero outcome are excluded from that column. Columns (1) and (2) are weekly spending on and weekly sales of the study-covered packaged products, respectively. Column (3) adds sales from fresh categories, self-consumption of any products, and other income to packaged-product sales. Column (4) is customers per day; column (5) is hours open per day; column (6) is monthly operating costs; and column (7) is the total value of store-owned assets. Monetary values are in USD PPP. Panels A and B, the Size (cont.) interaction rows and the standard errors are as in Table A13. Back to Section 6.3.2. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A15. Treatment effects accounting for neighborhood exposure

	Weekly wholesale expenditure			Products (#)		
	Full sample	Small	Large	Full sample	Small	Large
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: No exposure adjustment						
Treatment	12.18	46.22***	-11.28	3.16	5.50**	0.71
	(19.38)	(13.73)	(40.36)	(2.61)	(2.61)	(4.77)
Panel B: Exposure to all nearby stores						
Direct effect	22.99	81.87***	-20.14	7.45*	13.08**	0.95
	(34.47)	(25.91)	(57.89)	(4.38)	(5.09)	(6.75)
Exposure	-4.56	-13.59**	4.66	-1.83	-2.98*	-0.12
	(10.94)	(6.73)	(19.99)	(1.56)	(1.56)	(2.81)
Implied cluster effect	7.07	32.94**	-4.53	1.05	2.37	0.54
	(21.55)	(13.82)	(47.59)	(3.35)	(3.02)	(6.68)
Panel C: Exposure outside own cluster						
Own-cluster effect	7.61	34.11**	-5.72	1.30	2.82	0.57
	(21.00)	(13.56)	(44.96)	(3.20)	(2.92)	(6.25)
Exposure	-4.17	-11.90**	4.23	-1.67	-2.59*	-0.11
	(10.00)	(5.81)	(18.14)	(1.42)	(1.33)	(2.57)
Control mean	323.07	233.03	419.68	36.96	28.36	45.34
Observations	1,183	1,150	1,150	996	963	963

Notes: The unit of observation is the store. Columns (1)–(3) report weekly wholesale expenditure from all sources; columns (4)–(6) report distinct products in receipt-verified supplier purchases over the wholesale diary, restricted to stores with such purchases. Monetary values are in USD PPP. Panel A omits exposure. Exposure counts treated census corner stores within 450 meters, including (Panel B) or excluding (Panel C) the own cluster. Exposure is measured in standard deviations of the recentered instrument (12.7 stores). Both 2SLS specifications use the same instrument: exposure minus its expectation conditional on own treatment. Columns (1) and (4) include stratum fixed effects. Other columns interact treatment and, where included, exposure with the focal store's size indicator (floor area above the control-group median), with stratum-by-size fixed effects. Direct effects hold total measured exposure fixed; cluster effects include treatment of neighbors inside the cluster. Panel B assumes equal spillovers within and across clusters. The implied cluster effect equals the direct effect plus the exposure coefficient times the mean own-cluster count in exposure units. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). The standard error of the implied cluster effect uses the delta method, treating counts as fixed. Back to Section 6.3.3. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A16. Effect of entrant access on incumbent expenditure components: supplier visits, orders per visit, and order size (logs)

	All incumbent purchases			
	Total exp. (1)	Total visits (2)	Orders/visit (3)	Exp./order (4)
Panel A: Full sample				
Treatment	0.083	0.016	0.042**	0.025
	(0.055)	(0.028)	(0.020)	(0.034)
Control mean	5.206	1.657	-0.492	4.041
Observations	1,169	1,169	1,169	1,169
Panel B: Heterogeneity by store size				
Treatment × small	0.211***	0.058	0.070**	0.083*
	(0.076)	(0.039)	(0.033)	(0.044)
Treatment × large	-0.038	-0.035	0.023	-0.026
	(0.085)	(0.042)	(0.023)	(0.056)
Control mean (small)	4.958	1.555	-0.583	3.986
Control mean (large)	5.454	1.765	-0.401	4.090
p-value, equal effects	0.034	0.107	0.248	0.131
Size (cont.) interaction: coef	-0.180	-0.051	-0.017	-0.112
Size (cont.) interaction: p-val	0.001	0.131	0.332	0.000
Observations	1,137	1,137	1,137	1,137

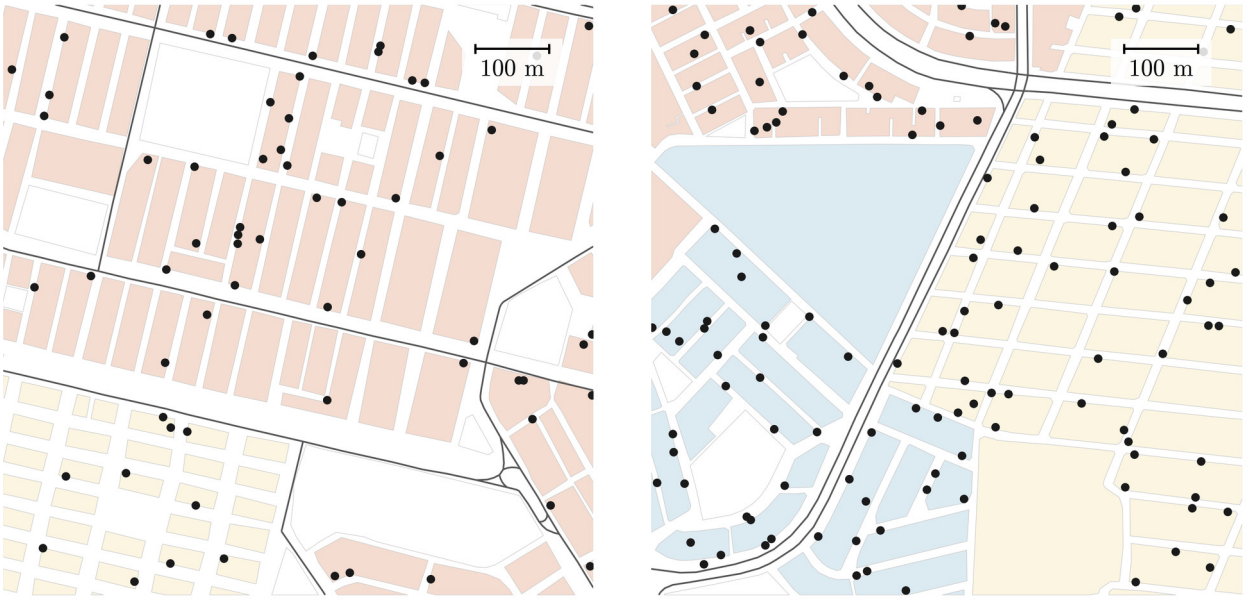
Notes: The unit of observation is the store. The sample is stores with at least one incumbent-supplier order. Each outcome is in logs. Column (1) is weekly expenditure with incumbent suppliers; column (2) is total weekly supplier visits; column (3) is orders per visit; column (4) is expenditure per order. The coefficients in columns (2), (3) and (4) sum to the coefficient in column (1). An order counts as a successful visit. Orders are from the wholesale diary and visits without purchase, excluding the entrant's, are from the store survey. Monetary values are in USD PPP. Panels A and B, the Size (cont.) interaction rows and the standard errors are as in Table 5. Back to Section 6.4.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A17. Effect of entrant access on non-price supplier benefits, store–supplier pairs

	Free goods		Trade credit	
	All pairs	Buyers	All pairs	Buyers
	(1)	(2)	(3)	(4)
Panel A: Full sample				
Treatment	0.001	0.008	0.000	0.001
	(0.001)	(0.009)	(0.000)	(0.004)
Control mean	0.012	0.193	0.003	0.040
Observations	130,130	8,312	130,114	8,295
Panel B: Heterogeneity by store size				
Treatment × small	0.002**	0.026**	0.001**	0.011*
	(0.001)	(0.013)	(0.000)	(0.006)
Treatment × large	0.000	0.003	-0.000	-0.008
	(0.002)	(0.014)	(0.000)	(0.006)
Control mean (small)	0.009	0.162	0.002	0.040
Control mean (large)	0.016	0.219	0.003	0.041
p-value, equal effects	0.320	0.226	0.046	0.033
Size (cont.) interaction: coef	-0.001	-0.014	-0.000	-0.006
Size (cont.) interaction: p-val	0.339	0.241	0.369	0.289
Observations	130,130	8,287	130,114	8,272

Notes: The unit of observation is the store–supplier pair. The sample is every assigned store crossed with every named incumbent supplier. Columns (1) and (3) indicate whether the store received free goods from, or bought on trade credit from, that supplier, with zeros for pairs with no purchase; columns (2) and (4) restrict to pairs with at least one purchase. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum and supplier fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size and supplier-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Back to Section 6.4.1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

B Additional figures



Panel A: Sparse neighborhood

Panel B: Dense neighborhood

Figure B1. Corner store density in two neighborhoods

Notes: Each panel shows an 800-meter square of the study area at about the 20th (Panel A) and 80th (Panel B) percentile of store count. Black dots are corner stores. Census blocks are shaded by socioeconomic status: poor and very poor (red), lower-middle-class (yellow), and upper-middle-class (blue). Dark gray lines are major roads, and thin outlines are block boundaries. Store locations are from the 2018–2019 INEI corner store census, socioeconomic status and block boundaries from INEI, and major roads from OpenStreetMap. Back to Section 2.1.



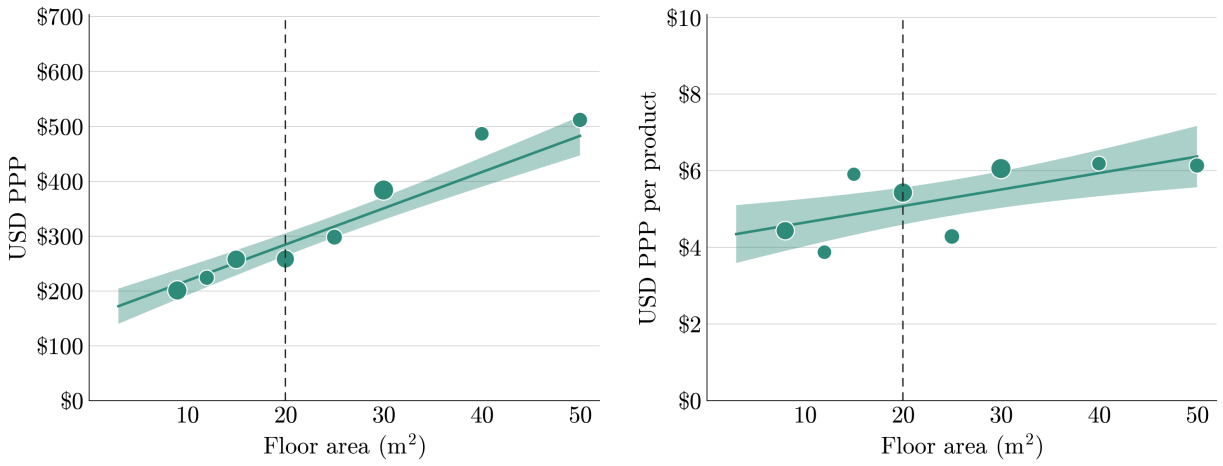
Panel A: A smaller store



Panel B: A larger store

Figure B2. Corner stores in the study area

Notes: This figure shows two corner stores in the study area, photographed during fieldwork. Back to Section 2.1.



Panel A: Weekly wholesale expenditure

Panel B: Expenditure per product

Figure B3. Store scale and floor area

Notes: Panel A plots weekly wholesale expenditure from the wholesale diary against floor area among 470 control stores. Panel B plots weekly receipt-verified supplier expenditure per distinct product among the 393 control stores of Figure 1, with products defined as in that figure. Markers (green) report outcome means within quantile bins, placed at the median floor area in each bin. Larger markers denote bins containing more stores. The solid lines are linear fits to individual stores using uncapped floor area, and the shaded areas are their pointwise 90% confidence intervals from conventional OLS standard errors. The fitted slope is 6.60 USD PPP per m^2 ($R^2 = 0.16$) in Panel A and 0.04 USD PPP per m^2 ($R^2 = 0.02$) in Panel B. The dashed vertical line marks the median floor area among control stores. Both outcomes are winsorized at their 99th percentile among the control stores plotted. Floor area is top-coded at 50 m^2 for plotting only. Back to Section 2.2.

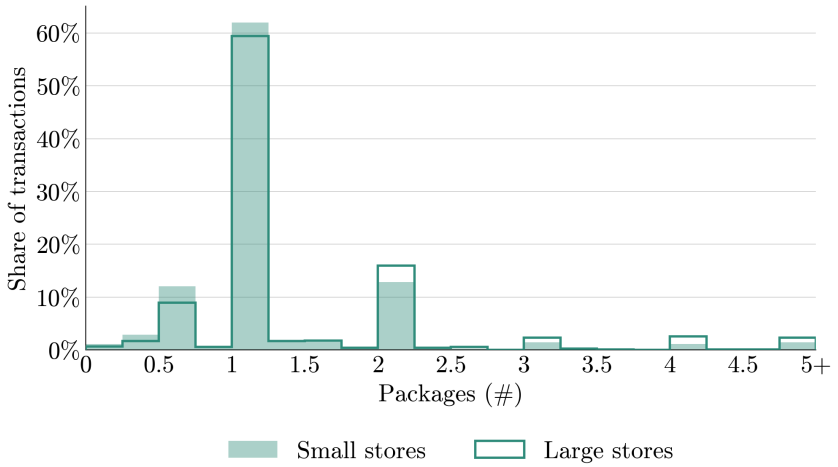
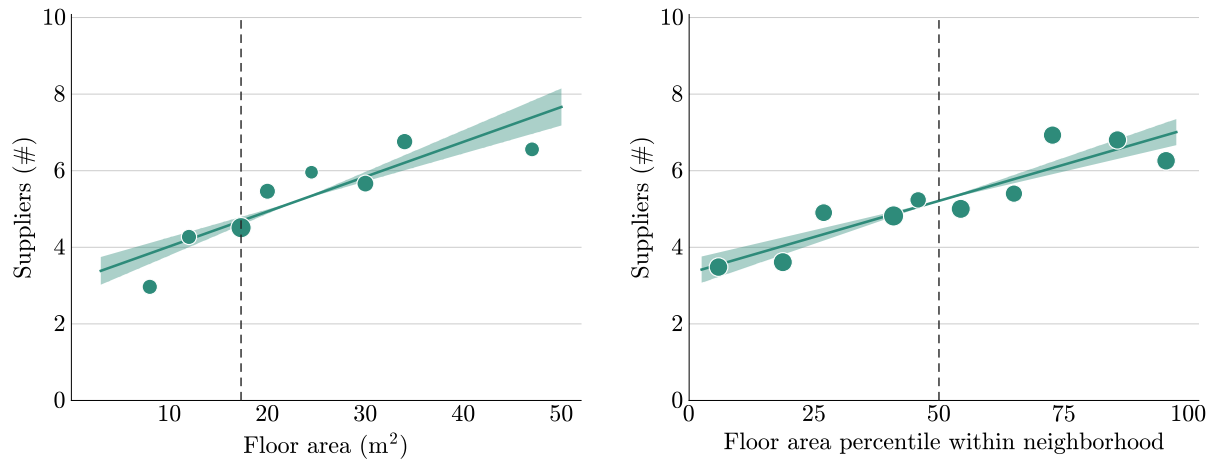
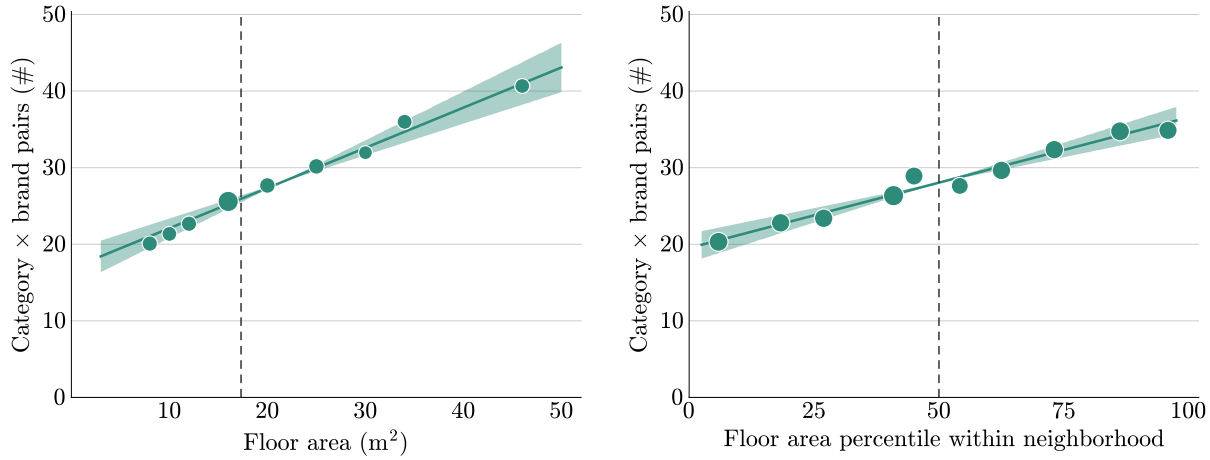


Figure B4. Package quantities by store size

Notes: The figure plots modal-package equivalents purchased per transaction among 13,446 transactions at 391 control stores, covering 993 products. The package and transaction definitions follow Table 7. Filled bars are small stores and outlined bars large stores. Purchases of five or more packages are pooled in the final bin. Back to Section 2.2.



Panel A: Number of suppliers



Panel B: Number of category × brand pairs

Figure B5. Supplier relationships and product variety by store size

Notes: Panel A plots the number of suppliers with reported visits in the previous two weeks, among nine named suppliers; Panel B plots the number of category × brand pairs across nine categories. Both panels use consenting complementary-survey stores with observed floor area and complete responses: 333 in Panel A and 403 in Panel B. Panel A excludes one enumerator's interviews with inconsistent supplier reporting. Markers (green) report neighborhood-adjusted outcome means within quantile bins, placed at the median floor area or percentile in each bin; larger markers denote bins containing more stores. The solid lines are store-level linear fits with neighborhood fixed effects, and the shaded areas are their pointwise 90% confidence intervals from neighborhood-clustered standard errors. The fitted slopes in Panel A are 0.09 suppliers per m^2 ($R^2 = 0.13$) and 0.04 suppliers per percentile point ($R^2 = 0.17$), respectively. In Panel B, they are 0.52 pairs per m^2 ($R^2 = 0.22$) and 0.17 pairs per percentile point ($R^2 = 0.21$); R^2 is calculated within neighborhoods. The dashed vertical lines mark the median floor area in the common observed-area sample (17.3 m^2) on the left and the 50th percentile within neighborhood on the right. Floor area is top-coded at 50 m^2 for plotting only; regressions use individual-store observations and uncapped floor area. Back to Section 2.2.

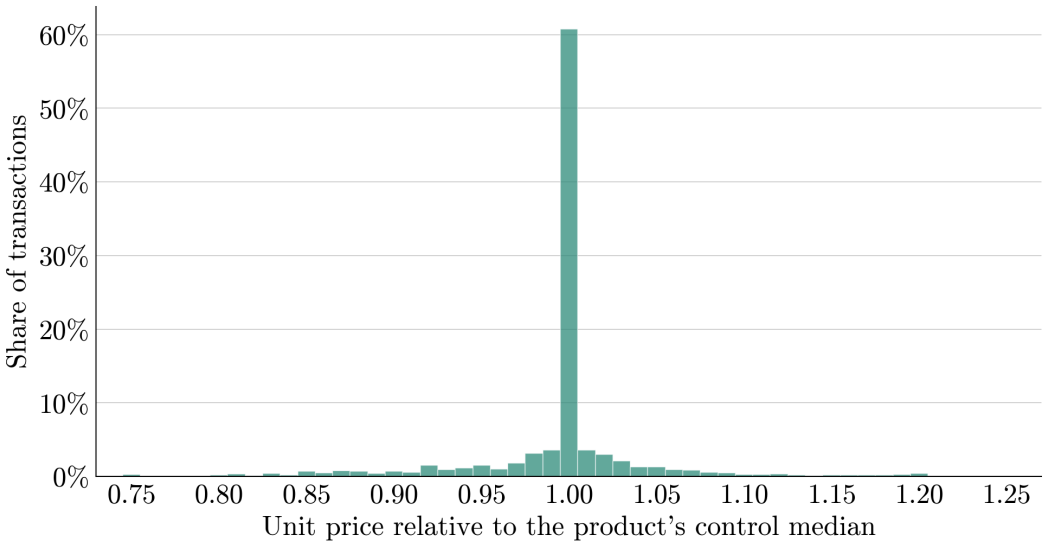
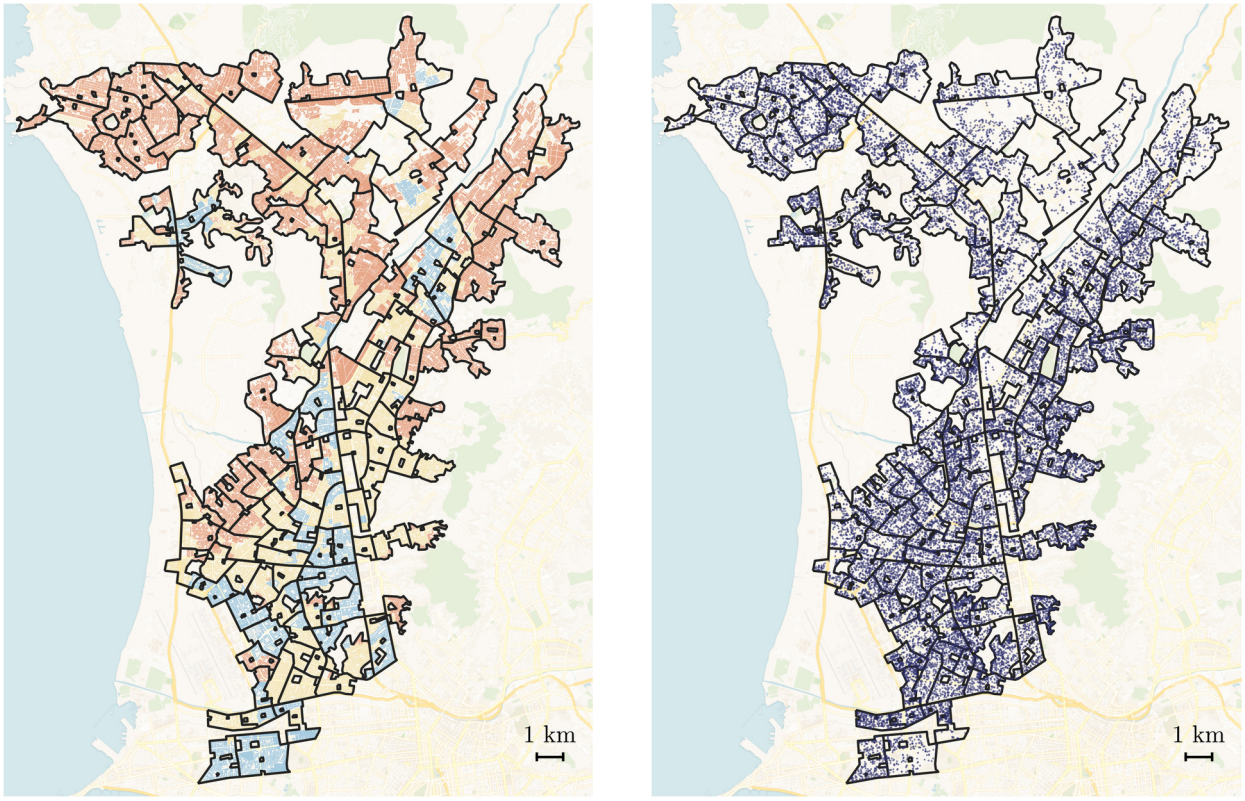


Figure B6. Wholesale price dispersion around the product median

Notes: This figure shows the distribution of wholesale unit prices relative to the median price paid for the same product by control stores, over 13,204 transactions in 993 products bought by control stores from incumbent suppliers. Products are defined and restricted as in Table 7, which excludes prices outside the 1st and 99th percentiles of the relative price. Bars are one percentage point wide, with one centered on the median. Back to Section 2.3.



Panel A: Socioeconomic status

Panel B: Corner stores

Figure B7. Study area: socioeconomic status and corner stores

Notes: The figure shows the study area in northern Lima. Panel A colors census blocks by socioeconomic status (SES): poor and very poor (red, 38% of the study area's population), lower-middle-class (yellow, 39%), and upper-middle-class (blue, 23%). Panel B outlines the 119 strata of the study area (black) and plots all 25,562 corner stores within them (navy). SES and block boundaries are from INEI. Store locations are from the 2018–2019 INEI corner store census. Background map tiles are CartoDB (OpenStreetMap data). Back to Section 3.2.1.

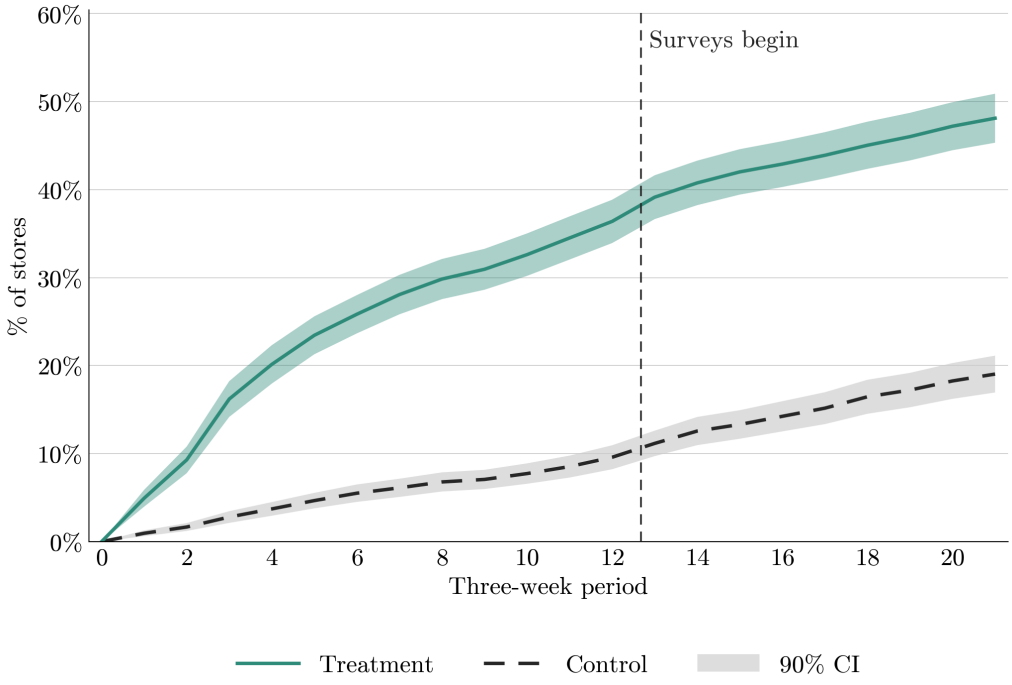


Figure B8. Entrant clients with any purchase by the end of each three-week period

Notes: This figure shows the number of distinct clients with a delivered purchase by the end of each three-week period, from the entrant's administrative client records, divided by the census-adjusted number of stores in the cluster. Period 1 begins July 1, 2024. The dashed vertical line marks the start of store surveys, and the last period shown is the last survey period. The solid green and dashed black lines are raw means across treated and control randomization clusters; shaded areas are pointwise 90% confidence intervals using standard errors clustered by stratum. Back to Section 3.3.

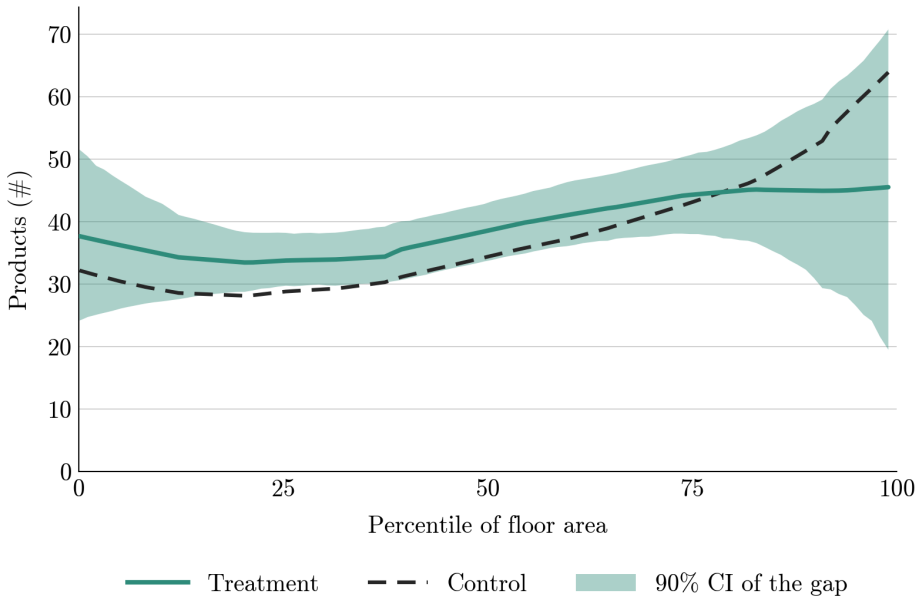


Figure B9. Product variety by floor area, treatment and control

Notes: This figure shows the number of distinct products a store carries, counted over receipt-verified supplier transactions as in Table 2, against store floor area, for control stores (dashed black line) and treated stores (solid green line). The horizontal axis is the percentile of floor area in the control-store distribution. The control line is a lowess smoother of the store-level outcome among control stores (bandwidth 0.8). The treatment line is the control line plus the smoothed treatment–control gap at that store size, and the shaded band is the pointwise 90% confidence interval of that gap from a bootstrap that resamples strata (500 replications). No fixed effects are absorbed. The sample is the stores with at least one receipt-verified supplier purchase (996 stores). Back to Section 6.2.1.

C International evidence on small retailers

I compile evidence from Latin America, Africa, and Asia to place the economic importance of corner stores in Peru in a broader context. The comparison distinguishes small, independently owned stores selling a general range of groceries from traditional retail more broadly, which also includes specialists, markets, and street vendors. Table C1 reports retailer counts and shares of businesses, together with the number of people per corner store, while Table C2 reports their shares of household spending or retail sales. No single source provides comparable estimates across countries. I therefore combine economic censuses and business and household surveys with published industry estimates, using the available categories to approximate these two groups.

I calculate business shares using the closest available business population, with the same denominator for both retailer groups within each country. For spending shares, I use household surveys or sales reports and combine the channels classified as traditional retail. The spending basket varies with the source, from food or consumer packaged goods to total retail sales. I report ranges when the sources do not support a single estimate and mark entries with a dash when no defensible estimate is available. The country notes below document the sources, calculations, and assumptions behind each estimate. Differences in business populations, spending baskets, and geographic coverage mean that the tables establish broad magnitudes rather than precise differences across countries.

C.1 Sources for Table C1

Peru. Redacción Gestión (2024) reports ABP's estimate of 535,000 bodegas. ComexPerú (2025) estimates 6.1 million nonfarm unincorporated micro/small businesses, including unregistered firms, with 32% in commerce. Applying the 70% retail share of registered commerce in INEI (2025) gives 1.37 million traditional retailers. This projection assumes the same retail mix among unregistered businesses. ComexPerú labels the estimates 2024, although its chart cites the 2023 household survey.

Bangladesh. The 2024 economic census records 1.18 million general stores predominantly selling food, beverages, or tobacco and 4.35 million retail units among 11.7 million nonfarm economic units (BBS, 2026). Permanent, temporary, and home-based businesses are covered. Market-stall activities are classified separately from general stores. Supermarkets cannot be separated from general stores, and chain ownership is unobserved in broader retail. These activity categories therefore approximate the two store groups.

China. China Economic Net (2020) reports a Ministry of Commerce estimate of six million independent convenience stores selling consumer packaged goods. The business share uses the 2018 census total of 63 million individual nonfarm businesses (NBS, 2019). The store estimate was reported in 2020. A broader retailer count is unavailable because the published individual-business total combines wholesale and retail.

Colombia. Gómez Goyeneche (2026) reports NIQ's estimate of 453,750 neighborhood stores. DANE (2025a) supplies 2024 nonfarm microbusiness and commerce totals, including

unregistered firms. The traditional-retailer estimate applies the 85% retail share among urban commerce respondents in DANE (2026) to the microbusiness commerce total, assuming the same retail mix. Expanding for the 3% of residents outside survey coverage gives a business denominator of 4.3 million (DANE, 2025b), assuming equal businesses per resident and the same retail mix in uncovered areas.

Egypt. Upadhyay et al. (2022) estimate 400,000 neighborhood grocery stores. CAPMAS (2025) supplies the 2022–23 denominator of 3.86 million fixed-premises establishments and 1.78 million retail establishments with fewer than ten workers, used to approximate traditional retail. This food-and-nonfood count includes fixed market stalls but excludes mobile sellers. Small chain branches may remain. The grocery-store count is an industry estimate, rather than a census enumeration.

India. Sriram and Pandya (2024) report an estimated thirteen million kirana grocery stores. MoSPI (2025) supplies the 2023–24 denominator of 73.4 million unincorporated nonfarm businesses and its 27% retail-trade share, implying 19.9 million traditional retailers. The survey includes registered and unregistered businesses but excludes construction. Its retail category covers food and nonfood goods, excluding motor vehicles. Unincorporated ownership approximates traditional retail.

Indonesia. Modgil (2022) reports Aksesmu's estimate of 3.6 million warung grocery stores. The BPS census supplies 26.4 million nonfarm businesses and 10.8 million single-location retailers, including home businesses, market sellers, and mobile vendors (BPS, 2018, n.d.). Single-location retailers approximate traditional retail. The corner-store share combines the 2022 reported count with the older 2016 business denominator.

Kenya. MSEA (2024) reports about 600,000 microretail stores in its discussion of dukas. This is used as an approximate general-store count. The report does not separately tabulate specialists or market sellers. KNBS (2016) supplies the 2016 denominator of 7.41 million licensed and unlicensed MSMEs. Applying retail's 91% share of licensed trade to all trade gives 4.14 million retailers, assuming the same industry mix among unlicensed businesses. This projection includes food and nonfood retail.

Mexico. INEGI reports 574,141 grocery stores and 4.80 million covered businesses for 2018 (INEGI, 2023, n.d.). Subtracting multilocation firms from retail firms gives 1.97 million single-location retailers, the traditional-retail proxy. Fixed market premises are covered; itinerant vendors and separately sampled rural businesses are excluded. National population per covered grocery store is consequently an upper bound on people per store.

Philippines. DTI (2024) reports more than 1.3 million sari-sari stores. The approximate business share uses a policy estimate of seven million registered and unregistered MSMEs discussed in the 2021 budget hearings, as reported by Mercado (2021). This is a comparison across years using an approximate denominator. A comparable traditional-retailer count is unavailable.

South Africa. Reddy (2025) reports more than 150,000 spaza shops. Motsomi (2024) supplies the 2024 denominator of 3.91 million MSMEs, including informal and agricultural

businesses. Spazas are a narrow neighborhood-store format that excludes larger independent general dealers. Available trade totals also include wholesale and repairs, preventing a comparable broader retailer count.

Vietnam. AAFC (2022) reports Nielsen's estimate of more than 1.4 million neighborhood grocery stores in 2021. Markets are listed separately. The denominator is 5.07 million nonfarm household businesses in 2021, including registered and unregistered businesses (Le, 2024). Applying the approximately one-third retail share of household businesses reported by World Bank (2022) for 2020 gives about 1.7 million traditional retailers. This projection assumes that the survey's industry mix applies to the census population. The unusually high grocery-store share of this total may reflect differences between the industry estimate and survey classifications.

C.2 Sources for Table C2

Peru. Cheng (2025) reports a 36% bodega share of household spending on fast-moving consumer goods. The broader share in Kantar Worldpanel (2024) exceeds 80% and includes markets, specialists, independent minimarkets, and household purchases from wholesalers. The two estimates refer to different periods.

Bangladesh. Mahmud (2025) reports a 2–3% modern-retail share, including online sales. The traditional share is the remainder, assuming these sectors exhaust retail sales. Available roadside-shop estimates also include open-air markets and cannot identify corner stores separately.

Brazil. Calculations use monetary food-at-home purchases in IBGE's 2017–18 household expenditure survey (IBGE, 2023), with final survey weights and corrected, deflated purchase values. Corner stores comprise general grocers and neighborhood grocery stores; traditional retail adds food specialists, bakeries, markets, mobile sellers, and direct sellers. Nonmonetary acquisitions are excluded throughout; restaurant purchases are excluded from the retailer numerator. These are format proxies: independent supermarkets are excluded, while specialist chains may remain.

China. Agricultural Trade Office Beijing (2025) assigns 57.5% of grocery-retailer sales to local grocers, including market sellers, and 2.9% to food and drink specialists. Their sum gives the 60% traditional-format proxy. Specialist chains may remain. The source cannot identify corner stores separately.

Colombia. Gómez Goyeneche (2026) reports NIQ's 37% neighborhood-store share of consumer packaged-goods sales. Salinas (2025) reports grocery-retailer shares of 47% for local grocers, including markets, and 13% for specialists, giving 60% for traditional retail. Specialist chains may remain.

Egypt. The traditional-format share adds local-grocer and food, drink, and tobacco specialist sales in Abdel Halim (2025, Figure 1), giving 80% of grocery-retailer sales. Local grocers include market sellers and cannot identify corner stores separately. Specialists are treated as predominantly independent.

India. Worldpanel by Numerator (2025, p. 24) reports a 78% urban household share through kirana/paan-beedi stores, including tobacco kiosks. The 78–86% traditional-retail range assigns none or all of the chemist and residual channels to traditional retail. These are packaged-goods shares rather than shares of all food purchases.

Indonesia. Minot et al. (2015) survey households in Surabaya, Bogor, and Surakarta in 2010–11. Warungs receive 31% of food expenditure, wet markets 24%, and peddlers 11%, giving at least 66% through traditional retailers. Spending includes meals outside the home, and each food category is assigned to its principal outlet. The subtotal omits separately unidentified spending at semi-permanent stalls and excludes restaurants.

Kenya. Japan International Cooperation Agency and Deloitte Tohmatsu Financial Advisory LLC (2020, p. 58), drawing on industry information and interviews, estimate traditional retail at 70% of retail sales, including food and nonfood. The estimate includes markets and does not separate corner stores.

Mexico. Domínguez-Barreto et al. (2025) report survey-weighted shares of food and beverage purchases for home consumption in ENIGH 2022: 33% through neighborhood grocery stores and 84% after adding specialists, public and street markets, street vendors, and purchases from acquaintances. Meals consumed away from home, pet food, tobacco, and government food provision are excluded. Purchases from acquaintances are sales, rather than gifts.

Nigeria. Office of Agricultural Affairs, Lagos (2024, p. 5) assigns 25% of food-retail sales to small grocery and convenience stores, 72% to open-air markets, and 3% to supermarkets. The first two channels give 97%. The convenience category may include small chains. Online sales overlap these channels and are not added.

Philippines. Worldpanel by Numerator (2025, pp. 20, 28) assigns 39% of household consumer packaged-goods purchases to sari-sari stores and 3% to stores inside markets. The 42–69% traditional-retail range adds none or all of grocery, drug, convenience, and residual channels with unspecified ownership. Conflicting quarter labels in the report prevent assigning a precise observation period.

South Africa. Niavas et al. (2019) report traditional stores receiving 30% of purchases by value in BCG's 2018 consumer survey. The category includes spazas, general dealers, specialists, and liquor outlets. Market-vendor coverage is not separately identified. Most respondents had household income below USD 500 per month. A separate spaza share is unavailable.

Vietnam. Worldpanel by Numerator (2025, pp. 22–23) reports 49% through street-side provision stores in Hanoi, Ho Chi Minh City, Da Nang, and Can Tho, and 71% in rural areas. Adding wet markets and traditional specialists gives 61% and 87%. The upper endpoints include all mini-stores, whose ownership is unspecified. These are packaged-goods shares rather than shares of all food purchases.

Table C1. Corner stores and traditional retailers: prevalence among businesses

Country	Corner-store definition	Corner stores			Traditional retailers	
		Count (year)	Share of businesses (%)	People per store	Count (year)	Share of businesses (%)
		(1)	(2)	(3)	(4)	(5)
Peru†	*Bodegas*	535,000 (2024 report)	9	64	1,370,000 (2024)	22
Bangladesh	General food stores	1,182,742 (2024)	10	147	4,345,826 (2024)	37
China	Independent convenience stores	6,000,000 (2020 report)	10	235	—	—
Colombia†	*Tiendas de barrio*	453,750 (2026 report)	11	118	1,100,000 (2023, 2024)	26
Egypt	Koshks	400,000 (2022 report)	10	282	1,778,543 (2022–23)	46
India	Kiranas	13,000,000 (2024 report)	18	112	19,900,000 (2023–24)	27
Indonesia	Warung grocery shops	3,600,000 (2022 report)	14	77	10,801,260 (2016)	41
Kenya†	Dukas (microretail-store proxy)	600,000 (2024 report)	≈8	94	4,140,000 (2016)	56
Mexico	Grocery and general stores	574,141 (2018)	12	≤217	1,972,021 (2018)	41
Philippines†	Sari-sari stores	≥1,300,000 (2024 report)	≈19	≤90	—	—
South Africa†	Spazas	≥150,000 (2025 report)	≳4	≤432	—	—
Vietnam	Neighborhood grocers	≥1,400,000 (2021)	≈28	≤71	1,700,000 (2020–21)	33

Notes: Corner stores exclude market stalls. Traditional-retailer counts cover food and nonfood retail and also include specialists, markets, and street sellers within the stated universe. Counts and shares are approximate where sources or coverage differ. † Business denominators are restricted to micro, small, or medium enterprises (MSMEs). People per store uses World Bank population for the count/report year or the latest available earlier year (World Bank, n.d.). "Report" indicates an unspecified observation year. Business denominators, sources, and assumptions appear in Section C.1.

Table C2. Corner stores and traditional retailers: shares of consumption and retail sales

Country	Corner-store definition	Consumption or sales denominator	Corner stores		Traditional retailers	
			Share (%)	Period	Share (%)	Period
		(1)	(2)	(3)	(4)	(5)
Peru	*Bodegas*	Fast-moving consumer goods	36	Q1 2025	>80	2024 report
Bangladesh	—	Total retail sales	—	—	97–98	2025 report
Brazil	General grocery stores	Food-at-home purchases	22	2017–18	45	2017–18
China	—	Grocery-retailer sales	—	—	60	2024
Colombia	*Tiendas de barrio*	Consumer goods sales	37	2026 report	60	2024
Egypt	—	Grocery-retailer sales	—	—	80	2024
India (urban)	Kirana/paan-beedi stores	Household consumer packaged goods	78	Q3 2025	78–86	Q3 2025
Indonesia (three cities)	Warung grocery shops	Household food expenditure	31	2010–11	≥66	2010–11
Kenya	—	Total retail sales	—	—	70	2020 report
Mexico	*Tiendas de abarrotes*	Food-at-home purchases	33	2022	84	2022
Nigeria	Small grocery and convenience stores	Food-retail sales	25	2024 report	97	2024 report
Philippines	Sari-sari stores	Household consumer packaged goods	39	2025 report	42–69	2025 report
South Africa	—	Surveyed consumer-product purchases	—	—	30	2018
Vietnam: urban	Street-side provision shops	Household consumer packaged goods	49	Q3 2025	61–73	Q3 2025
Vietnam: rural	Street-side provision shops	Household consumer packaged goods	71	Q3 2025	87–91	Q3 2025

Notes: Shares use the basket and geographic coverage specified for each country. Traditional-retail estimates approximate the broader channel. Format-based sources do not fully identify ownership. Asian-panel ranges vary the classification of uncertain channels while excluding supermarket, online, and direct-selling channels. They are not confidence intervals. These panels cover consumer packaged-goods acquisitions, potentially including gifts or samples, and the year ending Q3 2025 except the Philippines. Country sources and assumptions appear in Section C.2.

References for Appendix C

AAFC. Vietnamese market overview. Agriculture and Agri-Food Canada, Exporting to Vietnam: A Guide for Canadian Businesses, August 2022. 18 August; Nielsen grocery-store estimate as of February 2021.

Abdel Halim, Samar. Retail foods annual. GAIN Report EG2025-0010, U.S. Department of Agriculture, Foreign Agricultural Service, Cairo, Egypt, April 2025.

Agricultural Trade Office Beijing. Retail foods annual. GAIN Report CH2025-0193, U.S. Department of Agriculture, Foreign Agricultural Service, Beijing, China, September 2025.

BBS. Economic census 2024: National report, volume I, detailed statistical tables. Bangladesh Bureau of Statistics, May 2026. May 18; Table S1, economic units by four-digit industry class.

BPS. Hasil pendataan usaha/perusahaan perdagangan besar dan eceran, reparasi dan perawatan mobil dan sepeda motor: Sensus ekonomi 2016-lanjutan indonesia. Technical report, Badan Pusat Statistik, Jakarta, 2018.

BPS. Sensus ekonomi 2016: Banyaknya UMK dan UMB menurut wilayah dan lapangan usaha, Indonesia. Badan Pusat Statistik, SE2016-Lanjutan online tables, n.d. Table 21; accessed 24 October 2026.

CAPMAS. Final results of the sixth economic census, 2022/2023. Technical report, Central Agency for Public Mobilization and Statistics, Cairo, 2025. Results released 27 October 2025; English title from the report foreword.

Cheng, Mauricio. Consumo FMCG en hogares peruanos creció 7.1% en valor durante el primer trimestre del 2025. Kantar Worldpanel, June 2025.

China Economic Net. Chengxin youxuan launches a small-store strategy to support community shops. *People's Daily Online*, December 2020. 11 December; in Chinese.

ComexPerú. Las micro y pequeñas empresas en el Perú: Resultados en 2024. Technical report, Sociedad de Comercio Exterior del Perú, Lima, 2025.

DANE. Encuesta de micronegocios (EMICRON) 2024: Boletín técnico y anexos. Technical report, Departamento Administrativo Nacional de Estadística, Bogotá, July 2025a. 30 July; accompanying statistical annex, table B.2.

DANE. Proyecciones de población: Series nacional y departamental por área, con base en el CNPV 2018. Departamento Administrativo Nacional de Estadística, 2025b. National series 2018–2070, released 18 July; departmental series 2018–2050, released 8 August.

DANE. Censo económico nacional urbano 2024: Resultados preliminares. Departamento Administrativo Nacional de Estadística, interactive results database, 2026. Sectoral module with reference year 2023; accessed 24 October 2026.

Domínguez-Barreto, Ana Paula, Irene Farah, Nancy López-Olmedo, Carolina Pérez-Ferrer, Yenisei Ramírez-Toscano, Brent A. Langellier, M. Arantxa Colchero, Juan A. Rivera-Dommarco, Tonatiuh Barrientos-Gutiérrez, and Dalia Stern. Changes in the healthfulness of food and beverage purchases from 2006 to 2022 by outlet type in Mexico. *BMC Medicine*, 23:205, 2025. doi: 10.1186/s12916-025-04036-8.

DTI. DTI pilots "tindahan mo, e-level up mo!" to digitalize sari-sari stores in the Philippines. Department of Trade and Industry, Philippines, June 2024. 26 June.

Gómez Goyeneche, David. Las tiendas de barrio concentran 37% de las ventas de consumo masivo en el país. *La República*, August 2026.

IBGE. Pesquisa de orçamentos familiares 2017–2018: Microdados e documentação. Instituto Brasileiro de Geografia e Estatística, July 2023. Microdata, documentation and expenditure translators; 13 July release.

INEGI. Las empresas en los Estados Unidos Mexicanos: Censos económicos 2019. Technical report, Instituto Nacional de Estadística y Geografía, Aguascalientes, 2023. 2023 copyright edition; originally catalogued in 2020; reference year 2018.

INEGI. Las tiendas de conveniencia o minisúper en México. Instituto Nacional de Estadística y Geografía, Censos Económicos 2019, n.d. Reference year 2018; undated minimonograph.

INEI. Perú: Estructura empresarial, 2024. Technical report, Instituto Nacional de Estadística e Informática, Lima, November 2025.

Japan International Cooperation Agency and Deloitte Tohmatsu Financial Advisory LLC. Africa: Data collection survey for value chain analysis for retail business in Africa. final report. Technical Report IM JR 20-014, Japan International Cooperation Agency, August 2020.

Kantar Worldpanel. Descubre el comportamiento real de compra de los hogares peruanos en el canal tradicional. Kantar Worldpanel, November 2024.

KNBS. Micro, small and medium establishment (MSME) survey: Basic report. Technical report, Kenya National Bureau of Statistics, Nairobi, September 2016.

Le, Duy Binh. Defining legal status for household businesses in Viet Nam: Key to formalization. Technical report, International Labour Organization, Geneva, 2024.

Mahmud, Tanvir. Exporter guide annual. GAIN Report BG2025-0008, U.S. Department of Agriculture, Foreign Agricultural Service, Dhaka, Bangladesh, June 2025.

Mercado, Neil Arwin. Only around 500k MSMEs in PH seen to receive gov't aid in 2022. Inquirer.net, September 2021. September 10.

Minot, Nicholas, Randy Stringer, Wendy J. Umberger, and Wahida. Urban shopping patterns in Indonesia and their implications for small farmers. In *International Conference of Agricultural Economists*, Milan, Italy, 2015. International Association of Agricultural Economists. August 9–14; conference paper, June 15 version.

Modgil, Shweta. How Aksesmu is leveraging Google Cloud to help Indonesian warung owners digitalise their businesses. *ETCIO Southeast Asia*, March 2022. 3 March.

MoSPI. Annual survey of unincorporated sector enterprises (ASUSE) results for 2023–24. Ministry of Statistics and Programme Implementation, Government of India, Press Information Bureau, January 2025. 29 January; reference period October 2023–September 2024.

Motsomi, Abel. FinScope MSME survey South Africa 2024. FinMark Trust, results presentation, September 2024. 10 September.

MSEA. The MSEA agenda: January–march 2024. Micro and Small Enterprises Authority, 2024.

NBS. Communiqué on major data of the fourth national economic census: No. 2. Technical report, National Bureau of Statistics of China, Beijing, 2019.

Niavas, Stefano, Nomava Zanazo, and Charlotte Slabbert. Routes to growth in South Africa through traditional trade. Boston Consulting Group, August 2019. 28 August.

Office of Agricultural Affairs, Lagos. Retail foods annual. GAIN Report NI2024-0014, U.S. Department of Agriculture, Foreign Agricultural Service, Lagos, Nigeria, October 2024.

Redacción Gestión. Asociación de bodegueros: 2,600 bodegas agremiadas cerraron en 2023 por extorsiones. *Gestión*, March 2024. 21 March.

Reddy, Nevolan. Retail foods annual: South Africa. GAIN Report SF2025-0032, USDA Foreign Agricultural Service, Pretoria, November 2025. 19 November.

Salinas, Ana. Retail foods annual. GAIN Report CO2025-0027, U.S. Department of Agriculture, Foreign Agricultural Service, Bogota, Colombia, November 2025.

Sriram, M. and Dhwani Pandya. Indians get hooked on 10-minute grocery apps, squeezing small retailers. Reuters, June 2024. June 11.

Upadhyay, Ameya, Efayomi Carr, and Stella Klemperer. Digitizing the corner shop: Egypt spotlight. Technical report, Flourish Ventures, 2022. Research in partnership with 60 Decibels and MaxAB.

World Bank. From the last mile to the next mile: 2022 Vietnam poverty and equity assessment. Technical report, World Bank, Washington, DC, 2022.

World Bank. Population, total. World Development Indicators, indicator SP.POP.TOTL, n.d. Accessed 25 October 2026.

Worldpanel by Numerator. Asia Pulse: Asia's FMCG market at a glance. Q3 2025 edition. Technical report, Worldpanel by Numerator, December 2025.

D Survey protocol

This appendix documents the corner store survey protocol summarized in Section 4.1.1.

D.1 Sampling

The study excludes 24 of the entrant's operating territories. Those near the city center were excluded because their wealthier neighborhoods host sharply different retail dynamics, with fewer and larger stores. The rest were deemed too unsafe for enumerators, on the recommendation of the surveying firm and the entrant's sales supervisor.

Within each stratum, four stores were sampled from the control cluster and six from the treatment clusters, three from each. In the seven strata with four clusters, two stores per treatment cluster were sampled instead. The resulting 40% control share exceeds the control share of clusters, which improves precision in the control-only descriptive analyses behind Section 2.

D.2 Eligibility

Stores were eligible for the survey if they had at least two coolers, restricting the sample to stores above a minimum scale. This choice was meant to maximize power to detect treatment effects on take-up, because the entrant's sales agents tended to skip the smallest stores in the pilot, at least in the weeks after rollout. Coolers served as the screen because they are visible from outside, allowing enumerators to assess eligibility without engaging the owner. The screen was designed to exclude only the smallest stores. The screen is unlikely to explain the main heterogeneity pattern: treatment effects are largest among small stores in the observed sample, while treatment-induced increases in coolers are concentrated among the largest stores.

D.3 Field search

Enumerators received a map per cluster marking the selected starting points, drawn from the census. Each starting point began a search: the enumerator approached the nearest eligible store and surveyed it if it was open and the owner agreed to participate. If the attempt failed, the enumerator moved to the next-nearest store to the starting point, up to three blocks away, until a survey was completed. If no store could be surveyed within three blocks, the search was abandoned and a new starting point was drawn at random from the census. Failed attempts were rare and stemmed mostly from refusals, though enumerator misreporting cannot be fully ruled out despite various audits. Anecdotally, refusals reflected a sharp rise in local extortion before the survey, which made owners wary of enumerators despite their efforts to build trust. Other failures involved a closed store or an absent owner. In those cases enumerators were instructed to move to another search and return later, rather than replacing the store immediately, where possible.

D.4 Cluster assignment

In a few instances, navigation errors led enumerators to survey stores outside their assigned cluster: 25 stores were surveyed in a different study cluster than assigned, and five fell outside the study area altogether and are dropped, leaving the final sample of 1,183 stores. The main analyses assign each store to the cluster where it is actually located. Results are qualitatively unchanged when stores are instead assigned to the cluster of their intended starting point.

D.5 Survey rounds

The retail audit was originally assigned to a separate mystery shopper. After data from that enumerator were found to have been falsified, they were discarded and the audit was moved into the main second-visit survey. Valid retail outcomes are therefore unavailable for round 1.

D.6 Receipt collection

Respondents received a fixed payment of about 3 USD PPP at the first visit and a gift worth about 10 USD PPP at the second, together with a payment of either 0.3 or 0.6 USD PPP for each receipt produced at either visit. On average, stores produced receipts for 61% of their supplier orders, but that average masks substantial heterogeneity: 14% of stores with supplier orders produced no receipt at all and 30% produced one for every order, leaving 56% who verified some of their purchases but not all. The per-receipt payment amounts were randomized to test whether stores were producing fewer receipts than they had at hand, perhaps because the payoff did not compensate for the effort of searching for them or safely storing them between visits. The payment amount does not significantly affect receipt production (Table D3).

Table D1. Corner store survey rounds

Round	Month	Strata (#) (1)	Stores (#) (2)	Stores (%) (3)	Cum. (%) (4)
1	March	28	280	24	24
2	April	36	360	30	54
3	May	28	280	24	77
4	June	14	139	12	89
5	July–August	13	129	11	100
Total		119	1,188	100	

Notes: The unit of observation is the survey round. Strata were randomly assigned to rounds, run as roughly monthly batches that overlap at the edges. Store counts cover all 1,188 surveyed stores under their original cluster assignment.

Table D2. Supplier orders and receipt verification

	Observations	Mean	Share zero	p25	p50	p75	p90
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Supplier orders							
1st visit	1,183	4.95	0.040	3.00	4.00	6.00	9.00
2nd visit	1,153	7.29	0.037	3.00	5.00	9.00	15.00
Total	1,183	12.06	0.011	7.00	10.00	14.00	22.00
Panel B: Receipt-verified supplier orders							
1st visit	1,183	3.15	0.271	0.00	2.00	4.00	8.00
2nd visit	1,153	5.61	0.262	0.00	3.00	8.00	15.00
Total	1,183	8.62	0.148	2.00	5.00	12.00	21.00

Notes: The unit of observation is the store. Panel A counts supplier orders and Panel B those for which the store produced a receipt, reported for the first visit (one week of purchases), the second visit (the two weeks between visits), and the two combined. Purchases at wholesale markets and modern stores are excluded.

Table D3. Effect of the randomized per-receipt payment

	Supplier orders			Receipt-verified orders		
	1st visit (1)	2nd visit (2)	Total (3)	1st visit (4)	2nd visit (5)	Total (6)
High amount, 1st visit	0.03	0.37	0.55	0.11	0.30	0.56
	(0.17)	(0.59)	(0.72)	(0.18)	(0.66)	(0.81)
High amount, 2nd visit	—	0.18	0.44	—	0.21	0.47
		(0.50)	(0.64)		(0.54)	(0.71)
1st × 2nd	—	0.00	-0.37	—	0.18	-0.18
		(0.83)	(1.06)		(0.92)	(1.20)
p-value, jointly zero	—	0.826	0.792	—	0.777	0.705
Low-amount mean	4.99	6.92	11.55	3.08	5.15	7.94
Observations	1,183	1,153	1,183	1,183	1,153	1,183

Notes: The unit of observation is the store. Columns (1)–(3) count supplier orders and columns (4)–(6) those for which the store produced a receipt, for the first visit (one week of purchases), the second visit (the two weeks between visits), and the two combined. The per-receipt payment was randomized to a high or low amount at each visit (0.6 vs 0.3 USD PPP), independently across visits. The rows report the difference for the higher amount at the first visit, at the second visit, and their interaction, from a regression of the outcome on the two indicators and their interaction with stratum fixed effects. Columns (1) and (4) carry the first-visit term only. The row labeled p-value reports the test that the payment terms are jointly zero, and Low-amount mean is the mean among stores assigned the low amount at every visit entering the regression. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

E Digitization and product harmonization

This appendix describes how receipt photographs become the transaction-level dataset behind the variety and wholesale-price outcomes.

E.1 Digitization

Enumerators photographed every receipt produced during the two visits. Each photographed receipt was digitized line by line, recording the product description, the quantity purchased, and the amount paid. Receipt formats vary widely across suppliers—in how discounts, taxes, and free units appear—and digitization recorded these components explicitly. Free goods (6.2% of transactions) enter as transactions with a positive quantity and a zero amount. A small number of digitized lines with no amount recorded are dropped. Orders without a receipt enter the diary with their reported total but no product detail, so the transaction-level data cover receipt-verified purchases. Product-level detail is therefore available almost exclusively for supplier orders: market purchases have no digitized product detail, and purchases at modern stores account for only 0.9% of transactions.

E.2 The product catalog

Each digitized line is matched to a product catalog built for this project. The catalog assigns every item a category, a brand, and a product identity, where a product is a specific brand–size–flavor combination. The diary spans 2,806 distinct products across 495 brands and 64 categories. Coarser definitions are available: pooling flavors within brand and size yields 1,997 products, and pooling package sizes as well yields 1,212. The variety outcomes and the transaction-level regressions use the finest definition. Results are unchanged when products are defined at the coarser levels. Digitized lines that cannot be matched to a catalog entry are flagged as "other" products (2.4% of transactions). They are excluded from the variety decomposition and from wholesale-price and product-composition analyses. Counting them as distinct products in the variety outcomes leaves the estimates unchanged. Similarly, orders from suppliers the owner could not name are assigned to a catch-all "other" supplier (2.5% of supplier orders), and each such order counts as a separate supplier in supplier-count outcomes, except a handful of orders recorded without any supplier information, which are excluded.

E.3 Prices, quantities, and packages

Quantities are recorded in units, but stores buy in packages whose size varies across products. For each product, the package size is defined as the modal quantity purchased, and the package price is the unit price scaled to that modal package. Wholesale price and quantity outcomes are expressed per package throughout. Every reference quantity used in construction—modal package sizes and benchmark mean and median wholesale prices—is computed from control-group incumbent-supplier transactions, excluding the entrant and purchases at markets and modern stores, so no reference responds to treatment. Free goods count toward variety outcomes but are excluded from wholesale-price construction.

F Neighborhood exposure

This appendix examines whether the growth of small stores in treated clusters comes partly at the expense of nearby stores. I use differences in the treatment of neighboring clusters to estimate how access to the entrant nearby affects a store's wholesale purchases and product variety, holding its own cluster's treatment fixed. I first define exposure and explain the estimation procedure, then discuss the results and their implications for expenditure reallocation.

F.1 Measuring neighborhood exposure

Because treatment is assigned at the cluster level, the experiment cannot separate the effect of a store's own access to the entrant from the treatment of other stores in its cluster without additional assumptions. I first examine exposure to treatment in other clusters, defined as the number of treated corner stores within 450 meters but outside the store's own cluster. The own-cluster treatment effect includes both the store's access and the treatment of its neighbors inside the cluster. Section F.3 then discusses the all-neighbor specification used in the main text and the assumption it requires to separate these effects.

The 450-meter radius captures nearby stores while allowing exposure to vary with treatment assignment for most surveyed stores. I count neighbors using the INEI corner store census, including stores outside the survey sample. Within this radius, surveyed stores have an average of 55 neighbors outside their own cluster, and 10% have no variation in exposure conditional on their own assignment (Table F1). A 300-meter radius focuses on closer competitors but raises the share without exposure variation to 23%. I report estimates using this smaller radius and a measure that gives greater weight to closer stores as robustness checks.

F.2 Identification and estimation

The number of treated stores nearby depends on both local store density and the experimental design. First, as Borusyak and Hull (2023) note, stores in denser areas have more potential treated neighbors, and these locations may differ in demand or other determinants of store outcomes. Second, the randomization protocol in this specific setting also creates a systematic exposure difference between treated and control stores. Because each stratum has exactly one control cluster, a store's assignment changes the treatment probabilities of neighboring stores in the same stratum. Being in the control group means all other clusters in that stratum are treated with certainty, whereas being in a treated cluster means the treatment probability in other same-stratum clusters is less than one. Control stores have an average of 41 treated neighbors outside their cluster, compared with 27 for treated stores (Table F1). Estimating spillovers therefore requires separating the variation generated by treatment assignment from these systematic differences in expected exposure.

I isolate the random variation by subtracting each store's expected exposure from its realized exposure, following Borusyak and Hull (2023). Expected exposure is the average number of treated neighbors across possible treatment draws the experiment could have

produced. I account for the experimental design issues described above by fixing the focal store's own assignment to compute expected exposure. This means that for control clusters, the only possible variation comes from neighboring stores in other strata. This *recentered exposure* measures whether a store happens to have more or fewer treated neighbors than expected under the experimental design. I use recentered exposure as an instrument for the actual number of treated neighbors. The instrument explains 42% of the variation in realized exposure outside the store's own cluster after accounting for own treatment and stratum fixed effects (Table F1).

I estimate the effect of neighboring treatment jointly with the effect of a store's own-cluster assignment by two-stage least squares:

$$Y_{ics} = \tau T_c + \beta X_i + \gamma_s + \varepsilon_{ics}. \tag{8}$$

Here Y_{ics} is the outcome for store i in cluster c and stratum s, T_c indicates own-cluster treatment, X_i is realized exposure, and γ_s denotes stratum fixed effects. Exposure is centered on its sample mean and expressed in standard deviations of the instrument, equivalent to 12.7 treated stores. Thus, β measures the effect of a one standard deviation increase in exposure, holding own treatment fixed. The coefficient τ measures the effect of treating the store's own cluster, holding outside exposure fixed.

To examine whether spillovers differ by store size, I interact both own treatment and exposure with the large-store indicator and include stratum-by-size fixed effects, as in the main analysis. Within each size group, I assume that exposure has the same effect on treated and control stores. Standard errors are clustered by stratum. Tables F2–F4 report the specifications for product variety, product availability and retail prices, and wholesale expenditure.

F.3 Results and interpretation

The results provide suggestive evidence that small stores buy less when more nearby stores in other clusters receive access to the entrant, consistent with business stealing across cluster boundaries. A one standard deviation increase in exposure reduces weekly wholesale expenditure by 12 USD PPP, or 5% of the control mean ($p = 0.043$; Table F4). The decline comes from purchases with incumbent suppliers, which fall by 13 USD PPP ($p = 0.015$). Small stores also buy 2.6 fewer products, an effect significant at the 10% level ($p = 0.054$; Table F2). The category estimate points in the same direction, at 0.5 fewer categories, although it is less precise ($p = 0.195$). The corresponding estimates for large stores are imprecise, as are the differences between the two size groups.[44]

These patterns suggest that part of the original treatment effects among small stores reflect losses at control stores exposed to treated competitors. Intuitively, control stores have more treated neighbors outside their own cluster by design, so the estimated negative exposure effect implies greater losses at control stores. Holding exposure fixed, the small-store effect on weekly wholesale expenditure falls from 46 to 34 USD PPP ($p = 0.013$ for

[44] Exposure has no statistically significant effect on availability of the six popular products or the retail price index for either size group (Table F3).

the latter). The corresponding effect on product variety falls from 5.5 to 2.8 products and becomes imprecise ($p = 0.336$).

The all-neighbor specification used in the main text separates own access from treatment of neighbors inside the cluster by assuming that exposure to treated neighbors has the same effect within and across cluster boundaries. Under this assumption, variation in treatment across cluster boundaries also identifies how treatment of neighbors inside the cluster affects a store. I count all treated census stores within 450 meters, including those in the store's own cluster, and instrument this measure using its deviation from expected exposure.

The all-neighbor estimates then suggest that small stores' gains from own treatment are partly offset by treatment of nearby stores (Table A15). Holding total measured exposure fixed, own treatment is estimated to increase small stores' weekly wholesale expenditure by 82 USD PPP ($p = 0.002$) and the number of products purchased by 13.1 ($p = 0.011$). Treating neighbors inside the same cluster offsets 49 USD PPP of this gain, leaving an implied cluster effect of 33 USD PPP ($p = 0.019$). The implied cluster effect on product variety is 2.4 products ($p = 0.436$). These are close to the 34 USD PPP and 2.8 product estimates under external exposure.

The negative small-store estimates also persist under alternative measures of proximity. Using a 300-meter radius or giving greater weight to closer stores yields similar reductions in product variety and weekly wholesale expenditure, although the expenditure estimates are less precise (Table F5, using external exposure).[45] The all-neighbor specification gives qualitatively similar results: small-store exposure effects remain negative with a 300-meter radius ($p = 0.140$ for expenditure and $p = 0.111$ for products) and distance weights ($p = 0.166$ and $p = 0.075$, respectively). The large-store effects remain imprecisely estimated.

These results provide suggestive evidence that gains at treated stores are accompanied by losses at nearby small stores. Accounting for differences in neighborhood exposure reduces the estimated small-store treatment effects, although the effects on wholesale expenditure and products purchased remain positive (though the latter loses statistical precision). These comparisons do not aggregate gains and losses across all stores, so they do not establish whether the wholesaler's entry expands total market activity.

[45] The distance-weighted measure includes stores within 450 meters, weighted by $\exp(-d/150)$ for distance d in meters. Each specification reports exposure effects per standard deviation of its own instrument.

Table F1. Neighborhood exposure and identifying variation

	450 m count (1)	300 m count (2)	450 m weighted (3)
Panel A: Neighborhood exposure			
Mean neighboring stores	54.77	18.40	7.50
Mean treated neighbors: control stores	40.94	13.62	5.57
Mean treated neighbors: treated stores	26.71	8.67	3.57
Share with no neighboring stores	0.008	0.056	0.008
Share with no exposure variation	0.101	0.232	0.100
Panel B: Instrument and first stage			
Instrument standard deviation	12.71	5.97	2.10
First-stage coefficient	1.06	1.03	1.05
	(0.05)	(0.05)	(0.06)
First-stage partial R-squared	0.42	0.39	0.41
First-stage F-statistic	387.85	378.21	354.64
Observations	1,183	1,183	1,183

Notes: The unit of observation is the store. Neighbors are census corner stores outside the surveyed store's own cluster. Column (1) uses a 450-meter radius, column (2) a 300-meter radius, and column (3) a 450-meter radius with neighbors weighted by $\exp(-d/150)$, where d is distance in meters. Recentered exposure is realized exposure minus its expectation over the randomization given the store's own assignment. Exposure has no variation when it cannot change across treatment assignments consistent with the surveyed store's own assignment. Both shares use all surveyed stores as their denominator. The first stage uses all surveyed stores and regresses realized exposure on recentered exposure and own treatment, with stratum fixed effects. Realized and recentered exposure are expressed in standard deviations of the instrument. The F-statistic tests that the coefficient on recentered exposure is zero, using a covariance matrix clustered by stratum. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024).

Table F2. Effect of own treatment and neighborhood exposure on product variety

	Products (#) (1)	Categories (#) (2)	Brands/ category (3)	Products/ brand (4)
Panel A: Full sample				
Treatment	1.30	0.47	0.04	-0.02
	(3.20)	(0.61)	(0.04)	(0.04)
Exposure	-1.67	-0.34	-0.02	-0.00
	(1.42)	(0.29)	(0.02)	(0.02)
Control mean	36.96	11.48	1.69	1.69
Observations	996	996	996	996
Panel B: Heterogeneity by store size				
Treatment × small	2.82	0.78	0.09*	-0.06
	(2.92)	(0.78)	(0.05)	(0.05)
Treatment × large	0.57	0.21	0.00	-0.01
	(6.25)	(1.02)	(0.06)	(0.06)
Exposure × small	-2.59*	-0.48	-0.02	-0.01
	(1.33)	(0.37)	(0.02)	(0.02)
Exposure × large	-0.11	-0.11	-0.01	-0.01
	(2.57)	(0.47)	(0.03)	(0.03)
Control mean (small)	28.36	9.79	1.61	1.69
Control mean (large)	45.34	13.15	1.76	1.70
p-value, treatment equal	0.741	0.656	0.246	0.486
p-value, exposure equal	0.365	0.526	0.635	0.978
Observations	963	963	963	963

Notes: The unit of observation is the store. The sample is the stores with at least one receipt-verified supplier purchase. The outcomes are the same as in Table 2. Exposure is the number of treated census corner stores within 450 meters and outside the store's own cluster, in standard deviations of its recentered instrument (12.7 stores), and is instrumented by that recentered measure. Panel A reports the coefficients on the store's own treatment and on exposure from a regression with stratum fixed effects. Panel B interacts both with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the rows labeled p-value test equality of the two treatment effects and of the two exposure effects. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F3. Effect of own treatment and neighborhood exposure on product availability and retail prices

	Availability		Retail prices		
	Any missing	Share missing	Price index (log)	Share below mode	Share above mode
	(1)	(2)	(3)	(4)	(5)
Panel A: Full sample					
Treatment	-0.066**	-0.012	0.000	-0.004	-0.003
	(0.030)	(0.008)	(0.002)	(0.007)	(0.012)
Exposure	-0.008	-0.000	0.001	-0.004	0.005
	(0.016)	(0.004)	(0.001)	(0.004)	(0.006)
Control mean	0.337	0.070	0.000	0.137	0.232
Observations	880	880	880	880	880
Panel B: Heterogeneity by store size					
Treatment × small	-0.046	-0.012	0.005	-0.028***	0.000
	(0.044)	(0.013)	(0.003)	(0.010)	(0.016)
Treatment × large	-0.108**	-0.011	-0.003	0.018	-0.002
	(0.048)	(0.012)	(0.003)	(0.013)	(0.020)
Exposure × small	0.004	0.001	0.001	-0.006	-0.004
	(0.021)	(0.005)	(0.001)	(0.005)	(0.007)
Exposure × large	-0.015	0.001	0.002	-0.001	0.020**
	(0.028)	(0.007)	(0.002)	(0.006)	(0.009)
Control mean (small)	0.429	0.095	-0.003	0.188	0.255
Control mean (large)	0.223	0.041	0.003	0.070	0.194
p-value, treatment equal	0.329	0.962	0.081	0.009	0.942
p-value, exposure equal	0.582	0.948	0.571	0.530	0.042
Observations	857	857	857	857	857

Notes: The unit of observation is the store. The sample is the stores covered by the retail audit. The outcomes are the same as in Table 3. Exposure is the number of treated census corner stores within 450 meters and outside the store's own cluster, in standard deviations of its recentered instrument (12.7 stores), and is instrumented by that recentered measure. Panel A reports the coefficients on the store's own treatment and on exposure from a regression with stratum fixed effects. Panel B interacts both with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the rows labeled p-value test equality of the two treatment effects and of the two exposure effects. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F4. Effect of own treatment and neighborhood exposure on wholesale expenditure

	All purchases				Receipt-verified	
	Total exp.	Entrant	Incumbent suppliers	Markets/ modern stores	Entrant	Incumbent suppliers
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full sample						
Treatment	7.61	5.21**	-0.01	2.41	2.58	0.41
	(21.00)	(2.54)	(19.58)	(8.27)	(2.13)	(18.87)
Exposure	-4.17	-2.14	-5.02	2.99	-2.32	-2.98
	(10.00)	(1.58)	(8.35)	(3.26)	(1.47)	(7.89)
Control mean	323.07	7.88	266.98	48.20	4.79	179.71
Observations	1,183	1,183	1,183	1,183	1,183	1,183
Panel B: Heterogeneity by store size						
Treatment × small	34.11**	7.11***	23.65*	3.35	4.93***	23.98
	(13.56)	(2.54)	(13.58)	(7.58)	(1.83)	(14.96)
Treatment × large	-5.72	4.26	-25.76	15.78	1.89	-29.08
	(44.96)	(5.65)	(40.98)	(15.87)	(5.25)	(38.77)
Exposure × small	-11.90**	-1.53	-13.14**	2.77	-1.42*	-10.57
	(5.81)	(1.07)	(5.35)	(2.99)	(0.82)	(6.55)
Exposure × large	4.23	-2.12	1.90	4.46	-3.09	3.58
	(18.14)	(2.47)	(15.83)	(6.41)	(2.37)	(15.65)
Control mean (small)	233.03	4.93	196.42	31.68	2.59	107.06
Control mean (large)	419.68	11.80	339.88	68.00	7.71	261.12
p-value, treatment equal	0.397	0.647	0.250	0.480	0.585	0.207
p-value, exposure equal	0.378	0.819	0.349	0.808	0.504	0.386
Observations	1,150	1,150	1,150	1,150	1,150	1,150

Notes: The unit of observation is the store. The outcomes are the same as in Table 4. Monetary values are in USD PPP. Exposure is the number of treated census corner stores within 450 meters and outside the store's own cluster, in standard deviations of its recentered instrument (12.7 stores), and is instrumented by that recentered measure. Panel A reports the coefficients on the store's own treatment and on exposure from a regression with stratum fixed effects. Panel B interacts both with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the rows labeled p-value test equality of the two treatment effects and of the two exposure effects. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table F5. Neighborhood exposure under alternative distance measures

	Products (#) (1)	Total exp. (2)
Panel A: 450 m radius		
Exposure	-1.67	-4.17
	(1.42)	(10.00)
Exposure × small	-2.59*	-11.90**
	(1.33)	(5.81)
Exposure × large	-0.11	4.23
	(2.57)	(18.14)
Panel B: 300 m radius		
Exposure	-1.67	-5.91
	(1.39)	(9.56)
Exposure × small	-2.44	-9.52
	(1.51)	(6.35)
Exposure × large	-0.82	-4.31
	(2.69)	(19.70)
Panel C: Distance-weighted, 450 m radius		
Exposure	-1.60	-5.59
	(1.40)	(10.51)
Exposure × small	-2.45*	-8.37
	(1.34)	(5.94)
Exposure × large	-0.40	-3.48
	(2.88)	(21.76)
Observations (full)	996	1,183
Observations (by size)	963	1,150

Notes: The unit of observation is the store. Column (1) is the number of distinct products bought over receipt-verified supplier transactions, over the stores with at least one receipt-verified supplier purchase; column (2) is weekly wholesale expenditure from every source, including wholesale markets and modern stores. Monetary values are in USD PPP. Panels A, B and C define exposure as treated census corner stores outside the store's own cluster within 450 meters, within 300 meters, and within 450 meters weighted by distance, each in standard deviations of its own recentered instrument and instrumented by it. In each panel, the Exposure row is from a regression of the outcome on the store's own treatment and exposure with stratum fixed effects, and the rows for small and large stores are from one that interacts both with an indicator for large stores and includes stratum-by-size fixed effects. Treatment coefficients are omitted. Observation counts apply to all panels. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

G Self-reported measures

G.1 Wholesale purchases and reported sales

Self-reported sales show little response to the wholesaler's entry, despite a substantial increase in small stores' wholesale purchases. The wholesale diary measures imply an increase of approximately 20% in small stores' weekly expenditure (Table 4), without a corresponding reduction in retail prices (Table 3). With approximately stable markups, this expansion should translate into a similar percentage increase in sales. Yet the self-reported sales estimates are small and imprecise (Tables A13 and A14). This appendix examines whether difficulties measuring financial outcomes can account for this discrepancy, using the wholesale diary to assess the accuracy of owners' reports.

The additional purchases are unlikely to reflect temporary inventory accumulation. The survey took place nine to thirteen months after rollout, when take-up had stabilized, and stores typically replenish products weekly. Table G1 shows no significant effect on capacity utilization: the estimated increase is 0.2 percentage points overall ($p = 0.795$), 0.8 among small stores ($p = 0.509$), and 1.2 among large stores ($p = 0.298$). Further, I find no evidence of increases in waste for either store size (it falls by 0.8 USD PPP per week among large stores, $p = 0.079$). Together, these patterns support the interpretation of the purchase response as a sustained expansion in store activity.

G.2 Reliability of reported financial outcomes

Obtaining precise financial reports from store owners proved difficult during fieldwork. Recall and reporting errors are common in surveys of small businesses (de Mel et al., 2009). In this setting, owners were also reluctant to disclose financial information because of distrust and concerns about extortion and government audits.[46] Across completed interviews, 5.6% of sales questions and 3.8% of purchase questions went unanswered. The responses obtained are heavily rounded: 98% of sales reports and 95% of purchase reports are multiples of S/50, while 89% and 78%, respectively, are multiples of S/100. These increments represent approximately 28 and 56 USD PPP.

Reported sales are also implausibly high relative to reported purchases. A retail markup of 30% (as is anecdotally reported both by sales agents and store owners) implies sales of 1.30 for each unit spent on wholesale purchases. Instead, the median within-store ratio is 2.6, with an interquartile range of 1.6–3.8 (Figure G1). While purchase timing can generate differences between sales and purchases in a particular week, frequent replenishment makes it difficult to explain such a large and systematic discrepancy through timing alone.

The wholesale diary allows a direct check of reporting accuracy by comparing each owner's stated weekly purchases with expenditure calculated from individual orders. I find that diary purchases explain only 39% of the variation in reported purchases among control stores and 38% among treated stores (Table G2). The discrepancies run in both directions: 51% of stores report purchases that differ from the diary by more than a factor of 1.5, and 27% by

[46] Iacovone and McKenzie (2022) describe similar difficulties among retailers in Bogotá.

more than a factor of two (Figure G2). A uniform adjustment to reported purchases would therefore leave substantial errors at individual stores.

Table G3 relates these discrepancies to the store, owner, and local characteristics used in the balance tests. The outcomes are the log ratio of reported purchases to purchases in the wholesale diary and its absolute value, measuring the direction and magnitude of disagreement, respectively. The characteristics are jointly associated with both outcomes (joint p-values: < 0.001 and 0.043), but explain only 5.2% and 2.4% of residual variation after accounting for treatment assignment and missing-value indicators.[47]

I argue that the reported-sales estimates are too noisy to reliably detect the increase implied by the wholesale diary. Applying a 30% markup, the treatment effects on wholesale expenditure among small stores imply an increase in weekly sales of approximately 60 USD PPP. Using the estimated standard error of the small-store sales effect, I calculate that the minimum effect detectable with 80% power at the 5% significance level is 98 USD PPP.[48] The estimated sales effect of 30 USD PPP has a confidence interval that includes the diary-implied gain. The absence of a statistically significant sales effect therefore does not rule out growth of this magnitude.

G.3 Treatment effects and reporting accuracy

I test whether the sales estimates change among stores whose reported purchases agree more closely with the diary, which may indicate more accurate reporting. This is a scoping exercise because the restrictions depend on outcomes that treatment may affect.

I consider two restrictions: reported purchases within a factor of two of diary purchases, and within a factor of 1.5. Inclusion in either restricted sample is similar across treatment and control stores. The estimated increase in small stores' weekly sales rises from 30 USD PPP in the full sample to 57 and 70 USD PPP in the two restricted samples (Figure G3). Since the tighter restrictions also reduce the sample, the estimates are not precise enough to reach conventional significance levels, at p = 0.150 and 0.163, respectively. Nevertheless, the small-store sales effects are much closer to the increase implied by the diary. The large-store sales estimates move in the opposite direction, from −53 USD PPP in the full sample to −160 and −322 USD PPP in the two restricted samples (p = 0.284 and 0.136, respectively). These opposite movements are consistent with reallocation toward small stores, although the estimates remain imprecise.

The estimates are also sensitive to outliers, particularly among large stores. Winsorizing outcomes at the 98th percentile of their control-group distribution changes the large-store sales estimate from −53 to 2 USD PPP per week and substantially reduces its standard error (Figure G3). Across the restricted samples, both winsorized and log estimates become more positive for small stores and more negative for large stores, but remain imprecise.

The diary-based results remain qualitatively similar across subsamples for both size groups, with growth concentrated among small stores (Table G4). Small stores' whole-

[47] Results are similar among stores with complete characteristics. With stratum fixed effects, the corresponding shares are 2.0% and 2.0%, with joint p-values of 0.091 and 0.239, respectively.

[48] The calculation uses the standard error of the small-store treatment effect from the main specification, clustered by stratum, with a two-sided test.

sale expenditure rises by 46, 31, and 34 USD PPP across the three samples, and product variety by 5.5, 3.8, and 5.1 products. Large-store estimates are not distinguishable from zero in any subsample. The similar pattern of diary-based responses supports the interpretation that closer agreement identifies more informative reports.

Table G1. Effect of entrant access on waste and capacity utilization

	Waste value (1)	Capacity utilization (2)
Panel A: Full sample		
Treatment	-0.30	0.002
	(0.36)	(0.009)
Control mean	3.20	0.830
Observations	1,183	1,183
Panel B: Heterogeneity by store size		
Treatment × small	-0.02	0.008
	(0.57)	(0.012)
Treatment × large	-0.84*	0.012
	(0.48)	(0.012)
Control mean (small)	3.77	0.769
Control mean (large)	2.60	0.896
Observations	1,150	1,150

Notes: The unit of observation is the store. Column (1) is the self-reported weekly value of wasted products, averaged across survey visits. Column (2) is the share of storage capacity in use. Monetary values are in USD PPP. Panels A and B and the standard errors are as in Table A13. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table G2. Self-reported purchases regressed on the wholesale diary

	Control		Treatment	
	Levels	Ranks	Levels	Ranks
	(1)	(2)	(3)	(4)
Panel A: All reported orders				
Coefficient	0.615***	0.550***	0.657***	0.582***
	(0.096)	(0.037)	(0.056)	(0.037)
R^2	0.387	0.303	0.379	0.339
Observations	462	462	698	698
Panel B: Receipt-verified orders				
Coefficient	0.592***	0.555***	0.656***	0.494***
	(0.108)	(0.038)	(0.069)	(0.039)
R^2	0.319	0.309	0.337	0.245
Observations	394	394	599	599

Notes: The unit of observation is the store. The dependent variable is self-reported purchases, the answer to a single question about spending over the last week, averaged across visits. The regressor is weekly purchases computed from the wholesale diary collected in the same interviews, in Panel A over all reported orders and in Panel B over the orders for which the store produced a receipt. Panel B is restricted to stores with at least one receipt-verified order. Columns (2) and (4) replace both variables by their ranks within the column's estimation sample. Monetary values are in USD PPP. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table G3. Observable correlates of reporting discrepancies

	Signed discrepancy		Absolute discrepancy	
	Coefficient	SE	Coefficient	SE
	(1)	(2)	(3)	(4)
Panel A: Local characteristics				
Corner stores	-0.048	(0.034)	-0.044*	(0.024)
Road density	-0.031	(0.029)	-0.010	(0.019)
Distance to major road	0.036	(0.030)	0.010	(0.020)
Distance to market	-0.011	(0.036)	-0.005	(0.022)
Slope	-0.048	(0.031)	-0.009	(0.018)
Altitude	0.064	(0.042)	0.011	(0.022)
Population	0.128***	(0.043)	0.051*	(0.027)
Mid SES	-0.017	(0.032)	-0.023	(0.022)
High SES	0.009	(0.037)	-0.009	(0.025)
Panel B: Store and owner characteristics				
Female	-0.047	(0.051)	0.017	(0.034)
Age	-0.038*	(0.021)	-0.021	(0.016)
Higher education	0.090	(0.054)	0.028	(0.044)
Has partner	-0.028	(0.048)	-0.038	(0.030)
Household size	-0.065***	(0.024)	-0.043**	(0.017)
Store at home	-0.053	(0.086)	-0.065	(0.059)
Years in business	0.029	(0.023)	0.013	(0.016)
Floor area	-0.052*	(0.027)	0.012	(0.016)
Treatment assignment	-0.067*	(0.040)	-0.046	(0.033)
Observations	1,154		1,154	
Outcome mean	0.172		0.535	
Joint test (p-value)	< 0.001		0.043	
Partial R^2	0.052		0.024	

Notes: The unit of observation is the store. The sample is restricted to stores with positive purchases in both measures. Columns (1) and (2) report coefficients and standard errors for the log ratio of reported weekly purchases to weekly purchases in the wholesale diary; columns (3) and (4) use its absolute value. Characteristics are defined as in Table A3. Continuous characteristics are standardized using their control-group mean and standard deviation; binary characteristics remain unscaled. Both regressions include all characteristics in Panels A and B and treatment assignment. Missing characteristics are set to zero after standardization, with missing-value indicators included as controls. The joint test is that the seventeen characteristic coefficients are zero, excluding treatment assignment and missing-value indicators. Partial R^2 is the additional share of residual variation explained by these characteristics, conditional on treatment assignment and missing-value indicators. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table G4. Effect of entrant access on wholesale expenditure and product variety by self-report agreement

	All stores		Within 2×		Within 1.5×	
	Wholesale exp.	Products (#)	Wholesale exp.	Products (#)	Wholesale exp.	Products (#)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full sample						
Treatment	12.18	3.16	-6.82	0.07	-31.51	-3.64
	(19.38)	(2.61)	(20.22)	(3.16)	(25.50)	(4.41)
Control mean	323.07	36.96	340.41	38.58	352.11	38.59
Observations	1,183	996	841	723	569	481
Panel B: Heterogeneity by store size						
Treatment × small	46.22***	5.50**	31.01**	3.82	33.61*	5.12
	(13.73)	(2.61)	(13.67)	(2.91)	(19.28)	(3.96)
Treatment × large	-11.28	0.71	8.45	-0.46	-63.22	-11.90
	(40.36)	(4.77)	(43.82)	(6.26)	(58.93)	(8.91)
Control mean (small)	233.03	28.36	255.52	28.30	270.73	27.48
Control mean (large)	419.68	45.34	398.83	47.08	402.99	47.26
p-value, equal effects	0.179	0.382	0.625	0.535	0.123	0.078
Size (cont.) interaction: coef	-71.86	-6.50	-33.10	-6.86	-65.18	-10.79
Size (cont.) interaction: p-val	0.020	0.119	0.453	0.236	0.235	0.136
Observations	1,150	963	801	677	523	428

Notes: The unit of observation is the store. In columns (1), (3) and (5) the outcome is weekly wholesale expenditure from every source, including wholesale markets and modern stores. In columns (2), (4) and (6) the outcome is the number of distinct products bought over receipt-verified supplier transactions, where a product is a brand–size–flavor combination, and the sample is the stores with at least one receipt-verified supplier purchase. Columns (3) and (4) restrict to the stores whose self-reported weekly purchases are within a factor of two of their wholesale diary total, and columns (5) and (6) to those within a factor of 1.5. Stores with a diary total of zero or a missing self-report are excluded from columns (3)–(6). Monetary values are in USD PPP. Panel A reports the coefficient on cluster-level treatment assignment from a regression with stratum fixed effects. Panel B interacts treatment with an indicator for large stores (floor area above the control-group median) and includes stratum-by-size fixed effects; the row labeled p-value tests equality of the two effects. The rows labeled Size (cont.) interaction: coef and Size (cont.) interaction: p-val report the change in the treatment effect per standard deviation of floor area (standardized on the control distribution) and its p-value, from the Panel A specification adding floor area and its interaction with treatment. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

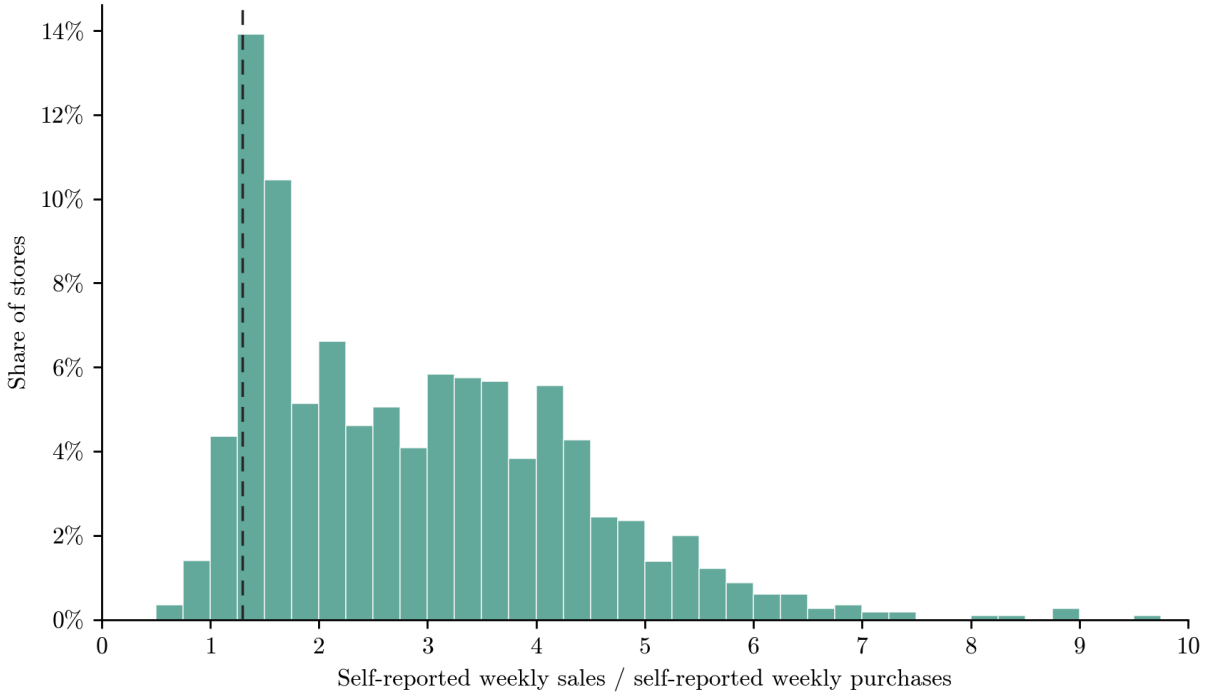
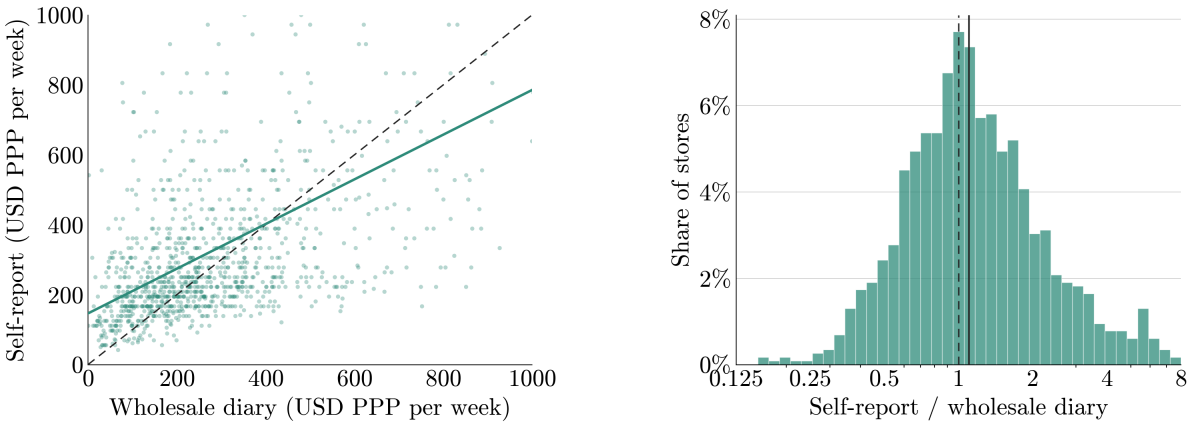


Figure G1. Reported sales relative to reported purchases

Notes: This figure shows the distribution of self-reported weekly sales divided by self-reported weekly purchases across 1,149 stores with observed sales and positive purchases. Both measures cover non-fresh products and average weekly reports across interviews. The dashed black line at 1.30 marks the ratio implied by a 30% retail markup. The horizontal axis runs from zero to ten, omitting two stores (0.2%) with ratios above ten. Bar heights are percentages of all 1,149 stores.



Panel A: Self-report against diary

Panel B: Distribution of their ratio

Figure G2. Self-reported purchases against the wholesale diary

Notes: This figure compares self-reported weekly purchases with the wholesale diary collected in the same interviews, pooling the treatment and control arms. Panel A plots one against the other and Panel B the distribution of their ratio. In both panels the dashed line marks perfect agreement; Panel A's solid line is fitted by least squares and Panel B's marks the median ratio. Self-reported purchases are the answer to a single question about spending over the last week, averaged across visits, and the diary is the weekly average of every order recorded. Both are in USD PPP per week, over the 1,154 stores with a positive value for both. 7% of stores fall outside the range plotted in Panel A and 1% outside Panel B.

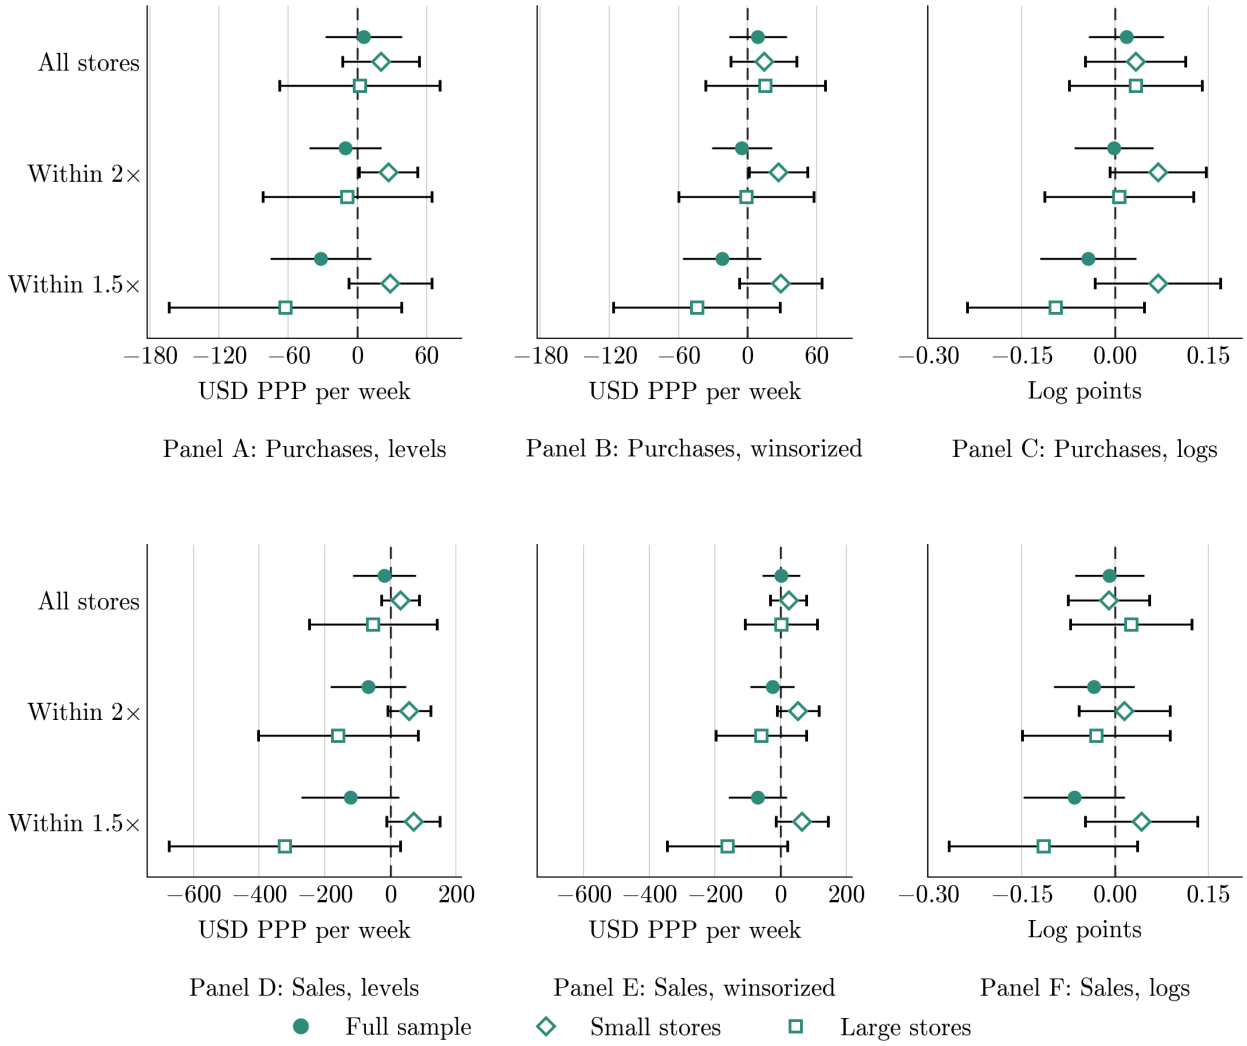


Figure G3. Effect of entrant access on self-reported store performance by self-report agreement

Notes: This figure shows store-level treatment effects on self-reported weekly purchases (Panels A–C) and weekly sales (Panels D–F), in USD PPP levels (Panels A and D), winsorized at the 98th percentile of their own control distribution (Panels B and E), or logs (Panels C and F). Each is estimated on three samples: all stores, the stores whose self-reported weekly purchases are within a factor of two of their wholesale diary total, and those within a factor of 1.5. Stores with a diary total of zero or a missing self-report are excluded from the two restricted samples. In the log panels, stores with a zero drop out. Filled circles are treatment effects from a regression of the outcome on cluster-level treatment assignment with stratum fixed effects. Open diamonds and squares are the effects for small and large stores from a regression that interacts treatment with an indicator for large stores, defined by the control-group median of store floor area, and includes stratum-by-size fixed effects. Horizontal bars are 90% confidence intervals using standard errors clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). The three samples have 1,183, 841 and 571 stores, respectively.

H Model derivations

This appendix derives the retail equilibria used in the main text, verifies the supplier's optimal wholesale prices and distribution choices before and after the wholesaler's entry (pre-entry and post-entry, for short), and provides the welfare comparisons. Throughout, I maintain Assumptions 1 and 2 unless stated otherwise.

H.1 Retail demand

This section derives the retail responses that determine supplier demand under one and two active stores.

One active store. If only store j is active, its demand is

$$q_j(p_j) = \max\left\{0, \min\left\{1, \frac{\nu - p_j}{\tau}\right\}\right\}.$$

When the store serves only part of the market,

$$\pi_j(p_j) = (p_j - w - h_j)\frac{\nu - p_j}{\tau},$$

so its optimal price and quantity are

$$p_j = \frac{\nu + w + h_j}{2}, \qquad q_j = \frac{\nu - w - h_j}{2\tau}.$$

The store instead serves the full market when

$$w \leq \nu - h_j - 2\tau,$$

in which case it charges

$$p_j = \nu - \tau, \qquad q_j = 1.$$

Because $h_L < h_S$, the relevant one-store outcome always has L active.

Two active stores with a common wholesale price. When both stores receive supply, store j faces demand, with ℓ denoting the other store,

$$q_j(p_j, p_\ell) = \max\left\{0, \min\left\{1, \frac{\nu - p_j}{\tau}, \frac{\tau + p_\ell - p_j}{2\tau}\right\}\right\}.$$

The outside option and the competing store each limit its customer region. If both stores are active while a gap of non-buyers separates their customer regions, each behaves as a local monopolist:

$$p_j = \frac{\nu + w + h_j}{2}, \qquad q_j = \frac{\nu - w - h_j}{2\tau}.$$

Hence total quantity is

$$q_L + q_S = \frac{2\nu - 2w - h_L - h_S}{2\tau}.$$

The customer regions just cover the full line at

$$w^{LS} = \nu - \tau - \frac{h_L + h_S}{2}.$$

At this wholesale price,

$$q_L^{LS} = \frac{1}{2} + \frac{h_S - h_L}{4\tau}, \qquad q_S^{LS} = \frac{1}{2} - \frac{h_S - h_L}{4\tau}.$$

Assumption 1 guarantees $q_S^{LS} > 0$. The associated retail prices satisfy $p_j^{LS} = \nu - \tau q_j^{LS}$ and $p_j^{LS} - w^{LS} - h_j = \tau q_j^{LS}$. At this just-covered boundary, a price cut attracts buyers from the rival at rate $1/(2\tau)$, while a price increase loses buyers to the outside option at rate $1/\tau$. The left- and right-hand profit derivatives are therefore

$$q_j^{LS} - \frac{p_j^{LS} - w^{LS} - h_j}{2\tau} = \frac{q_j^{LS}}{2} > 0, \qquad q_j^{LS} - \frac{p_j^{LS} - w^{LS} - h_j}{\tau} = 0.$$

Where demand and margins are positive, profit derivatives are nonincreasing within each branch and jump downward at transitions between branches. Prices with zero demand or nonpositive margins cannot improve on positive profit, so these conditions verify global best responses.

More generally, if store j pays wholesale price w_j, a just-covered equilibrium has $p_j = \nu - \tau q_j$ and the one-sided pricing conditions require

$$\tau q_j \le p_j - w_j - h_j \le 2\tau q_j, \qquad \frac{\nu - w_j - h_j}{3\tau} \le q_j \le \frac{\nu - w_j - h_j}{2\tau}.$$

At the common wholesale price w^{LS}, the upper quantity bounds sum to one and therefore pin down the displayed quantities and retail prices uniquely.

Two active stores with different total marginal costs. For the post-entry analysis, write $k_j = w_j + h_j$ for store j's total marginal cost. When both stores make positive sales and the consumer between them strictly prefers purchasing,

$$q_j = \frac{\tau + p_\ell - p_j}{2\tau}.$$

Store j's first-order condition is

$$q_j - \frac{p_j - w_j - h_j}{2\tau} = 0.$$

Solving the two pricing conditions gives

$$p_j = \tau + \frac{2(w_j + h_j) + (w_\ell + h_\ell)}{3},$$

and

$$q_j = \frac{1}{2} + \frac{(w_\ell + h_\ell) - (w_j + h_j)}{6\tau}.$$

Thus

$$p_j - w_j - h_j = 2\tau q_j.$$

At these prices, the consumer indifferent between stores obtains utility

$$\nu - \frac{k_L + k_S + 3\tau}{2} \geq \frac{\tau - (h_S - h_L)}{2} \geq 0,$$

because both wholesale prices are at or below w^L and Assumption 3 holds. The prices maximize profit against the covered-market demand curve, including its zero- and unit-demand bounds. Allowing consumers to abstain can only reduce deviation profits at nonnegative margins, while pricing below cost cannot improve on positive profit. Thus the proposed prices are global best responses in the original retail game whenever both shares are positive.

When the two stores pay the same wholesale price w^L, these expressions become

$$q_L^{LS,E} = \frac{1}{2} + \frac{h_S - h_L}{6\tau}, \qquad q_S^{LS,E} = \frac{1}{2} - \frac{h_S - h_L}{6\tau},$$

and

$$p_L^{LS,E} = w^L + h_L + \tau + \frac{h_S - h_L}{3}, \qquad p_S^{LS,E} = w^L + h_L + \tau + \frac{2(h_S - h_L)}{3}.$$

Since $w^L = \nu - h_L - 2\tau$,

$$p_L^{LS,E} = \nu - \tau + \frac{h_S - h_L}{3}, \qquad p_S^{LS,E} = \nu - \tau + \frac{2(h_S - h_L)}{3}.$$

The consumer indifferent between the stores obtains utility

$$\frac{\tau - (h_S - h_L)}{2}.$$

Assumption 3 therefore guarantees that the full market is covered. For a strict inequality, all consumers strictly prefer purchasing and the two pricing conditions determine the unique equilibrium. At equality, the lower quantity bounds for a just-covered equilibrium sum to one, which pins down the same quantities and retail prices.

If the difference in total marginal costs satisfies $k_\ell - k_j \geq 3\tau$, an equilibrium has $p_j = k_\ell - \tau$, $p_\ell = k_\ell$, and quantities $q_j = 1$, $q_\ell = 0$. The active store's left-hand profit derivative is one, while its right-hand derivative is

$$1 - \frac{k_\ell - k_j - \tau}{2\tau} \leq 0.$$

The other store cannot attract buyers at a positive margin. Since $k_\ell \leq w^L + h_\ell < \nu$, even the most distant consumer is willing to buy. Thus the active store's quantity is the covered-market share capped at one. Retail prices at unused offers that leave a store inactive need not be unique, but the quantities used in the supplier's optimization are unchanged.

H.2 Supplier optimization before the wholesaler's entry

I next compare the supplier's optimal wholesale price when it serves one and two stores.

One active store. For

$$w \leq \nu - h_L - 2\tau,$$

L serves the entire market, so supplier profit is

$$\Pi^L(w) = w - c - F.$$

Profit is increasing in w throughout this region. If the supplier raises the wholesale price above the full-coverage boundary, L serves only part of the market and supplier profit becomes

$$\Pi^L(w) = (w - c)\frac{\nu - w - h_L}{2\tau} - F.$$

Its derivative evaluated at

$$w^L = \nu - h_L - 2\tau$$

is

$$\left.\frac{d\Pi^L(w)}{dw}\right|_{w=w^L} = \frac{c + h_L + 4\tau - \nu}{2\tau} \leq 0,$$

where the inequality follows from Assumption 2. The partial-coverage profit function is concave in w, so the supplier does not want to raise the wholesale price above w^L.

Hence

$$w^L = \nu - h_L - 2\tau, \qquad q_L^L = 1,$$

and

$$\Pi^L = w^L - c - F.$$

Because $h_L < h_S$, supplying only S cannot generate more profit than supplying only L.

Two active stores. When both stores are active and a gap of non-buyers remains between their customer regions, supplier variable profit is

$$(w - c)\frac{2\nu - 2w - h_L - h_S}{2\tau}.$$

Its derivative at the full-coverage boundary is

$$\left.\frac{d}{dw}\left[(w - c)\frac{2\nu - 2w - h_L - h_S}{2\tau}\right]\right|_{w=w^{LS}} = \frac{c + \frac{h_L + h_S}{2} + 2\tau - \nu}{\tau}.$$

Assumptions 1 and 2 make this expression strictly negative, and the derivative decreases further as w increases while both stores remain active. For $c \leq w < w^{LS}$, total sales cannot exceed one, so gross profit is at most $w - c < w^{LS} - c$. Wholesale prices below marginal cost cannot improve on profitable operation. Thus

$$w^{LS} = \nu - \tau - \frac{h_L + h_S}{2}$$

is optimal conditional on serving both stores.

Supplier profit is therefore

$$\Pi^{LS} = w^{LS} - c - 2F.$$

If a higher wholesale price leaves only one store ordering, the supplier pays only one delivery cost and its profit is bounded by the optimized one-store profit. The global distribution choice therefore reduces to the optimized one- and two-store outcomes and zero profit from making no deliveries.

Comparison. The profit difference between the two-store and one-store outcomes is

$$\Pi^{LS} - \Pi^{L} = w^{LS} - w^{L} - F = \tau - \frac{h_S - h_L}{2} - F.$$

Defining

$$F^{pre} = \tau - \frac{h_S - h_L}{2},$$

the supplier serves both stores when $F < F^{pre}$ and only L when $F > F^{pre}$, provided operating through L remains profitable.

The latter requires

$$F < w^{L} - c.$$

Assumption 2 implies

$$w^{L} - c = \nu - c - h_L - 2\tau \geq 2\tau > F^{pre},$$

so the exclusion region

$$F^{pre} < F < w^{L} - c$$

is nonempty. For $F < F^{pre}$, serving both stores is profitable, since $\Pi^{LS} = w^{LS} - c - 2F > (w^{L} - c) - F^{pre} \geq 2\tau - F^{pre} > 0$. For $F > w^{L} - c$, $\Pi^{L} < 0$ and, since $F > w^{L} - c \geq 2\tau > F^{pre}$, $\Pi^{LS} = \Pi^{L} + F^{pre} - F < \Pi^{L}$, so the supplier serves no store. This establishes Proposition 1.

Active exclusion. At the one-store equilibrium,

$$w^{L} + h_L = \nu - 2\tau, \qquad w^{L} + h_S = \nu - 2\tau + (h_S - h_L).$$

Suppose both stores were offered w^{L} and chose their retail prices optimally. Assumption 1 gives $w^{L} + h_S < \nu$. If S sold nothing, a price slightly below $\min\{\nu, p_L + \tau\}$ would attract nearby buyers at a positive margin whenever this minimum exceeded $w^{L} + h_S$. Zero sales at S therefore require $p_L + \tau \leq w^{L} + h_S$. In that case L serves everyone and its margin satisfies

$$p_L - w^{L} - h_L \leq h_S - h_L - \tau < \tau.$$

A negative margin cannot be optimal because L could avoid losses by making no sales. At a nonnegative margin, a small retail-price increase loses buyers at rate at most $1/\tau$, so L's right-hand profit derivative is at least

$$1 - \frac{p_L - w^{L} - h_L}{\tau} > 0.$$

This contradicts optimal pricing by L. Hence S has positive sales in equilibrium. Its margin is also positive, since a zero-margin price could be raised slightly while retaining positive demand. The supplier nevertheless excludes it when $F > F^{pre}$ because the additional gross wholesale profit is smaller than the delivery cost.

H.3 Consumer welfare before the wholesaler's entry

With only L active,

$$p_L^L = \nu - \tau,$$

and every consumer purchases from L. Consumer surplus is therefore

$$CS^L = \int_0^1 (\nu - p_L^L - \tau x)\, dx = \frac{\tau}{2}.$$

With both stores active at the supplier's optimal two-store wholesale price w^{LS}, the marginal consumer is indifferent to purchasing. Utility therefore forms a triangle on each side of the market, with height τq_j^{LS} and base q_j^{LS}. Hence

$$CS^{LS} = \frac{\tau}{2}\left[(q_L^{LS})^2 + (q_S^{LS})^2\right].$$

Substituting the equilibrium quantities gives

$$CS^{LS} = \frac{\tau}{4} + \frac{(h_S - h_L)^2}{16\tau}.$$

Therefore

$$CS^L - CS^{LS} = \frac{\tau}{4} - \frac{(h_S - h_L)^2}{16\tau} > 0,$$

where the inequality follows from Assumption 1.

Adding the second store reduces travel distances, but before the wholesaler's entry it also allows the supplier to raise the wholesale price from w^L to w^{LS}. The resulting retail-price increase dominates the travel savings in aggregate. The consumer at $x = 1$ obtains zero utility with only L active and $\tau q_S^{LS} > 0$ with both stores, so consumers close to S gain from the second store.

Store profits also fall when both stores are served: L earns τ under exclusion, since $p_L^L - w^L - h_L = \tau$, while the two stores together earn $\tau[(q_L^{LS})^2 + (q_S^{LS})^2] = \frac{\tau}{2} + \frac{(h_S - h_L)^2}{8\tau} < \tau$.

H.4 Equilibrium after the wholesaler's entry

Suppose the entrant offers

$$w_E = w^L$$

and stores break ties in favor of the incumbent. I now impose Assumption 3.

Both stores. If the incumbent supplies both stores at any $w \leq w^L$, both buy from it. The full market is covered, so total quantity is one and incumbent profit is

$$\Pi^{LS,E}(w) = w - c - 2F.$$

This is increasing in w. A wholesale price above w^L loses both stores to the entrant. Hence

$$w^{LS,E} = w^L, \qquad \Pi^{LS,E} = w^L - c - 2F.$$

The entrant receives no orders.

Only L. If the incumbent supplies only L at wholesale price $w \leq w^L$, S can buy from the entrant at w^L. Assumption 3 and the retail-pricing results in Appendix H.1 imply that the market remains covered for every relevant w. Wholesale prices below c yield a negative wholesale margin and cannot improve on profitable operation, while wholesale prices above w^L receive no orders. It therefore suffices to optimize over $c \leq w \leq w^L$.

As long as both stores make positive sales,

$$q_L(w) = \frac{1}{2} + \frac{w^L - w + (h_S - h_L)}{6\tau}.$$

If the incumbent lowers w enough that L captures the entire market, $q_L = 1$. Hence

$$q_L(w) = \min\left\{1, \frac{1}{2} + \frac{w^L - w + (h_S - h_L)}{6\tau}\right\}.$$

The boundary at which L captures the full market is

$$w = w^L + (h_S - h_L) - 3\tau.$$

For wholesale prices above this boundary, incumbent profit before the delivery cost is

$$(w - c)\left[\frac{1}{2} + \frac{w^L - w + (h_S - h_L)}{6\tau}\right].$$

Its derivative is

$$\frac{3\tau + w^L + c + (h_S - h_L) - 2w}{6\tau},$$

and its second derivative is

$$-\frac{1}{3\tau} < 0.$$

At $w = w^L$, the derivative is

$$\frac{3\tau + (h_S - h_L) - (w^L - c)}{6\tau}.$$

At the full-capture boundary, the right-hand derivative is

$$\frac{9\tau - (h_S - h_L) - (w^L - c)}{6\tau}.$$

Concavity therefore gives three cases for the optimal wholesale price when the incumbent supplies only L:

$$w^{L,E} = \begin{cases} w^L, & w^L - c \le 3\tau + (h_S - h_L), \\ c + \dfrac{3\tau + (w^L - c) + (h_S - h_L)}{2}, & 3\tau + (h_S - h_L) < w^L - c < 9\tau - (h_S - h_L), \\ w^L + (h_S - h_L) - 3\tau, & w^L - c \ge 9\tau - (h_S - h_L). \end{cases}$$

At this wholesale price, L sells

$$q_L^{L,E} = \begin{cases} \dfrac{1}{2} + \dfrac{h_S - h_L}{6\tau}, & w^L - c \le 3\tau + (h_S - h_L), \\ \dfrac{3\tau + (w^L - c) + (h_S - h_L)}{12\tau}, & 3\tau + (h_S - h_L) < w^L - c < 9\tau - (h_S - h_L), \\ 1, & w^L - c \ge 9\tau - (h_S - h_L). \end{cases}$$

Thus optimized incumbent profit when supplying only L is

$$\Pi^{L,E} = \begin{cases} (w^L - c)\dfrac{3\tau + (h_S - h_L)}{6\tau} - F, & w^L - c \le 3\tau + (h_S - h_L), \\ \dfrac{\left[3\tau + (w^L - c) + (h_S - h_L)\right]^2}{24\tau} - F, & 3\tau + (h_S - h_L) < w^L - c < 9\tau - (h_S - h_L), \\ w^L - c + (h_S - h_L) - 3\tau - F, & w^L - c \ge 9\tau - (h_S - h_L). \end{cases}$$

In each case the incumbent's gross profit, $\Pi^{L,E} + F$, is below its pre-entry level $w^L - c$: in the first, L sells less than the full market at the same margin; in the second and third, at most the full market at a lower margin.

Supplying only S cannot generate more profit. At any common incumbent wholesale price, replacing S with L lowers the supplied store's handling cost and weakly raises the quantity sold at the same wholesale margin and delivery cost.

Distribution choice. Subtracting the optimized L-only profit from the profit obtained by supplying both stores gives

$$\Pi^{LS,E} - \Pi^{L,E} = F^{post} - F,$$

where

$$F^{post} = \begin{cases} (w^L - c)\dfrac{3\tau - (h_S - h_L)}{6\tau}, & w^L - c \le 3\tau + (h_S - h_L), \\ (w^L - c) - \dfrac{\left[3\tau + (w^L - c) + (h_S - h_L)\right]^2}{24\tau}, & 3\tau + (h_S - h_L) < w^L - c, \\ & w^L - c < 9\tau - (h_S - h_L), \\ 3\tau - (h_S - h_L), & w^L - c \ge 9\tau - (h_S - h_L). \end{cases}$$

I next verify that

$$F^{post} > F^{pre}$$

throughout the maintained parameter region.

In the first branch, Assumption 2 implies $w^L - c \geq 2\tau$, so

$$\begin{aligned} F^{post} - F^{pre} &\geq 2\tau \frac{3\tau - (h_S - h_L)}{6\tau} - \left[\tau - \frac{h_S - h_L}{2}\right] \\ &= \frac{h_S - h_L}{6} > 0. \end{aligned}$$

In the second branch, F^{post} is increasing in $w^L - c$, since

$$\frac{\partial F^{post}}{\partial (w^L - c)} = \frac{9\tau - (w^L - c) - (h_S - h_L)}{12\tau} > 0.$$

It is continuous at the boundary with the first branch, so the same strict inequality continues to hold.

In the third branch,

$$F^{post} - F^{pre} = 2\tau - \frac{h_S - h_L}{2} > 0.$$

Hence

$$F^{post} > F^{pre}.$$

The interval

$$F^{pre} < F < F^{post}$$

is therefore nonempty. For any F in this interval, the incumbent excludes S before the wholesaler's entry but supplies both stores after it.

Supplying both stores is also profitable throughout this region.

The incumbent can always offer only L the wholesale price w^L. At this wholesale price, L sells

$$q_L(w^L) = \frac{1}{2} + \frac{h_S - h_L}{6\tau}.$$

Optimizing the one-store alternative can only raise gross profit, so the gain from supplying both stores satisfies

$$F^{post} \leq (w^L - c)\left[1 - q_L(w^L)\right] = (w^L - c)\left[\frac{1}{2} - \frac{h_S - h_L}{6\tau}\right] < \frac{w^L - c}{2}.$$

Hence $F < F^{post}$ implies

$$\Pi^{LS,E} = w^L - c - 2F > 0.$$

Thus exiting is dominated, and the entrant receives no orders in the realized inclusion equilibrium. For $F > F^{post}$, serving only L yields higher profit than serving both stores, and the incumbent exits once F exceeds its gross profit from serving L alone. This establishes Lemma 1 and Proposition 2.

H.5 Consumer welfare after the wholesaler's entry

Under Assumption 3, the post-entry retail equilibrium is unique and

$$q_L^{LS,E} = \frac{1}{2} + \frac{h_S - h_L}{6\tau}, \qquad q_S^{LS,E} = \frac{1}{2} - \frac{h_S - h_L}{6\tau}.$$

Retail prices are

$$p_L^{LS,E} = \nu - \tau + \frac{h_S - h_L}{3}, \qquad p_S^{LS,E} = \nu - \tau + \frac{2(h_S - h_L)}{3}.$$

Because the full market is covered, aggregate consumer surplus is

$$\begin{aligned} CS^{LS,E} &= \int_0^{q_L^{LS,E}} (\nu - p_L^{LS,E} - \tau x)\, dx \\ &\quad + \int_{q_L^{LS,E}}^1 (\nu - p_S^{LS,E} - \tau(1-x))\, dx \\ &= \nu - p_L^{LS,E} q_L^{LS,E} - p_S^{LS,E} q_S^{LS,E} - \frac{\tau}{2}\left[(q_L^{LS,E})^2 + (q_S^{LS,E})^2\right]. \end{aligned}$$

Before the wholesaler's entry, total retail payments equal $p_L^L = \nu - \tau$ and total travel costs equal $\tau/2$. Substituting the post-entry prices and quantities into these two expenditure components gives

$$\begin{aligned} p_L^{LS,E} q_L^{LS,E} + p_S^{LS,E} q_S^{LS,E} - p_L^L &= \frac{h_S - h_L}{2} - \frac{(h_S - h_L)^2}{18\tau}, \\ \frac{\tau}{2} - \frac{\tau}{2}\left[(q_L^{LS,E})^2 + (q_S^{LS,E})^2\right] &= \frac{\tau}{4} - \frac{(h_S - h_L)^2}{36\tau}. \end{aligned}$$

The first line is the increase in retail payments; the second is the saving in travel costs.

Substituting equilibrium prices and quantities gives

$$CS^{LS,E} = \frac{3\tau}{4} - \frac{h_S - h_L}{2} + \frac{(h_S - h_L)^2}{36\tau}.$$

Since

$$CS^L = \frac{\tau}{2},$$

the aggregate change in consumer surplus is

$$CS^{LS,E} - CS^L = \frac{\tau}{4} - \frac{h_S - h_L}{2} + \frac{(h_S - h_L)^2}{36\tau}.$$

Letting $z = (h_S - h_L)/\tau$, consumer surplus increases if and only if

$$z^2 - 18z + 9 > 0.$$

The two roots are

$$9 - 6\sqrt{2} \qquad \text{and} \qquad 9 + 6\sqrt{2}.$$

Assumptions 1 and 3 imply $0 < z \leq 1$, so the relevant condition is

$$h_S - h_L < (9 - 6\sqrt{2})\tau.$$

The gains are heterogeneous across consumers. Consumers who continue purchasing from L face the same travel distance as before the wholesaler's entry but pay a retail price higher by

$$p_L^{LS,E} - p_L^L = \frac{h_S - h_L}{3},$$

and therefore lose.

A consumer who switches to S gains

$$p_L^L - p_S^{LS,E} + \tau(2x - 1) = -\frac{2(h_S - h_L)}{3} + \tau(2x - 1).$$

Thus switchers sufficiently close to S gain from the wholesaler's entry. When

$$h_S - h_L < (9 - 6\sqrt{2})\tau,$$

these travel gains dominate the retail-price losses in aggregate. The same comparison applies when the incumbent keeps w^L while excluding S (case (i) of Lemma 1) or exits, since both stores then pay w^L.

Exclusion with a lower wholesale price. In cases (ii) and (iii) of Lemma 1, L's total marginal cost $k_L = w^{L,E} + h_L$ falls below $w^L + h_L$, while S's remains $k_S = w^L + h_S$. When both stores are active, the retail prices $p_j = \tau + \frac{2k_j + k_\ell}{3}$ derived in Appendix H.1 fall for both stores relative to the inclusion equilibrium, so every consumer's best option improves. When L captures the entire market, $p_L = k_S - \tau = \nu - 3\tau + (h_S - h_L)$, which is below $p_L^L = \nu - \tau$ by Assumption 1, and a consumer at x obtains $\nu - p_L - \tau x = 3\tau - (h_S - h_L) - \tau x$. This exceeds $\nu - p_L^{LS,E} - \tau x$ by $2\tau - \frac{2(h_S - h_L)}{3}$ and $\nu - p_S^{LS,E} - \tau(1 - x)$ by $3\tau - 2\tau x - \frac{h_S - h_L}{3}$, which is smallest at $x = 1$, where it equals $\tau - \frac{h_S - h_L}{3}$. Both differences are positive under Assumption 1. Consumer surplus in cases (ii) and (iii) therefore weakly exceeds $CS^{LS,E}$, and the condition above is sufficient for the wholesaler's entry to raise aggregate consumer surplus in every post-entry outcome.

I Complementary survey

I.1 Neighborhood selection

The complementary survey links consumers' shopping reports to the corner stores available in their neighborhoods. I fielded it in July and August 2026 in 30 control clusters randomly sampled from the experiment, stratifying by district, and randomly selected one residential block within each cluster. The survey yielded 576 consumer interviews and a listing of 518 corner stores, averaging 19 consumers and 17 stores per neighborhood. Eligible blocks had no more than 10% of the area formed by the central block, adjacent blocks, and surrounding streets outside the cluster.

I.2 Store listing and survey

The field team identified all corner stores in the sampled block and neighboring blocks, adding other nearby stores when needed to account for differences in block size or shape. Enumerators mapped these stores before interviewing consumers, recording their locations, names, and appearance.

They also attempted a survey of all listed stores, with repeated visits as needed to reach owners and complete interviews. 19 stores remained closed throughout fieldwork, preventing observation of floor area and variety. Among the remaining 499 stores, 18% declined the owner survey. Owners received a small gift worth approximately 2 USD PPP for completing the survey. Table I1 reports store characteristics and shopping patterns, both overall and separately for low- and high-variety stores.

Surveyed owners reported the number of brands in nine categories covering leading products across beverages, dairy, pantry staples, and household supplies, and enumerators verified these counts. Table I2 lists the categories and reports their availability and brand counts. The survey also records availability and retail prices for the same six popular products as the main experiment. I construct the same retail price index, averaging each available product's log price relative to its geometric mean in the survey, as described in Section 4.2.[49]

Before requesting consent, enumerators recorded a five-point observational variety score. This provides information about stores that declined the owner survey. The variety score runs from 1, "can only buy a few basic products," to 5, "can find any product I need." Its rank correlation with verified brand counts is high at 0.60 (Figure I1). I impute brand counts at the 92 stores without complete counts using the mean count among surveyed stores with the same score. For demand estimation, I then map these nine-category counts into predicted counts across all categories, as described in Appendix J.1. I also record floor area as a measure of store size, which includes both selling and storage space.[50]

[49] For each product, I exclude nonpositive prices and prices strictly outside its first and ninety-ninth percentiles before constructing the index.

[50] I correct one enumerator's discrete size scale by matching its pooled distribution to that of the other enumerators, assigning tied values the mean over the corresponding percentile interval. I apply the correction separately to floor and selling area, then impute missing floor area from selling area recorded before consent. The fitted relationship is floor area $= 0.839 + 1.063\times$ selling area, with $R^2 = 0.901$.

I.3 Consumer survey

After mapping the stores, enumerators began at a corner of the sampled block and attempted to survey every household on that block. They interviewed the person who answered the door if they were at least 18 years old, with no restrictions on gender or their role in household shopping. The target was 20 completed interviews per block. When necessary to meet their targets, enumerators continued across the street and then farther out into the neighborhood. 80% of surveyed households are in the sampled block or within 20 meters of it, and an additional 13% are within 50 meters.[51] Consumers received a small gift worth approximately 2 USD PPP for completing the survey.

Surveyed consumers reported the two neighborhood corner stores they had visited most often during the preceding week. For each, they reported the number of visits, average expenditure per visit, and products bought on their most recent visit. They also reported purchases at other corner stores, markets, and supermarkets, along with demographics and how often they handled household shopping. When interviewing consumers, enumerators had printed maps of the neighborhood stores and were familiar with their names, owners, and appearance. This allowed them to match shopping reports to actual stores, linking consumers' destinations to store characteristics. Figure I2 illustrates two surveyed neighborhoods. Table I3 reports respondents' demographics and shopping patterns.[52]

To calculate walking times, I construct routes from household and store GPS locations, checked against field photographs, and convert distance to round-trip minutes at 75 meters per minute. Consumers' self-reported walking times to their chosen stores are 1.74 times calculated times on average. I use this ratio to express mean trip duration, the distance–variety trade-off, and welfare equivalents in perceived walking minutes. This changes the units of time without changing predicted store choices or the proportional welfare gain. Figure I3 compares calculated and reported times for named stores. Among consumers reporting different times to their two stores, the reported and calculated measures identify the same closer store in 82% of cases.

I.4 Hypothetical store choices

I also implemented a hypothetical store-choice exercise to assess consumers' preferences for variety while holding other store characteristics fixed. Respondents viewed images of a low- and a high-variety store and allocated their weekly visits between them. They were instructed to assume that retail prices, their relationship with the owner, and all other aspects of store quality were the same. The images were generated through AI prompting using real store photographs and the survey's five-point variety scale. Respondents' average ratings of 2.5 and 4.7 closely match the intended difference in variety.

These questions were added after the first week of fieldwork, when the shopping survey had already been completed in some clusters. To obtain responses in those clusters, I sampled

[51] The remaining households are 72 meters away on average. Excluding consumers more than 50 meters away does not affect the results.

[52] Demand estimation uses 563 respondents. The remaining 13 report no visits to mapped corner stores, have incomplete visit counts, or visited only stores without observed variety and floor area.

new blocks and sent enumerators back to administer the hypothetical-choice questions. These respondents did not complete the shopping survey because the stores in their blocks had not been mapped. The exercise includes 395 respondents from the shopping survey and 146 from these additional visits. Results are similar when the latter are excluded.

At equal distance, respondents allocate 75% of visits to the high-variety store, corroborating the strong preference for variety in observed shopping choices. In the distance-varying task, round-trip walking takes ten minutes to the high-variety store and two minutes to the low-variety store. The high-variety share falls to 66%.

The ten-minute round trip to the high-variety store is longer than almost all consumers' observed shopping trips. Respondents appear not to have internalized how long this walk would be relative to their usual trips. Their answers also suggest limited attention to the distance change: 27% maintained their previous allocation, and 23% increased the share of visits to the now more distant store. These questions were asked in passing during the survey, without incentives tied to respondents' choices. I therefore use the equal-distance task to corroborate preferences for variety, but rely on observed shopping choices to quantify the distance–variety trade-off.

Table I1. Store characteristics and patronage by variety

	All stores (1)	Low variety (2)	High variety (3)	p-value (4)
Panel A: Store characteristics				
Floor area (m^2)	21.5	17.6	25.7	< 0.001
	(1.5)	(1.1)	(1.8)	
Nine-category brand count	28.01	20.18	36.25	< 0.001
	(0.61)	(0.40)	(0.50)	
Categories carried	8.33	7.82	8.83	< 0.001
	(0.09)	(0.16)	(0.03)	
Retail price index	1.003	1.006	1.000	0.227
	(0.003)	(0.004)	(0.003)	
Store consented to survey	0.816	0.793	0.840	0.339
	(0.017)	(0.022)	(0.035)	
Stores	499	256	243	
Panel B: Shopping patterns				
Weekly visits per store	7.44	5.00	10.01	< 0.001
	(0.36)	(0.44)	(0.67)	
Weekly expenditure per store (USD PPP)	59.10	38.10	81.22	< 0.001
	(3.49)	(3.96)	(6.38)	
Received a visit	0.543	0.480	0.609	0.012
	(0.029)	(0.038)	(0.038)	
Expenditure per visit (USD PPP)	7.95	7.64	8.12	0.261
	(0.26)	(0.31)	(0.34)	
Distinct products in most recent purchase	1.94	1.99	1.91	0.446
	(0.09)	(0.14)	(0.08)	
Visit reports	968	374	594	

Notes: The unit of observation is the store, except for expenditure and products per visit, which use consumer–store purchase reports. Brand counts sum distinct brands within the nine survey categories. Low- and high-variety stores have counts at or below and above the demand-sample median, respectively. Floor area and brand counts include imputations. Categories carried exclude stores that did not consent to the survey. The retail price index is exponentiated before averaging across stores with observed prices. Weekly visits and expenditure sum sampled consumers' reports at each store, including zero visits. Expenditure per visit is weighted by visits with valid expenditure reports. Visit reports count all named-store reports, including those with missing expenditure or product information. Column (4) tests equality between low- and high-variety stores. Monetary values are in USD PPP. Standard errors, in parentheses, are clustered at the neighborhood level.

Table I2. Availability and brand counts in the nine survey categories

	Store share	Mean brands, if carried	SD
	(1)	(2)	(3)
Water	0.995	3.17	[1.37]
Beer	0.771	2.49	[1.41]
Canned fish	0.919	3.20	[1.77]
Soft drinks	0.998	5.33	[2.22]
Milk	0.919	2.71	[1.31]
Toilet paper	0.975	3.25	[1.30]
Detergent	0.936	4.08	[1.85]
Instant coffee	0.931	2.87	[1.21]
Cooking oil	0.885	2.88	[1.47]

Notes: The unit of observation is the store. The sample includes 407 stores with complete owner-reported, enumerator-verified brand counts. Column (1) is the share of stores with at least one brand in the category. Columns (2) and (3) report the mean and standard deviation (in brackets) of brand counts among stores carrying the category.

Table I3. Consumer characteristics and shopping patterns

	Mean (1)	SD (2)
Panel A: Consumer characteristics		
Age (years)	48.9	[16.4]
Female	0.648	[0.478]
Handles most household shopping	0.684	[0.465]
Panel B: Weekly shopping patterns		
Corner-store visits per week	7.4	[4.0]
Corner-store expenditure (USD PPP)	57.2	[51.2]
Market visits per week	3.1	[2.5]
Market expenditure (USD PPP)	92.3	[109.6]
Supermarket visits per week	1.0	[1.2]
Supermarket expenditure (USD PPP)	43.8	[108.8]

Notes: The unit of observation is the consumer. The survey includes 576 consumers, and each statistic uses available responses. Columns (1) and (2) report the mean and standard deviation (in brackets) of each characteristic. Handles most household shopping equals one for respondents who attend to household shopping more than half the time. Corner-store visits use the 563 consumers in the demand sample and include other corner stores. Corner-store expenditure combines spending at named stores and other corner stores, with the latter valued at the respondent's mean expenditure per named-store visit where needed. Market and supermarket expenditure are reported visits times average expenditure per visit. Monetary values are in USD PPP.

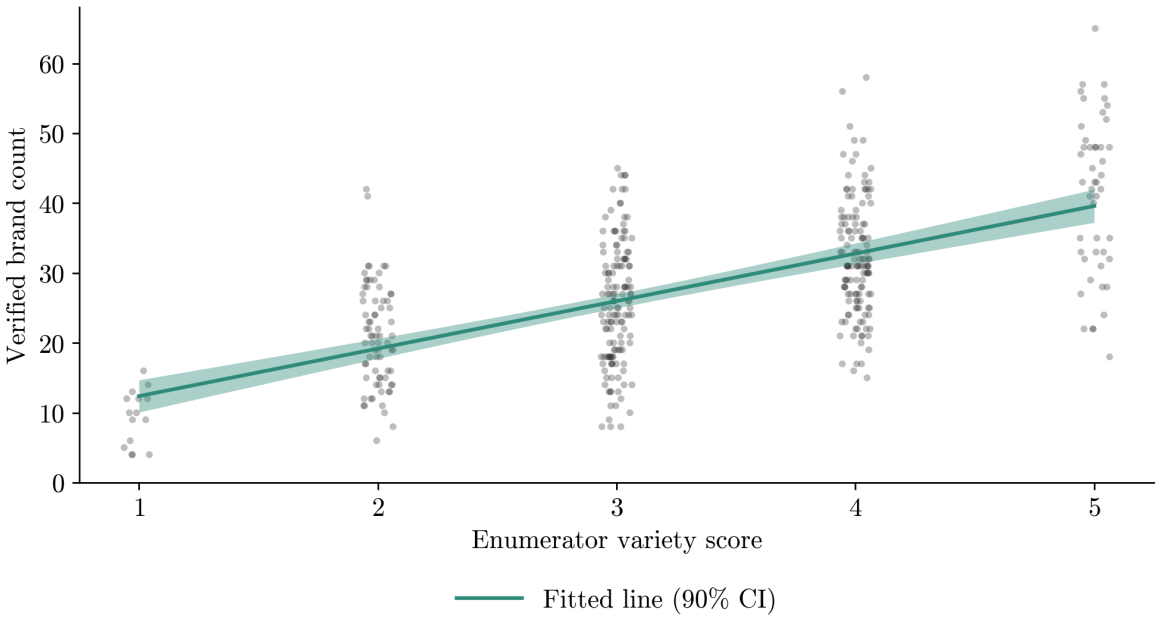


Figure I1. Enumerators' variety scores and verified brand counts

Notes: This figure plots owner-reported, enumerator-verified brand counts across the nine survey categories against the observational score recorded before requesting consent, for stores with complete verified brand counts in the demand sample. Horizontal jitter separates overlapping points. The fitted OLS line and its 90% confidence band use standard errors clustered by neighborhood. The rank correlation is 0.60, and the regression R^2 is 0.39. Brand imputation uses means within each score, rather than this fitted line.



Figure I2. Households and corner stores in two surveyed neighborhoods

Notes: This figure shows surveyed households and listed stores in the two neighborhoods, with variety classified by brand count as in Figure 5. Filled green circles denote high-variety stores, open black circles low-variety stores, and triangles households. Sampled blocks are shaded gray. Both panels use the same geographic scale and are centered on the sampled block, with stores outside the frame omitted. Census blocks are colored by socioeconomic status from INEI: poor and very poor (red), lower-middle-class (yellow), and upper-middle-class (blue). Roads are from OpenStreetMap.

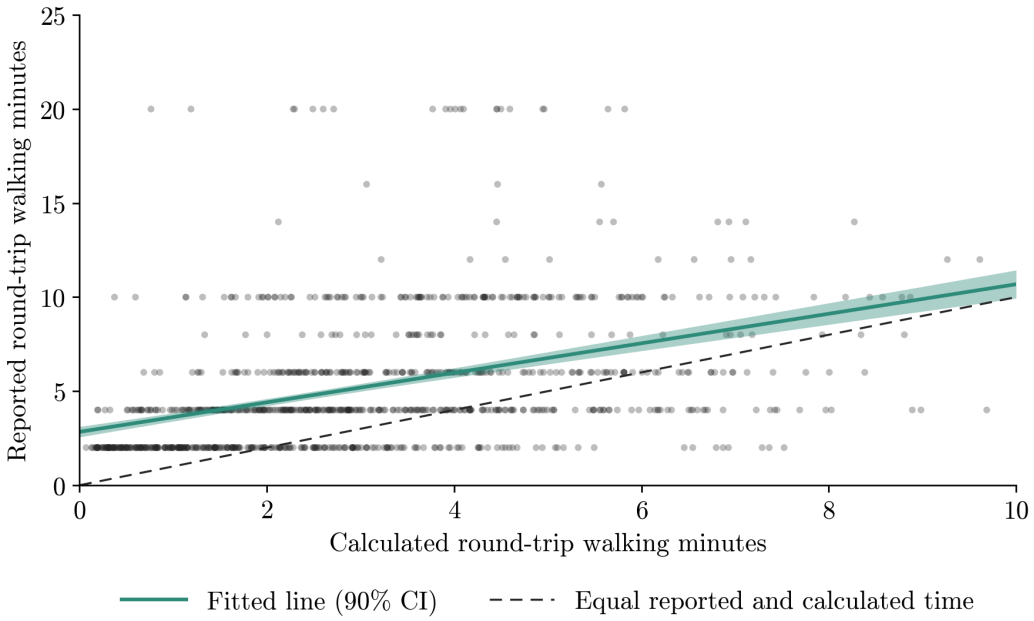


Figure I3. Calculated and reported walking times

Notes: This figure compares calculated and reported walking times over 978 named-store reports from 566 consumers with positive visits and valid calculated and reported walking times. Both axes report round-trip minutes. Calculated times use GPS routes and a walking speed of 75 meters per minute. The fitted OLS line and its 90% confidence band use standard errors clustered by consumer. The dashed line marks equality. Points above ten calculated or 25 reported minutes are omitted from the plot. The fitted line and statistics use the full sample. The rank correlation is 0.53, and the regression R^2 is 0.17.

J Demand estimation and welfare quantification

This appendix explains how I use consumers' weekly visits to neighborhood corner stores to estimate their preferences for proximity and product variety. I combine these preferences with the experimental effects of the wholesaler's entry to quantify consumer welfare gains. I first describe demand estimation, then explain the variety counterfactual and welfare calculation.

J.1 Estimating informed store choice

I estimate consumers' preferences for proximity and product variety from their weekly visits to neighborhood corner stores. The complementary survey records visits to the two stores each consumer visited most often and links these destinations to a mapped set of nearby alternatives. For the mapped stores, I observe or impute brand counts across nine leading product categories and floor area. I also record retail prices where available and calculate round-trip walking times from each household. These data allow me to estimate how consumers trade off product variety against the cost of reaching a store and its other characteristics. I take the number of trips as given and model the allocation of visits across stores, without an outside option. Table 8 reports the main specification and the versions with store random and fixed effects. Section 8.2.2 discusses their results.

To put demand estimation and the counterfactual on the same category coverage, I map the nine-category counts into predicted counts across all categories. Brands are counted separately within each category in both surveys. Among experimental controls with positive counts in the selected categories, I estimate

$$\log N_j = \delta_0 + \delta_1 \log n_j^{(9)} + u_j,$$

where N_j counts category × brand pairs across all categories and $n_j^{(9)}$ counts them within the nine survey categories. The regression uses 372 controls and gives an intercept of 0.862 (SE: 0.078) and a slope of 1.011 (SE: 0.034), with standard errors clustered by stratum. For each complementary-survey store, I set $n_j = \exp(\hat{\delta}_0 + \hat{\delta}_1 \log n_j^{(9)})$, using its observed or imputed nine-category count. These predictions exponentiate fitted log counts, without a residual-smearing adjustment. The slope is close to one, so the mapping largely preserves proportional differences in variety across stores.[53]

Variety attracts consumers by improving the match between their tastes and the products available at a store. On each shopping occasion, consumers know their product tastes before choosing a destination and buy one unit of the product they like best among those the chosen store carries. Store j carries n_j products, which I proxy with its predicted all-category brand count. I assume a common ordering of availability, so a store carrying n_j products offers products $1, \ldots, n_j$. Stores carrying more products therefore carry all the products available at stores carrying fewer products. Product tastes η_{ik} are independent Type-I extreme-value draws with scale $\sigma_\eta > 0$, redrawn on each occasion. Consumer i's best available product

[53] The mapping transfers a relationship estimated from wholesale purchases to counts of brands on shelves. It also extrapolates beyond part of the estimation sample: 365 of the 499 complementary-survey stores exceed the experimental controls' ninety-fifth percentile of nine-category counts, and 36 exceed their maximum.

at store j gives utility $M_i(n_j) = \max_{k \le n_j} \eta_{ik}$. The same product has the same taste shock wherever it is carried.

Consumers weigh this product match against travel costs and their general perception of each store's retail price level. Idiosyncratic tastes for stores also allow consumers to choose different destinations even when they face similar distances and product variety. Utility from store j is $U_{ij} = v_{ij}(\eta) + \epsilon_{ij}$, where

$$v_{ij}(\eta) = M_i(n_j) - \tau d_{ij} - \alpha \tilde{p}_j + X_j' \gamma.$$

Here d_{ij} is round-trip walking time, τ is the cost per minute, and $\tilde{p}_j$ is the retail price index, with price sensitivity α.[54] The vector X_j contains floor area in levels and an indicator for consent to the store survey, with coefficients γ. Store tastes ϵ_{ij} are independent Type-I extreme-value draws, independent of product tastes and redrawn on each shopping occasion. I normalize their scale to one without loss of generality, fixing the units of utility.

Under this taste distribution, expected product-match utility increases with log variety. Independence across products implies

$$\Pr\{M_i(n_j) \le m\} = \exp\left(-n_j e^{-m/\sigma_\eta}\right), \qquad \mathbb{E}[M_i(n_j)] = \sigma_\eta(\log n_j + \gamma_E),$$

where γ_E is Euler's constant. Greater dispersion in product tastes therefore makes additional variety more valuable. The coefficient on log variety in expected match utility is σ_η. Since I proxy for variety with brand counts, I report this parameter as the coefficient on log brand counts. Estimation uses the full distribution of realized product matches as described below.

Consumers observe their product tastes before choosing, but I do not observe these tastes in the survey. I therefore integrate over their possible realizations to calculate store-choice probabilities. For each simulated draw, I generate the best product matches jointly across all stores in the neighborhood, preserving the common tastes for products carried by multiple stores. Conditional on these matches, the Type-I extreme-value store tastes give logit choice probabilities. Averaging these probabilities across draws gives

$$\widehat{P}_{ij} = \frac{1}{R} \sum_{r=1}^{R} \frac{\exp\left[v_{ij}(\eta^{(r)})\right]}{\sum_{\ell \in \mathcal{J}_i} \exp[v_{i\ell}(\eta^{(r)})]}, \tag{J.1}$$

where $\mathcal{J}_i$ contains the mapped stores available to consumer i, and R is the number of simulation draws. Predicted brand counts determine which products each store carries and hence the joint distribution of the simulated matches.[55]

[54] The retail price index averages log relative prices for the same six popular products used in the main experiment and enters utility directly. I set missing indices to zero and include the consent indicator to retain stores that declined the survey.

[55] To generate matches efficiently, I order distinct brand counts and draw the best match among the additional products introduced at each count. The maximum across these successive groups gives the match at each store. For a group of size Δn, its maximum divided by σ_η is distributed as $\log \Delta n + g$, where g is a unit-scale Type-I extreme-value draw. This construction also accommodates fractional imputed counts. I approximate the product-taste integral with 2,048 scrambled Sobol draws, held fixed throughout estimation. A local numerical check with 4,096 draws implies parameter changes below 0.031 standard errors.

I choose the preference parameters that make the observed weekly visit counts most likely, using simulated maximum likelihood (Train, 2009). Let y_{ij} be consumer i's reported visits to store j, and let $T_i = \sum_j y_{ij}$ be total weekly visits to the named stores. All mapped stores retained in the demand sample enter as alternatives, including those with zero reported visits. Conditional on T_i, the visit counts have a multinomial likelihood. Omitting terms that do not depend on the parameters, I maximize

$$\widehat{\ell} = \sum_i \sum_{j \in \mathcal{J}_i} y_{ij} \log \widehat{P}_{ij}. \tag{J.2}$$

Because product tastes change across shopping occasions, I average the choice probabilities over tastes before forming the likelihood for weekly counts. Standard errors are clustered by neighborhood and incorporate uncertainty in the estimated mapping, as described in Appendix J.3.

For robustness, I allow stores to differ in persistent appeal beyond their observed characteristics. I add a store random effect ω_j to utility, shared by all consumers and fixed across shopping occasions. This differs from ϵ_{ij}, which captures a consumer's taste for a store on a particular occasion. I assume ω_j follows a Normal distribution with mean zero and variance σ_ω^2. These effects are independent across stores and independent of observed store characteristics and product tastes. The parameter σ_ω measures the dispersion of persistent store appeal.

Estimating this specification requires integrating the shared store effects over all consumers' choices within a neighborhood. For a given vector of effects ω, I first calculate $P_{ij}(\omega)$ by averaging over product tastes as above. I then form the joint likelihood of the neighborhood's visit counts and integrate over the Normal distribution of ω:

$$L_g = \int \prod_{i \in g} \prod_{j \in \mathcal{J}_i} P_{ij}(\omega)^{y_{ij}} \, dF_g(\omega), \tag{J.3}$$

where F_g is the joint distribution of store effects in neighborhood g. Integrating after combining consumers' choices preserves the fact that they share the same store effects. I maximize the sum of neighborhood log likelihoods and again cluster standard errors by neighborhood, incorporating mapping uncertainty.[56]

As a check on the distance coefficient, I also estimate a destination logit with store fixed effects. This specification exploits differences across consumers in their distance to the same stores, holding all store attributes fixed. Store fixed effects absorb variety and other store attributes, so this specification cannot identify the variety coefficient needed for the welfare calculation.

J.2 Constructing the variety counterfactual

I simulate the experiment in the complementary survey's store sample by applying the estimated treatment effect on log category × brand counts across all categories and its interaction

[56] I evaluate the store-effect integral with a Laplace approximation around the conditional mode. The random-effects specification uses 512 product-taste draws, held fixed throughout estimation.

with floor area. The experimental regression includes treatment, floor area, their interaction, and stratum fixed effects (following the specification with continuous floor area reported in the bottom rows of the main tables). For a store with baseline brand count n_j and floor area S_j, I construct

$$n_j^{\text{cf}} = n_j \exp(a + bS_j).$$

Here, a and b are the coefficients on treatment and its interaction with floor area in column (1), Panel B, of Table J1.[57] Locations, retail prices, floor area, and consumer trip totals remain fixed, isolating the welfare effect of expanded variety.

Applying these estimates requires connecting the surveys' measures of product variety and accounting for remaining differences in their store-size distributions.

First, the mapping aligns category coverage, but the diary records three-week brand purchases while the complementary survey records brands carried on shelves. Suppliers typically replenish stores weekly, so the diary covers several replenishment cycles. Restricting to the same nine categories, an additional square meter of floor area is associated with 1% more brands purchased among experimental controls and 1.6% more brands carried by complementary-survey stores (Table J2). These similar gradients support using purchases to proxy brands carried, but I transfer proportional rather than absolute changes in brand counts using treatment effects estimated in logs.

Second, the mapping uses the selected categories to predict broader product variety. The main experiment's purchase records allow me to assess how closely these measures track one another. Among experimental controls, brand counts for the nine categories in the complementary survey have a rank correlation of 0.90 with total brand counts and 0.90 with total product counts, the main variety outcome in Table 2. Their relationship with floor area is also similar: an additional square meter is associated with 1% more nine-category brands, 1.3% more brands overall, and 1.4% more products (Table J2). These comparisons support using the selected categories to capture broad differences in product variety. To assess sensitivity to the mapping itself, I also estimate demand directly on nine-category counts while retaining the all-category experimental response (the version labeled Without mapping in Table J3).

Third, I use treatment effects on brands across all categories to capture the broader expansion in product variety and retain the full experimental sample. Restricting categories discards information on purchases outside the selected categories and excludes 53 of the 996 stores. The nine-category estimates are smaller, though not significantly so. Neither the average effect nor its size interaction is significant at the 5% level, whereas both are significant using all brands in the full sample (Table J1). Comparing the measures on the same restricted sample shows that sample exclusion accounts for much of the change in the estimates. The differences between the estimates are not significant at the 5% level. I also report a version that uses nine-category counts in both demand estimation and the experimental response, dispensing with the mapping and holding category coverage fixed across the two steps (Table J3).

Finally, I assess sensitivity to differences in store size across surveys. The complementary

[57] The preferred fitted response is negative for 24 stores with floor area above 44.9 m^2. I apply the estimated responses without truncation.

survey lists all corner stores around each sampled block, whereas the main experiment's sampling criteria excluded very small stores (Appendix D.2). Nevertheless, the surveys' floor-area distributions overlap substantially. The preferred counterfactual transfers the experimental response at a given floor area. I also report a version that transfers the response by each store's size percentile within its survey sample, to assess sensitivity to the remaining differences (Table J3).

J.3 Welfare calculation and inference

I quantify consumer welfare by comparing expected maximum utility under the baseline and counterfactual product offerings. Conditional on product tastes, the extreme-value distribution of store tastes gives expected maximum utility equal to the log sum of exponentiated store utilities, up to a constant. This constant cancels between the two cases. Averaging the difference over product tastes gives the per-trip gain

$$\begin{aligned} G_i &= \mathbb{E}_{\eta,\epsilon}\left[\max_{j\in\mathcal{J}_i} U_{ij}^{\text{cf}} - \max_{j\in\mathcal{J}_i} U_{ij}\right] \\ &= \mathbb{E}_{\eta}\left[\log\sum_{j\in\mathcal{J}_i}\exp\big(v_{ij}^{\text{cf}}(\eta)\big) - \log\sum_{j\in\mathcal{J}_i}\exp(v_{ij}(\eta))\right]. \end{aligned} \tag{J.4}$$

Counterfactual utility uses the product offerings constructed in Appendix J.2. In the specification with store random effects, I also average over their estimated population distribution, holding each effect fixed between the baseline and counterfactual.

I express welfare gains as travel-cost equivalents. Multiplying G_i by the consumer's weekly visits to the named stores, T_i, gives the weekly utility gain, and dividing by the travel-cost parameter τ gives its equivalent in round-trip walking minutes. Predicted baseline walking is $D_i = T_i \sum_{j\in\mathcal{J}_i} P_{ij} d_{ij}$, so baseline travel costs are τD_i. I report the mean weekly equivalent, $\mathbb{E}_i[T_i G_i/\tau]$, and express the utility gain as a percentage of mean predicted baseline travel costs:

$$100\,\frac{\mathbb{E}_i[T_i G_i/\tau]}{\mathbb{E}_i[D_i]}.$$

This equivalent values improved access to products in walking units and need not equal the change in actual walking when consumers switch stores.[58]

I also express the welfare gain as a uniform proportional reduction in retail prices. Holding the product-price benchmarks used to construct the index fixed, a reduction c changes every store's index by $\log(1-c)$. It therefore raises utility on each trip by $-\alpha\log(1-c)$. Equating the resulting average weekly gain to the gain from expanded variety gives

$$c = 1 - \exp\left(-\frac{\mathbb{E}_i[T_i G_i]}{\alpha\mathbb{E}_i[T_i]}\right).$$

The confidence intervals account for uncertainty in estimated consumer preferences, the mapping between variety measures, and the experiment's effect on variety. I resample experimental strata and re-estimate the mapping and treatment responses jointly, preserving their

[58] Using consumers' perceptions of walking time, the main estimate of 1.9 calculated minutes per week corresponds to 3.3 perceived minutes. Appendix I.3 describes the conversion.

covariance. For the complementary survey, shared Normal perturbations of neighborhood scores preserve covariance across the demand specifications. I draw these perturbations independently of the experimental resamples. For each draw, I update demand parameters using their local response to the mapping, reconstruct baseline and counterfactual brand counts, and recompute welfare.[59] The fifth and ninety-fifth percentiles of the resulting distributions give the 90% confidence intervals. Each alternative counterfactual uses the corresponding experimental response and demand specification within the same draws.[60] The joint Wald tests in Table J3 use the covariance across welfare estimates within these common draws. Choice sets, trip totals, floor area, and imputed nine-category counts remain fixed.

The welfare gain is robust to allowing persistent differences in store appeal, omitting the mapping, transferring the experimental response by relative size, and using nine categories throughout; the last yields a lower point estimate, but a joint test does not reject equality across specifications (Table J3). The main specification gives a gain equivalent to 10% of baseline weekly travel costs, compared with 8.7% with store random effects and 8.7% under the relative-size counterfactual. Omitting the mapping gives 10.2%. Using nine-category counts throughout gives 5.9%, about 41% below the preferred estimate. This is equivalent to 1.1 round-trip walking minutes per week or a 3% reduction in retail prices. The 90% confidence intervals for welfare gains relative to baseline travel costs exclude zero in all specifications except the version using nine categories throughout (its CI is [-0.1%, 14.4%]). A joint test does not reject equality across specifications.

[59]The local response is obtained by differentiating the demand first-order conditions with respect to the mapping coefficients and checked against re-estimation. Positive taste, travel-cost, and store-effect scale parameters use the log parameterization from estimation, and all retail-price coefficient draws are retained.

[60]Each specification uses 1,999 draws. Numerical integration uses 512 product-taste draws for the main specification, 2,048 for the nine-category specification, and 512 joint product and store draws for random effects, with selected draws checked at higher resolution.

Table J1. Experimental effects on alternative brand measures

	All brands, full sample (1)	All brands, matched sample (2)	Nine categories, matched sample (3)
Panel A: Average treatment effect			
Treatment	0.142	0.121	0.081
	(0.060)	(0.062)	(0.054)
p-value	0.020	0.052	0.138
Panel B: Treatment response by floor area			
Treatment	0.317	0.254	0.203
	(0.091)	(0.096)	(0.087)
Treatment × floor area (m^2)	-0.0071	-0.0053	-0.0048
	(0.0028)	(0.0030)	(0.0027)
p-value for size (cont.) interaction	0.014	0.076	0.077
Stores	996	943	943
Panel C: Equality with column (1)			
p-value: average treatment effect	—	0.387	0.115
p-value: size (cont.) interaction	—	0.079	0.194

Notes: The unit of observation is the experimental store. The sample includes stores with positive receipt-verified supplier purchases and complete variety and floor-area measures. The outcome is the log count of category × brand pairs purchased over three weeks, including free goods. Columns (1) and (2) use all categories. Column (3) uses the nine survey categories. Columns (2) and (3) restrict to stores with positive counts in those nine categories. Singleton strata are excluded. Panel A includes treatment and stratum fixed effects. Panel B adds floor area and its treatment interaction. Standard errors, in parentheses, are clustered at the stratum level, following de Chaisemartin and Ramirez-Cuellar (2024). Panel C tests equality of each column's average effect and size interaction with column (1), using 9,999 common stratum-bootstrap resamples.

Table J2. Floor area and variety across the two surveys

	Log count per m^2 (1)	Percent per m^2 (2)	Stores (3)
Panel A: Experimental controls			
All products	0.0140	1.41	394
	(0.0032)		
All brands	0.0130	1.30	394
	(0.0028)		
Nine-category brands	0.0099	0.99	372
	(0.0025)		
Panel B: Complementary survey			
Nine-category brands	0.0154	1.55	403
	(0.0020)		

Notes: The unit of observation is the store. Each row regresses the log variety count on floor area in square meters. Panel A uses experimental controls with positive counts and complete variety and floor-area measures, from receipt-verified supplier purchases over three weeks. Products are distinct brand–size–flavor combinations. Brand counts sum distinct brands within categories, using all categories or the nine survey categories as indicated. Panel B uses complementary-survey stores with complete verified shelf counts and observed floor area. The nine categories are listed in Table I2, with instant coffee excluding ground coffee, coffee substitutes, and infusions. Column (2) reports $100[\exp(\text{coefficient}) - 1]$. Standard errors, in parentheses, are clustered at the stratum level in Panel A and at the neighborhood level in Panel B.

Table J3. Consumer welfare across demand specifications and counterfactuals

	Equivalent walking minutes (1)	Baseline travel costs (%) (2)	Retail price reduction (%) (3)
Panel A: Demand and mapping			
Main specification	1.9	10.0	5.0
	[0.6, 3.9]	[3.1, 20.5]	[0.7, 19.8]
Store random effects	1.8	8.7	4.8
	[0.6, 3.5]	[2.8, 16.8]	[-16.0, 23.2]
Without mapping	1.9	10.2	5.1
	[0.6, 4.0]	[3.1, 21.0]	[0.7, 20.7]
Panel B: Alternative variety counterfactuals			
Size percentiles within each survey	1.7	8.7	4.4
	[0.4, 3.7]	[2.1, 19.2]	[0.2, 18.4]
Nine categories throughout	1.1	5.9	3.0
	[0.0, 2.8]	[-0.1, 14.4]	[-0.7, 14.1]
Joint equality: p-value	0.519	0.530	—
Consumers	563	563	563

Notes: The unit of observation is the consumer in the demand sample. Column (1) reports equivalent round-trip walking minutes per week, using visits to named stores. Column (2) expresses the mean weekly utility gain as a percentage of mean predicted baseline travel costs under the corresponding demand model. Column (3) reports the uniform proportional retail-price reduction producing the same average weekly utility gain, holding the price-index benchmarks fixed. Brand counts sum distinct brands within categories. Panel A applies the all-category experimental response by floor area in levels. Its first two rows estimate demand on all-category counts predicted from nine-category counts using the log-log mapping. The third estimates demand directly on nine-category counts. In Panel B, the first row uses the main demand specification and applies the response by experimental-control floor-area percentile to floor-area percentiles within the complementary survey. The second uses nine-category counts in both demand estimation and the experimental response, with floor area in levels. Locations, retail prices, floor area, trip totals, and the consumer and store samples are held fixed. Store random effects are integrated over their estimated population distribution. Brackets contain 90% confidence intervals incorporating demand, mapping, and experimental-response uncertainty. The row labeled Joint equality reports joint Wald tests of equality across the five estimates, accounting for their shared data.